Composite Density Displays

Steven Lynch

ABSTRACT

Over the past ten years the science of seismic visualization has progressed on all fronts except one, the display of the actual seismic data itself. We can display and animate huge 3D volumes in real time but the seismic itself is still displayed as some combination of colour/greyscale variable density with a wiggle-trace overlay.

In the first part of this paper I introduce a new type of seismic display called a Composite Density display. These new displays are the geophysical equivalent of the computer game designer's bump-mapped images. They are constructed by combining a typical variable density display with what is essentially a shaded relief map of the seismic. The resultant display is, as I will show, very similar to the truly three-dimensional SeisScape® display but are much less resource intensive.

In the second part of the paper I tackle the question of "What makes a better display?" It won't take the reader long to realize that composite density displays are different but are they better? This is very hard to quantify since seismic is such a visual science and so the readers will have to make up their own minds. But to help, I present three sets of examples that cover a wider range of seismic situations.

The first example set, shows a seismic section that has been badly corrupted by migration artifacts. The variable density section shows that there are problems with the section, but the migration artifacts are very hard to see. By contrast, when the composite density's light source is oriented a certain way, the artifacts pop out of the section and are obvious to even the most unsophisticated viewer.

In the second set of examples, I tackle the nebulous subject of perception. By that I mean, "Can we use the composite density display to better understand what we can already see". The examples that were chosen, an in-line from the Stratton survey and a time slice from a Peruvian 3D are both free of processing artifacts. In both cases they were chosen because what is visible on the composite density display is also visible on the variable density display. By choosing these examples, I want the user to decide for themselves if the addition of lighting to the composite density display makes it easier to understand.

In the final set of examples, I compare the variable density and composite density views of two highly faulted seismic lines. In both cases, orienting the composite density's light source appropriately greatly enhances the viewer's ability to discern the faults. The first example has the light source perpendicular to the fault whereas the second has it parallel. This shows how important it is that the user be able to manipulate the light source in real time.

INTRODUCTION

Anyone who has been in geophysics for a while will readily acknowledge the tremendous impact that the past decade's explosion in computer power has had on the industry. I started my career in the late 70s, before the advent of personal computers and workstations. Now, every time I walk into a geophysicist's office I am amazed at the range of technology that they have on their desks. Foremost among these technologies is the ability to display and visualize entire seismic projects, in real time. Some of the software that is available, either on the PC or workstation platform, is simply amazing. We never thought it would get like this in the 70s, but then again, we never imagined that computers would progress the way they have.

Strangely though, we have moved forward on all fronts except one. We can now display huge seismic projects in ways we never imagined in the 70's but when it comes to displaying the seismic itself we display it exactly the same way. The typical modern computer-generated seismic display is some combination of colour/greyscale variable density with a wiggle-trace overlay. In fact, the last real advancement in seismic display took place in the late 1970s (Taner et al., 1979) when variable density displays went from greyscale to colour.

There are a number of reasons for this. Firstly, variable density displays are actually quite good. They are compact and are ideally suited to the limited display area of a computer monitor. They are also quick to produce and quick to display. Since we are not particularly adept at seeing colours, however, they are not very good at highlighting subtle amplitude anomalies or at showing character changes. To get around this we often overlay them with a wiggle-trace. But this approach is limited to small areas of the section because wiggle-traces require more surface area to be effective and so they are not well suited to computer monitors.

A second reason for our lack of progress is that we have the wrong mind-set about seismic data. We record seismic data in traces and have become used to thinking of a seismic line that way. We think of it as being a collection of individual traces rather than as being a surface. This concept has been ingrained in many of us because the seismic display that we became familiar with was the wiggle-trace display and it is strictly trace based.

A third reason for our lack of progress is the old adage "necessity is the mother of invention". If we believe this then where is the necessity for a better seismic display? We can see reflections on both the wiggle-trace and variable density displays. We can see character quite well on a wiggle-trace display; we can view large data sets on a computer monitor with a variable density display, so what more is there? What more do we need?

And the final reason for our lack of progress is simple. What more could we do? Given the immense size of some seismic data sets, what other types of display would be practical. It is possible that a single 2D seismic line may exceed 100-million samples. What other type of display is possible on such a large data set?

In this paper, I try to answer some of these questions by taking advantage of some of the advances in graphics technology that were developed originally for the computer gaming industry. As it turns out, one of the main problems that we have with seismic data displays, that being the immense size of a typical seismic data set, parallels a similar problem in the gaming world.

Game designers have been chasing reality for some time but ran into the problem that to model a typical game scene in true 3D requires too many computer resources. To be truly realistic, the scenes require too much data to render quickly. Consequently, they have developed techniques to produce the same realistic effects using far fewer resources.

The principal technique that they use is called bump-mapping because it makes flat surfaces look folded or three-dimensional. I use it here to make a variable density seismic section look like it is really a three-dimensional surface. The resultant display is called a composite density display and as I will show in the examples section, it has the capability to better communicate seismic information to the viewer. It also shows subtle trends in the data that may not be visible on a variable density.

SEISSCAPE: AN INTRODUCTION TO 3D SEISMIC

In the fall of 1999 I was given a crash course introduction into the emerging science of PC graphics cards. At the time I was developing geophysical modelling software and had little, if any, interest in the subject of PC graphic cards. I gave up on computer games at around the Zork stage, so I wasn't even remotely interested in the trials and tribulations of the graphic card industry. But this was to change through the incessant whining of my 13-year-old son and his heart felt (and surprisingly well thought out arguments) as to why his six-month-old graphic card was archaic and ruining his life.

As a normal parent I totally ignored his arguments until he used the magic term giga-flops. At that point I came out of sleep mode and paid attention because Giga-flop is not a term you expect to hear from a 13-year-old. I still didn't hear most of what he said but I did pick up that the board he wanted was $249 and was rated at 4 GF. For me this was a shock because the last time I priced out giga-flops was in 1986 and the cost was $5 million per GF. Obviously I was as old, decrepit and out of touch as Sean was fond of telling me.

This conversation not only got him the board he wanted but it alerted me to the explosive progress that was being made in graphics technology. Eventually this took me into a fascinating experiment with three-dimensional seismic displays. When I say three-dimensional I am not talking about what we normally refer to as 3D seismic but rather the three-dimensional display of what we call 2D seismic.

To expand on this, when we talk about seismic data we typically use terms like 2D/3D/4D seismic. We understand what these terms mean but the truth is we are one dimension short. 2D seismic, for example, is not two-dimensional. Even if we ignore any bends in the seismic line and display it as a straight line, it is still a three-dimensional surface, the dimensions being trace, sample, and amplitude.

