A Case Study in Multi-Sensory Investigation of Geoscientific Data

Chris Harding[1,3], Ioannis A. Kakadiaris [1,2], John F. Casey[1,3] and R. Bowen Loftin[4]

[1] Virtual Environments Research Institute, Univ. of Houston, Houston, TX, USA

harding77019@yahoo.com

[2] Department of Computer Science, University of Houston, Houston, TX, USA

ioannisk@uh.edu

[3] Department of Geosciences, University of Houston, TX, USA

jfcasey@uh.edu

[4]VMASC, Old Dominion University, Suffolk, VA, USA

bloftin@odu.edu

Abstract: In this paper, we report our ongoing research into multi-sensory investigation of geoscientific data. Our Geoscientific Data Investigation System (GDIS) integrates three-dimensional, interactive computer graphics, touch (haptics) and real-time sonification into a multi-sensory Virtual Environment. GDIS has been used to investigate geological structures on the high-resolution bathymetry data from the Mid-Atlantic Ridge. Haptic force feedback was used to precisely digitize line features on three-dimensional morphology and to feel surface properties via varying friction settings; additional, overlapping data can be perceived via sound (sonification). We also report on the results of a psycho-acoustic study about the absolute recognition of sound signals, and on the actual feedback that we have received from a number of geoscientists during a recent major geoscience conference.

1. Introduction and Background

In recent years, the natural resource industry has recognized three-dimensional visualization and modeling of geoscientific data as playing an important part in the exploration and development of natural resources. Several research projects have demonstrated that the use of Virtual Environments has the potential to improve productivity and lower costs in areas such as petroleum exploration [3], [5]. Most current virtual environments focus entirely on improving the user's comprehension of geoscientific data by using true three-dimensional (stereo graphic) environments, large screens and cooperation between different disciplines. However, the fields of haptic force feedback devices and real-time sound synthesis have matured sufficiently in recent years to allow research of the integration of touch and sound into visual virtual environments. Both technologies have seen some application in the geoscientific domain [1], [2].

The aim of our research is to integrate the ability to feel and interact with data via touch, and to analyze data via hearing into interactive three-dimensional systems to give the user the advantage of working with multiple, overlapping data properties simultaneously. Although the visual sense is still the main channel, presenting other aspects of data simultaneously through touch and sound could lead to enormous benefits if we succeed in mapping data from its scientific domain into a useful representation of touch and sound (the reader is referred to [4] for further information). This paper focuses on a novel sonification technique and its integration with visualization and force feedback to create highly interactive virtual environments for investigating geoscientific data.

In our experience, a demonstration of GDIS to potential users is a vital research tool. After geoscientists gathered first hand experience with this new technology, their feedback lead to valuable research contributions. The system has been very well received by participating geoscientists. Their input helps to open up new areas of potential geoscientific applications for mapping, displaying and perceiving complex multi-attribute data sets simultaneously.

2. The Geoscientific Data Investigation System

We have created a demonstration prototype called GDIS (Geoscientific Data Investigation System, which allows the multi-sensory investigation of geoscientific surface data, on which several geophysical properties has been mapped (e.g., gravimetric and magnetic data). The term *investigation* is used to emphasize that three major senses (stereo vision, touch and sound) are used for what is usually called *vis*ualization and modeling. GDIS uses all three senses to simultaneously explore different, overlapping surface properties and accurately digitize lines on the surface.

Easy access to the development of human-computer interaction via touch has become possible in the last five years with the development of haptic force feedback systems such as Sensable Technology's PHANToM. We employ the Desktop PHANToM in our system (Figure 1), which can project a point force of up to 6.5 N within its three-dimensional workspace (16 cm x 13 cm x 13 cm). The PHANToM creates the illusion of touching solid objects with a virtual fingertip and the feel of physical effects such as attraction, repulsion, friction and viscosity. The stylus at the end of the arm gives the user force feedback and direct interaction with the data in three dimensions. Using this "three-dimensional force feedback mouse" allows not only the interrogation of three-dimensional data within a virtual space, but also provides input constraints for fine-grained interaction. In our research, we have concentrated on the haptic rendering of *surface data* (which is a common data type across the geosciences) onto which three-dimensional lines can be digitized. This allows the user to explore minute features of the surface's morphology and to use the tactile feedback to model line and polygonal structures. In addition, we translate surface properties into different friction values in order to make it noticeably easy or difficult to move the stylus tip over certain parts of the surface.