It is that third dimension, amplitude, which we have a tendency to ignore. This is due to the fact that our traditional seismic displays, the wiggle-trace and variable density displays, do not have a three-dimensional nature. When you think of it, they are both flat and because of that we tend to think of seismic as existing in a plane. But it doesn't; seismic is, in reality, a three-dimensional surface. We have just never looked at it that way.

Not that I realized this right away. As with most of the industry I was comfortable with wiggle-trace and variable density displays and had little motivation to look beyond them. That not withstanding, once I realized the capabilities of the new graphic cards it became clear to me that we could go beyond. It was clear that we could view small seismic lines, at least, as a three-dimensional surface.

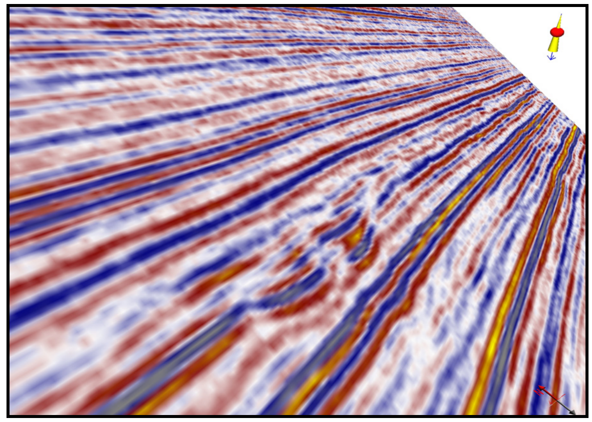

FIG 1. Variable density display of a small channel.

Following up on this idea, in early 2000, I developed a technique to do just that. I called the resultant display a SeisScape® display, see Figures 1 and 2, and I presented it at a CSEG luncheon talk in Calgary in November 2000.

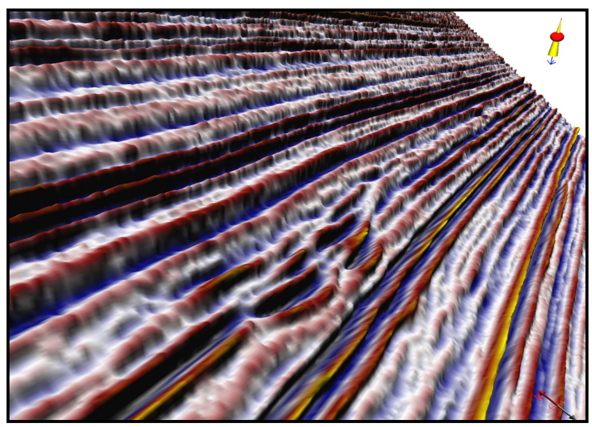

FIG2. SeisScape® display using the same colour palette. Note how much easier it is to follow the subtle events above and below the channel.

What I learned by studying SeisScape® displays was fascinating. When I started the project I did so not realizing that there was much more to seismic than I had seen in the past. But, it turns out that viewing seismic as a true three-dimensional surface has tremendous implications for our ability to perceive faults, amplitudes and subtle character changes. Compare, for example, the image of the small channel on Figures 1 and 2. On the variable density display of Figure 1 the channel is evident but there is nothing inherent in the display that tells us how high the amplitudes of the events in and around the channel are. We can make a mental correlation of colour to height, providing we understand the colour palette. But that requires an imaginative effort that is very difficult to do.

The SeisScape® display of the same channel is a different matter. Because of its three-dimensional nature, and because the image reacts to light, we can easily see how high the amplitudes are and we can also better follow the subtle events both above and below the channel. With the variable density, the channel is flat and lifeless whereas it springs to life on the SeisScape® display. Even with this simple example it is easy to conclude that the information content in a SeisScape® display is much higher than in a variable density display.

SeisScape®, though, is not the subject of this paper and I am only introducing it to build the framework for what I really want to talk about.

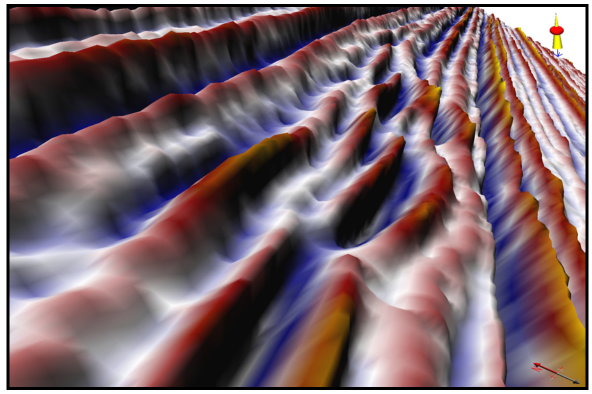

FIG 3. Close-up view of a channel

SeisScape® has a problem and it is one that is mirrored in the world of game design. A SeisScape® display is a fully three-dimensional representation of a seismic section. As such, the amount of resources required to produce it can be prohibitive.

In the world of computer graphics, three-dimensional objects are modelled using triangles. Figures 3 and 4 show a close-up of the channel shown in Figure 2. Figure 4 shows how the display is constructed geometrically by using triangles.

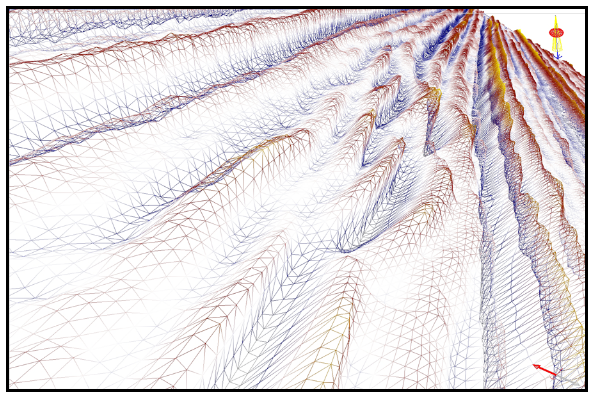

FIG 4. The same view as Figure 3 but displayed in wire-frame mode showing the triangles that make up the display.

Almost all three-dimensional objects that you see in a computer-generated scene are made up of these triangles.

In the case of a seismic display the number of triangles required to display an image is:

$$Tr = 2 \times (Nt - 1) \times (Ns - 1),$$

where

Tr = number of triangles,

Nt = Number of traces,

and Ns = Number of samples.

In the case of a large line with, for example, 10,000 traces and 4,000 samples a SeisScape® display would require almost 80 million triangles to render it completely. This is beyond the capabilities of even the most advanced graphic card.

There are simplification techniques that we can use to reduce the number of triangles. But even if we use them, SeisScape® displays are still very resource intensive and so are, at the moment, most useful for viewing small areas of a line rather than the line as a whole.

But all is not lost. Fortunately, this problem of needing too many triangles to model a display is the same problem that game designers ran into in their search for the holy grail of gaming, immersive reality. In trying to fully model a complex 3D scene, game designers quickly learnt that trying to model every geometrical object in the required detail was, as with a SeisScape® display, computationally prohibitive.

In the remainder of this paper I will introduce the reader to one of the techniques that game designers have used to solve the problem and I will to show how it can be applied to the display of seismic data.

BUMP MAPPING

In the middle to late 90s the computer gaming industry faced a challenge as it tried to produce ever more realistic computer generated scenes. Scenes in early computer games looked very much like cartoon drawings. This was because the objects in the scene did not react to light. Detail could be added by texturing an object. Walls, for example, could be textured with brick-like texture to make them appear realistic. But, a real brick wall reacts to light by reflecting and scattering it. It looks different depending on how the light hits it. Without simulating these lighting effects a textured brick wall would always look like a cartoon, flat and unrealistic.

You run into the same problem when you try to make a character in a scene look realistic. For example, how do you make a reptile's scaly skin react to light? You could always, in theory, model it as a true three-dimensional surface, like SeisScape® does with seismic data. Modelling the shape of the reptile requires relatively few triangles. But to bring the beast to life it must be scaly and the scales must react to light. Game designers found that they couldn’t do it because modelling every scale required several orders of magnitude more triangles.

This is where the expression “Necessity is the mother of invention” comes into play. Game designers needed a level of realism that they couldn’t get. So they turned to a technique called bump-mapping, which was first introduced by Blinn in 1978. It is a technique that makes a flat surface look uneven or warped. It works on the principal that the visual cortex is very proficient at interpreting patterns of light and shade. It turns out that we can tell a tremendous amount about the shape and texture of an object by simply looking at how it reacts to light. We don’t even need to see the actual three-dimensional shape for our visual cortex to figure out what it is. All we really need is light.

This idea of tricking the visual cortex by using shading is nothing new. If you have ever used a geographical system you will recognize this concept as being nothing more than shaded relief maps. Shaded relief maps were introduced to display elevation data (Batson et al., 1975) and have been used for years in other disciplines to highlight structures. In a recent short note in Geophysics, Barnes even used shaded relief to help visualize time slices. He made the mistake of saying, though, that this type of display was of little use for vertical displays. As we will see later, he was wrong.

Bump-maps are an extension of the shaded relief concept. They use what is essentially a shaded relief image to modulate the color of a texture. Since the visual cortex is well adapted to interpreting patterns of light and shade, this tricks it into thinking that the texture is, in fact, a three-dimensional object. And that is how it presents it to us.

To understand how bump-mapping works we first need to understand how lighting is calculated in the computer. A full treatment of the subject can be found in Real Time Rendering 4.3.1 and 5.7.5, (Akenine-Moller, 1999). The basic equation for calculating lighting in a computer is based on Lambert's Law which states that for a completely diffuse material (no shininess) the amount of reflected light is calculated as the cosine of the angle between the incoming light and the surface's normal vector:

$$i = n \bullet l = \cos(\phi), \tag{1}$$

where

i = amount of reflected light,

n = surface normal vector,

l = light direction vector,

and ϕ = angle between n and l.

The first step in the bump-mapping sequence is to produce a texture using this lighting equation. First the surface normals are calculated for every point (sample) on the seismic surface. Then, using a set lighting direction, a texture is created where each pixel is calculated using equation (1). This is the shaded relief map. There are many ways to make this map. Usually it is done algorithmically in the CPU. There is a new technique available now, and which I used here, called Dot3 bump-mapping, where it is done entirely in the GPU (graphics processing unit).

Each pixel in our final texture is made up from the corresponding pixel in both the variable density display and the shaded relief map. The pixel in the shaded relief map contains an intensity only value that varies from 0.0 to 1.0. Each pixel in the variable density display represents an RGB colour value. It has three 8-bit fields that vary from 0 to 255, one each for the red, green and blue element. In this way, black is represented as RGB(0, 0, 0), white is RGB(255, 255, 255), red is RGB(255, 0, 0), etc.

To obtain the final bump-map display each of the three colour components of the variable density's RGB value is multiplied by the corresponding intensity value from the shaded relief texture. In this way, the final image is darker were the seismic surface is facing away from the light. The effect of this can be seen in Figure 8.

Figures 5 to 8 show how the seismic version of a bump-mapped display is constructed. Figure 5 is our starting point, a standard variable density display. Figure 6 is the shaded relief display calculated with a vertical light source. Figure 7 is the same shaded relief display but with the light direction from the bottom left. Finally, Figure 8 is the bump-mapped image formed by combining Figure's 5 and 7.

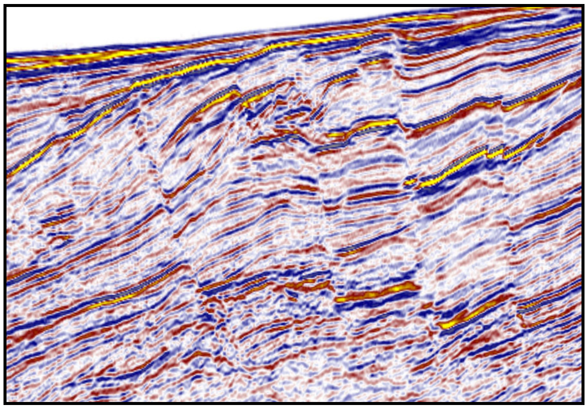

FIG 5. Variable density display of a series of faults.

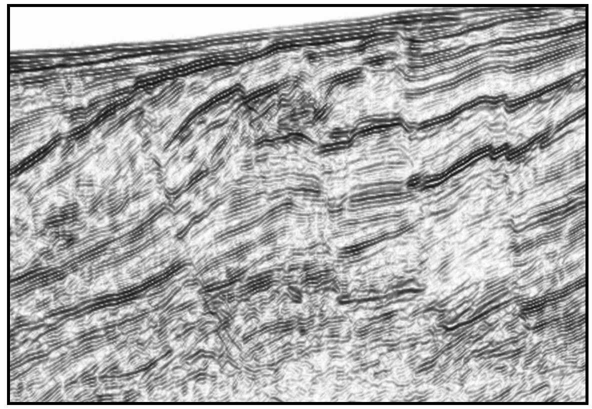

FIG 6. Shaded relief (light density) display illuminated from directly above.

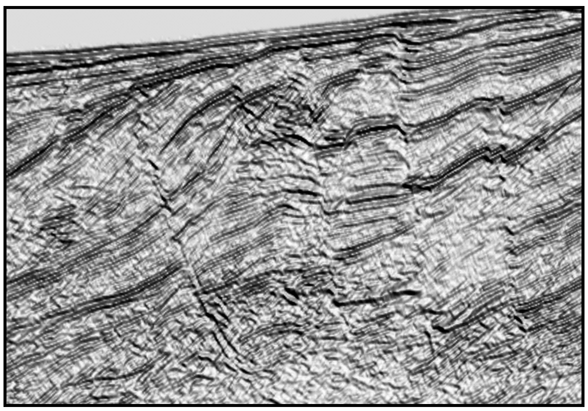

FIG 7. The same light density display but with the illumination from the bottom left.

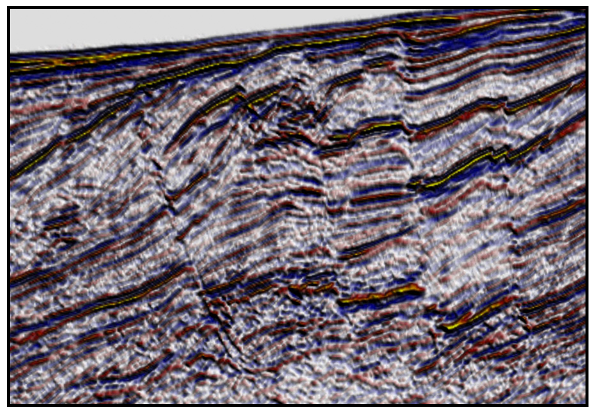

FIG 8. A bump-mapped (composite density) display formed by combining the variable density of Figure 5 with the light density of Figure 7.

As you can see, both the shaded relief displays and the final bump-mapped display look strikingly three-dimensional.

Since the term bump-mapped is non-intuitive and in keeping with the term variable density I am going to introduce two new terms, one each for the shaded relief map and the final bump-mapped image. I am going to introduce the term light density for the shaded relief map and the term composite density for the final bump-mapped image.

Efficiency: Composite Density vs. SeisScape®

One of the great advantages of bump-mapping is that you can now use textures to produce the same three-dimensional effect as you would if you modelled the object in three dimensions. We shouldn't expect that the two will be perfect copies but as you will see, they will be close. The best way to show this is to compare a bump-mapped composite density display with a fully three-dimensional SeisScape® display. I do this next and compare two things: the visual content of the two displays and the resources that are required to produce them.

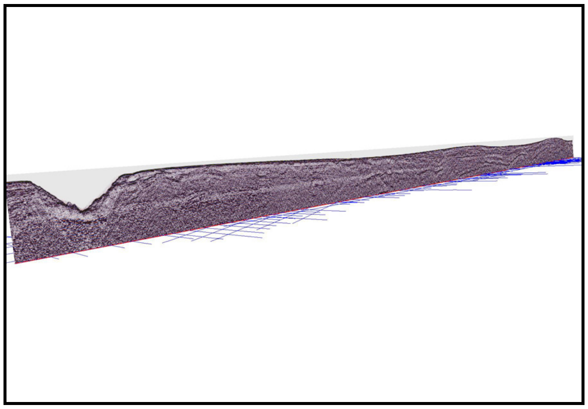

FIG 9. A Composite density display of a 22,000-trace line.

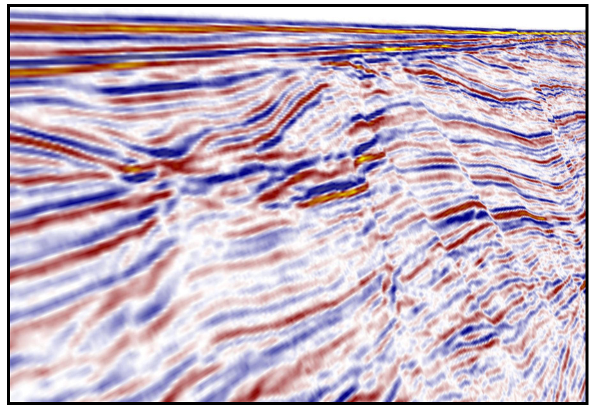

FIG 10. Variable density display zoomed into a small area in the middle of the line.

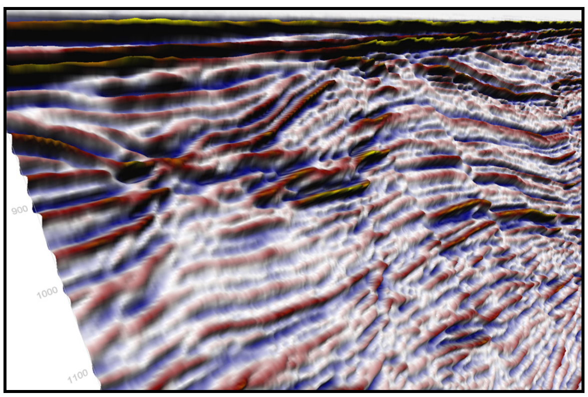

FIG 11. SeisScape® display of the same area. The SeisScape® scene is clipped to contain only the data in the view.

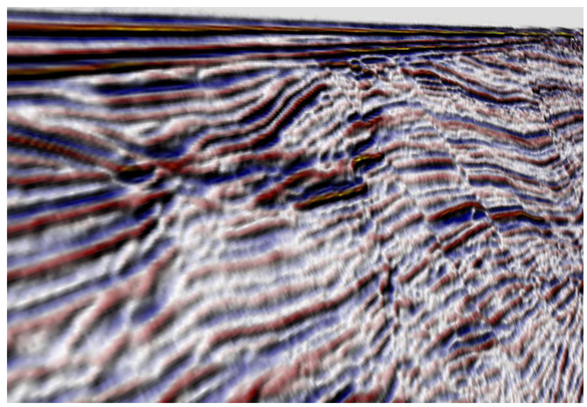

FIG 12. Composite density of the same area. The colour palette, lighting and height are the same as in FIG 11.

Because I want to compare the resources and functional limitations of the two displays, I chose the largest 2D line that I could find. The seismic line is from Peru and contains 22,000 traces with 2,000 samples per trace. Figure 9 shows the entire line as a composite density display. The area that I will focus on is from roughly the center of the line.

As you can see from Figures 9 and 10, both the SeisScape® display and the composite density display show the true three-dimensional nature of the seismic even though the composite density display is flat. Although amplitude variations are, due to its 3D nature, easier to see on the SeisScape® display, it is surprising how much of the information you can see on the corresponding composite density image. At a casual observance, the two look virtually identical. Even though the composite density display is not as three-dimensional as the SeisScape® display, when you compare it to the variable density display of Figure 10 it has almost the same visual impact.

I chose this line because it illustrates the greatest advantage of the bump-mapping approach, animation speed. To produce the true 3D SeisScape® display, I had to clip the data so that the actual display contained 800 traces with 2,000 samples per trace. This resulted in a final 3D surface that had 3.2 million triangles and which rendered at approximately 8 frames per second.

Compare this animation speed with the results of animating the composite density display of the entire 22,000 traces (Figure 9). In this case I was able to animate the lighting on this display and render it at the hard-coded maximum of 60 frames per second. In fact, there was no noticeable degradation in rendering speed between rendering this line and much smaller lines.

As an aside, readers may wonder what type of graphics hardware I used in this test. I used an NVidia GeForce4 Ti 4200 with 128 MB of on-board memory. At the time of writing, this board is about one year old and is considerably less capable than the current top-of-the-line gaming cards. The CPU speed is irrelevant here because the entire bump-mapping and rendering is done in the GPU, the CPU remains idle.

What this example shows is that bump-mapping is a reasonable compromise when it comes to communicating amplitude information. A bump-mapped display, being two-dimensional, is not as realistic a true three-dimensional display and cannot communicate as much visual information. But when you compare Figure's 11and 12 it is obvious that the differences between the composite density display and the SeisScape® display are, for this example, very small. Considering that the composite density is hundreds or even thousands of times less resource intensive than the SeisScape® display, it is a reasonable compromise and it opens the door to viewing seismic as a three-dimensional surface.

EXAMPLES

So far we have answered the question "What more could we do?" but now we need to address the question of "Why would we want to do it?" Considering that variable density displays have served us well for almost a generation, do we really need the composite density display at all? It is certainly a better-looking display but is it a better display?

That is a very difficult question to answer quantitatively. As Feagin wrote in 1981, "How can we know for sure that one type of display is superior to another – and under what conditions?" It is not enough to simply produce a different display or even a better-looking display, we need a display that shows more information. But how can we decide what is better?

The answer is; by looking at it! Seismic is a visual science and the only way to decide if one display is better than another is to look at comparisons and decide for ourselves. To this end I show three sets of examples. The first set focuses on the composite density display's ability to reveal things that we can't see on a variable density display. The second set shows that even if we see the same things on both displays our visual cortex prefers the composite density display. The final set uses composite density display to address a practical problem, fault-plane identification.

In the following section I make my attempt at quantifying composite density displays by examining three questions:

1. Can we see things we couldn't see before?

2. If we can, are the things that we are seeing relevant?

3. Can we better understand what we are seeing?

Processing Artifacts

Any reader who is familiar with seismic data will already be familiar with variable density displays and may be wondering why we need a new display at all. We have been using variable density display for years and it is almost inconceivable that there is pertinent information in a section that the variable density display doesn't show. So I am going to start the examples section by addressing this question head on. For my first example, I examine a processed seismic section and show how effective both the light density and composite density displays are at highlighting subtle effects that are missed by the variable density display.

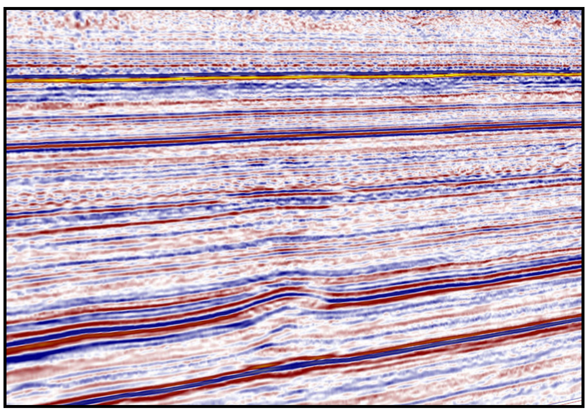

FIG 13. Variable density display of a Zama reef. Note the inconsistent amplitudes and sinusoidal events in the middle of the section.

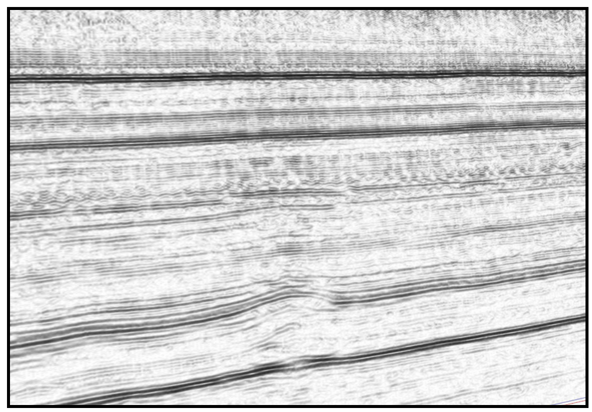

FIG 14. Light density of the same section, illuminated from the noon position (directly above, shining down).

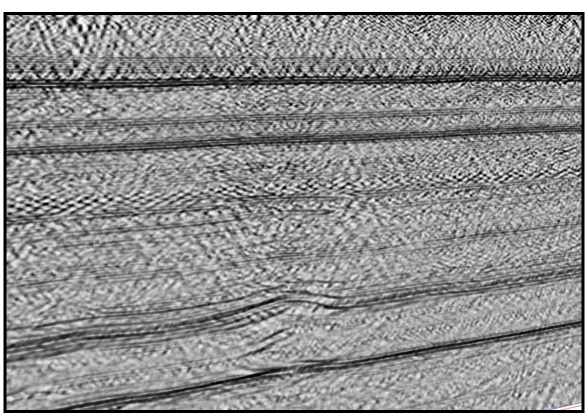

FIG 15. The same light density but illuminated from the lower left. Note the prominent migration artifacts.

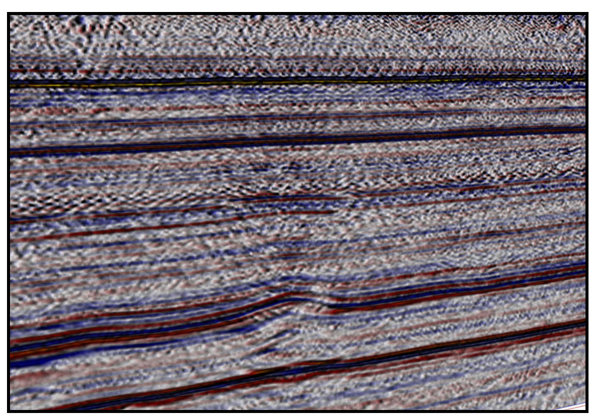

FIG 16. Composite density formed by combining Figures 13 and 15. Again note the high angle migration artifacts.

The data set shown in Figures 11 to 14 is from the Zama Lake region of Alberta, Canada and highlights a prominent reef. An examination of the variable density display shows that the data appears to be somewhat erratic. The amplitudes along most of the events are choppy and inconsistent and there are several events that display and almost sinusoidal pattern. Even if one does not know the geology of the region, it is clear that amplitude effects we are seeing are not caused by any geological event and must therefore be related to the processing.

The actual culprit, though, is hard to see from the variable density display of Figure 13. It is still not evident on the light density display of Figure 14 (illumination pointing from directly above). However, once we start to move the light source to different positions the cause of the section's problems become readily apparent. As you can see from the light density and composite density displays of Figures 15 and 16 (lighting from the lower left), the section is cut through with migration artifacts (smiles).

This is an obvious example, and an experienced processor would not need to use the light density displays to detect that the section had problems with the migration. But an interpreter, who may only get to see this final section many years after the line was processed, may not be aware of not only what the problem is but also what magnitude the problems are.

In this case we are making use of a well-known property of shaded relief maps, that lighting tends to accentuate features that are perpendicular to the light direction. In Figures 15 and 16 the light is oriented perpendicular to the migration artifacts and so they are highlighted in the display.

The fact that the artifacts pop out of the display is a direct result of what we can think of as the visual cortex's pattern recognition engine. When the visual cortex is fed data from the optic nerves it uses its 500 million years of training to detect patterns in the data stream. And this is the problem with this data example. Although the smiles are present in the variable density display, there is nothing in our evolutionary history that has trained our Visual Cortex how to detect them from a colour-only input.

But when we feed it a data stream that contains light intensity only, it knows exactly what to make of it because it has been doing it since eyes first evolved.

I want to end this section on a personal note. I started my career in processing and I still have numerous friends in the processing industry who may not be overly amused at having the flaws in their algorithms so blatantly exposed. To them I add this note.

I started my career with Gulf Calgary in the fall of 1977. At that time Alberta was undergoing a boom in seismic acquisition; so much so that experienced operators were in short supply. As a consequence, the observer's notes that we relied upon to do our line geometries were at the best of times erratic and at the worst, inspired works of fiction.

Gulf was then using a weathering technique called FARR (first arrival refraction), which we, in processing, quickly adapted to check our geometry. I noticed that reversed polarity geophones were easily identifiable on the FARR displays and gave us a way to do the geometry without relying so heavily on the observer's notes. I proposed the idea to management that because of the difficulties we were having with line geometries, that they consider reversing the polarity, in the field, of every 48th geophone.

Unfortunately, being an industry neophyte, I was oblivious to the difficulties of flipping the polarity on a DFS V at -40^0c. The idea was accepted; Gulf started ordering their field crews to reverse every 48th geophone (blizzards and frostbite be damned) and I became, overnight, the single most hated person in the history of field acquisition.

I can only hope that processors have a less violent bent than jug hounds!

Perception

Having shown that a seismic section contains information that we can't see on a variable density display, the next issue concerns the things that we can see. We have just seen that our visual cortex has a hard time deciphering variable density displays, but so far only in the context of picking out things that we couldn't see before. But what implications does this have for the variable density display in general?

With that in mind, I present two examples (devoid of processing artifacts), one an in-line from a 3D survey and the other, a time slice (from a different survey). I have chosen these two examples specifically because there aren't any obvious processing artifacts. In both examples, what you can see on the composite density display you can also see on the variable density display. But the question now is: "From an interpreter's perspective, which one is easier to interpret?"

This is an important question and it takes us right back to the nebulous subject of perception. Seeing is one thing, understanding is another.

Interpreting seismic data is far more than abstractly timing horizons. A good interpreter must put themselves into the geological context of the data they are interpreting. To help them do this when we process seismic data we attempt to provide an illusion that they are actually underground; that what they are looking at is actually real and not just abstract colours on a page. We are, it is true, only showing the seismic representation of geology. It may be in time, it may be filtered, but we would still like it to look real. In fact, we spend an inordinate amount of time developing processing techniques to make the data represent the geology as closely as possible.

But, after we have done that, can we really convey the illusion without light?

This gets back to the subject of perception and the visual cortex's pattern recognition engine. It is very hard to quantitatively determine which of the displays in Figures 17 to 20 is better, but I will argue that it is very easy to know. All you have to do is sit back, relax and then look at the images one at a time and examine how they make you feel. If this sounds like something from the 60s I apologize, but I will argue that when you relax and focus on Figures 17 and 19 you will feel confused. When you do the same thing with Figures 18 and 20 you will feel relaxed.

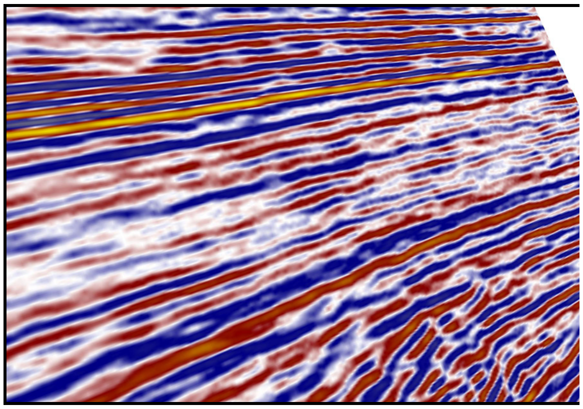

FIG 17. Variable density of an in-line from the Stratton dataset.

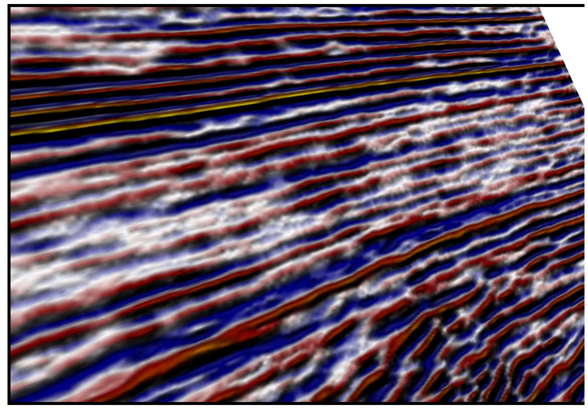

FIG 18. Composite density of the same section. Note how much easier it is to follow the subtle events in the middle of the section.

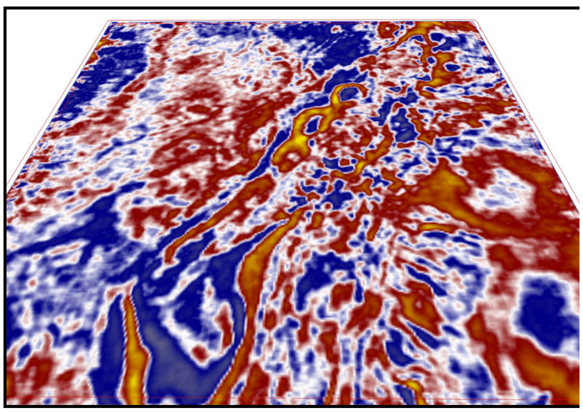

FIG 19. Variable density time-slice from a Peruvian 3D.

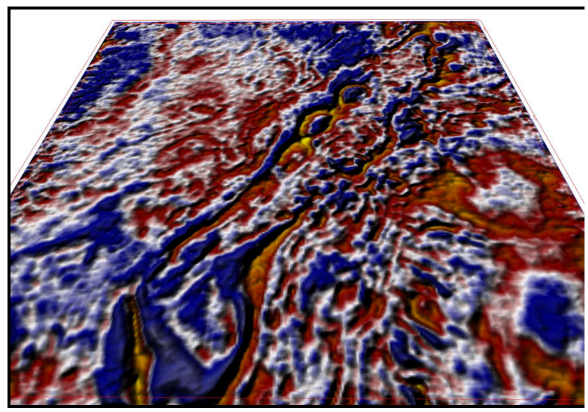

FIG 20. Composite density of the same data. Note how the lighting makes the display less confusing by conveying the illusion of height.

This is a phenomenon that I first noticed with SeisScape® displays, and I put it down to the fact that seismic data in 3D looks very similar to natural landscapes (thus the name SeisScape®). As such they elicit the same emotional response.

Perception, as I have said, is a nebulous subject and very hard to quantify. However, in the context of trying to provide the aforementioned illusion to geology, I will make the claim that the composite density display is far more efficient.

Fault-Plane Illumination

The first two examples that I have shown focused first on processing artifacts which no one wants to see and second on perception which no one understands. So for my final example, and as a break from the esoteric, I will return to reality and show something of direct use. In this case, the use of composite density displays in fault-plane identification.

A quick search of the literature will show that this is an important topic. There has been a lot of research done in the past few years trying to develop techniques to highlight faults and help in their identification. Coherence displays, for example, do a good job of highlighting faults but as with all attribute displays they have a problem. When you look at a coherence, or any other type of attribute display, you lose the seismic itself. What we would like to produce is a seismic display that shows faults with the same clarity.

With that in mind I now show two examples of complex faulting and show how composite density displays can be used to accentuate the faulting.

One thing to keep in mind is that the composite density displays shown here are hampered by the fact that unlike the static variable density display, the composite density display is dynamic. A composite density display relies heavily on the orientation of the light source, which may be moved in real time. As such, the following composite density displays should be considered as being only one possible look at the data. By moving the light source it is possible to highlight the faults better at the expense of the reflections and vica-versa. I have chosen the above displays as representative of an orientation that displays both with equal clarity.

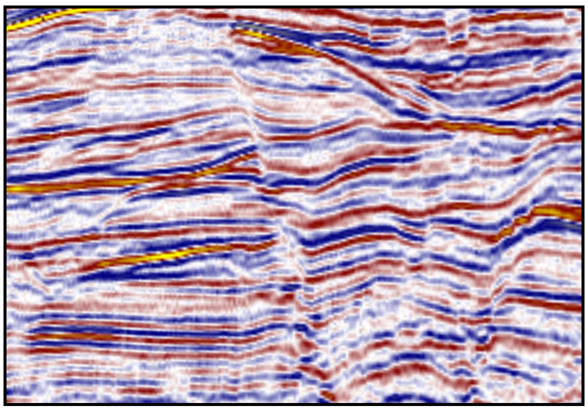

FIG 21. Fault Example 1 – Variable density. We are comparing the fault in the middle of the section.

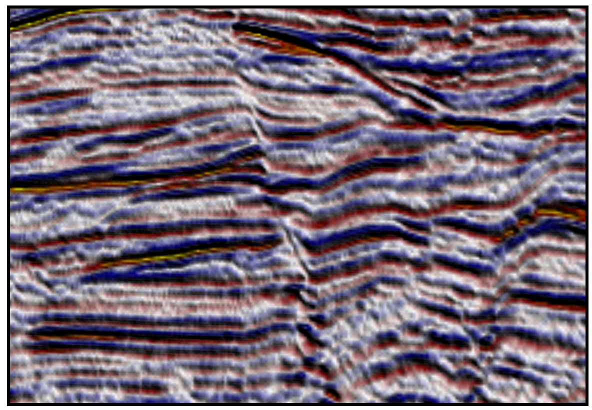

FIG 22. Fault Example 1 – Composite density, lighting from the lower left. Note the appearance of the fault-plane reflections.

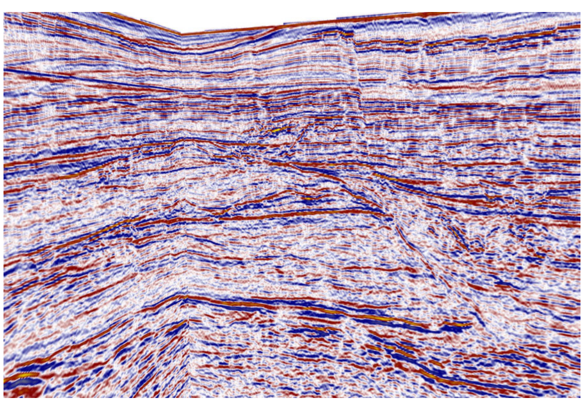

FIG 23. Fault Example 2 – Variable density. Tie between two faulted lines. Note the large fault, centre right.

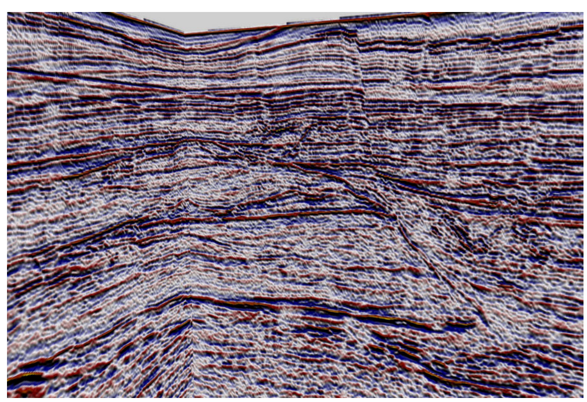

FIG 24. Fault Example 2 – Composite density, lighting from the upper left. Note how the fault is enhanced.

In both examples, the faults are much easier to see on the composite density display. This is particularly evident when looking at Figure 22 because actual fault-plane reflections pop out where there aren't any at all on the variable density display. Figure 22 is illuminated from the lower left. Notice how the illumination has brought out the reflections from the main fault in the middle of the section.

The reader should not be too surprised to see this, as it is a well-known phenomenon of shaded relief displays that light sources tend to enhance events perpendicular to it. That is what is happening in this case. The light source is essentially perpendicular to the observed fault-plane reflections and so we should not be too surprised to see the fault-plane reflection pop out of the background.

The situation is completely different on Figure 24 however. Figures 23 and 24 show a tie between two faulted lines. In both figures there is a large fault that starts in the middle of the image and extends almost to the lower right edge. This fault, once again, is considerably enhanced on the composite density display. What is surprising is that unlike Figure 22 where the lighting was perpendicular to the fault-plane, the lighting here is from the upper left and so is almost parallel to the fault.

This is an interesting, if confusing, observation. Shining the light perpendicular to it enhances one fault; shining the light parallel to it enhances the other. This takes us into the subject of fault-plane illumination, which is too complex to go into in this paper. However, I will make note of some of the observations that I have made to date.

In theory, lighting will highlight features that are perpendicular to the direction of the light. This might lead one to believe that by shining a light in a particular orientation, faults might magically appear. I have discovered that this is not necessarily the case. Due to their geological nature faults are complex seismic objects, far more seismically complex than a horizon.

Depending on the type of fault, the throw on the fault, and the geological sequence, a single fault may represent itself seismically as regions of no amplitude, positive amplitude, and negative amplitude, all on the same fault-plane. Because of this it is crucial to be able to move the light source around quickly to experiment with different light positions. Certain positions may highlight reflections from the fault itself, as in Figure 22 whereas others can pick up reflections from the termination of events (Figure 24). In either case, however, I have so far seen only a weak correlation between the ability to see a fault and the light being perpendicular to it. In fact, in most instances as in Figure 24, faults magically appear when the light source is shone parallel to it. This, though, is the subject of further research.

What this final example illustrates is that by using lighting it may be possible to accentuate faults on the seismic display itself. For these examples I have used a light density display generated from the seismic data itself to produce the composite density images. But it is also possible to combine the variable density display with light density displays produced from other data sources, such as coherence type displays. This may be even more effective at bringing out faults on the seismic sections themselves.

CONCLUSIONS

In this paper I have introduced a new seismic display, the composite density display, which is the geophysical equivalent of a game designer's bump-mapping. Composite density displays are produced dynamically, by modulating the colours of a variable density display with the intensity values taken from a shaded relief (light density) display.

The graphic processing unit calculates the light density display during each rendering cycle thus allowing the user to animate the light source in real time and observe the effect.

Because it contains a lighting component, the composite density display tricks the viewer's visual cortex into thinking it is looking at a three-dimensional surface. The visual cortex interprets the display's patterns of light and shade as being undulations in a surface and so that is how we perceive it. As a result, composite density displays take on many of the same characteristics of the fully three-dimensional SeisScape® display but at much less computational cost.

The use of lighting in the composite density displays makes it not only a different display but also a better display. In my first set of examples I showed that lighting could highlight processing anomalies that are present but too subtle to see on a variable density display. In this first example we clearly identified a series of migration artifacts that were causing the data to look choppy and sinusoidal.

In the second set of examples, I used the lighting to show how the visual cortex is much better able to understand what it is seeing.

My final examples illustrate the usefulness of lighting in fault-plane identification. In both examples, fault-planes are much clearer on the composite density displays than they are on the variable density displays. Interestingly though, in one example the light was perpendicular to the fault whereas in the other it was parallel. This shows that it is very important to have real time control of the lighting because it is hard to predict what will show up at any given light orientation.

Seismic displays have remained essentially static since the evolution of the variable density display in the mid 1970's. Composite density displays use techniques developed initially for game developers and are just one possible new way to look at seismic data. As game development continues we should be able to borrow some more of their techniques to help us better visualize and understand our own immersive reality.

ACKNOWLEDGEMENTS

I would like to thank Divestco Inc for permission to publish this paper. I also want to thank PeruPetro for permission to show the Peruvian 2D and 3D data sets and SeisView Surveys for permission to show the Zama Lake data.

I would also like to credit Dennis Meisenger of Divestco Inc for suggesting the name composite density display and Larry Lines of CREWES for his encouragement.

And finally, I would like to thank my now 18-year-old son Sean for forcing me onto this path in the first place.

REFERENCES

Akenine-Möller, T. and Eric H., 1999, "Real-Time Rendering", A.K. Peters Press, chap 4-5.

Barnes, A. E., 2003, Shaded relief seismic attribute: Geophysics, **68**, 1281-1285.

Batson, R. M., Edwards K. and Eliason, E.M., 1975, Computer generated shaded relief images: J. Research, U.S. Geol. Surv., **3**, 401-408.

Blinn, J., 1978, Simulation of wrinkled surfaces: Computer Graphics (SIG-GRAPH '78 Proceedings), 286-292

Taner, M.T., Koehler, F., and Sheriff, R.R., 1978, Complex seismic trace analysis: Geophysics, **44**, 1041-1063

Feagin, F. J., 1981, Seismic data display and reflection perceptibility: Geophysics, **46**, 106-120.

Lynch, S., 2000, "Ancient evenings, seismic visualization using very old techniques: CSEG Recorder, **25**, 5-6