The use of non-speech sound to interactively explore data (scientific sonification -"the use of data to control a sound generator" [6]), has been under investigation since the early 1990s. Although recent advances in real-time sound synthesis have made this technology more widely available to researchers, there are only very few guidelines on mapping data into sounds, and there seems to be relatively little research that is related to geoscientific data [3]. It seems to be clear, however, that sound needs to not only be effective but also pleasant (or at least not disturbing). As the way different people react to sound can be very different, sound mapping in an application may need to be highly customizable to a specific user. The human hearing system is very well suited to detect even small changes in sounds (e.g., has very good temporal and pitch resolution) but has difficulties determining absolute values.

One advantage of sonification is that the user's eyes are free to process visual data while hearing a different set of data. We have integrated a novel "sound map" into the visual rendering of surfaces, giving the user the ability to listen to a local surface property while simultaneously visually observing other properties.

3. Absolute Recognition of sound—a psycho-acoustic user study

We conducted a psychological study into the use of sound to convey data in an absolute way. We use the term *absolute* signal (as opposed to a *relative* signal) to refer to a signal that is evaluated outside the context of previous or following signals. Each signal is initially defined as corresponding to a certain concept, in our case the numbers one to five. The test subjects were asked first to remember the sound and later to recognize it again—making a connection back to the numeric definition given earlier.

It is well known that the human hearing system is quite capable of discriminating even small changes in audio signals in terms of pitch or tempo [7]. However, the absolute recognition of signals is considered much more difficult and usually dependent on musical abilities and training (the “perfect pitch” phenomena is an extreme example for the absolute recognition of sound). Because absolute recognition is acknowledged to be a difficult problem, our study concentrated on a fairly simple setup. In our study, we defined an audio signal by three simple parameters. These parameters, pitch (frequency), instrument (timbre) and tempo (tone repeat rate) are used to generate (musical) tones via a MIDI (Musical Instrument Digital Interface) synthesizer.

The goal of our study was to determine if the subjects could be trained to reliably recognize a certain small set of audio signals, and connect them back to an equal number of numeric values. We chose a small number of integer values to represent a partition of the overall data set in a number of “bins”. This so-called “ballpark setup” connects, in our case, five logically progressive musical notes to a sequence such as “very low, low, medium, high and very high” in order to give the user a rough idea about local data values rather then their actual, precise values. This technique was used in GDIS in addition to the more traditional relative sonification, where a change sound symbolizes a general, relative change in the sonified data set; for example a rise in pitch could be used to sonify a rise in temperature.

3.1 Setup of the experiment

The idea of the ballpark scenario raised several interesting questions: How should the audio signals (using pitch, tempo and instrument) be composed to allow for a short training and easy recognition? Should we use only single parameters (e.g., only pitch or only instrument) or some kind of combination? The study aimed to answer these questions by reducing the ballpark scenario to a simple “game” in which the test subject was guided through seven conditions and asked to recognize a series of audio signals from a set of five different sounds. We postulated, that, not only should the subjects be able to learn how to differentiate between audio signals, but also, that the subjects would be more successful in recognizing conditions where more than one parameter was changing (conditions 4 to 7). The composition of the five audio signals for each of the seven conditions is shown in the table below.

Condition	**Setup of the five audio signals**
1	*Pitch* only: single note C in five different octaves: C2, C3, C4, C5, C6, played by a grand piano.
2	*Instrument* only: single note C4 played by an church organ, a grand piano, a pan flute, a muted trumped and a cembalo.
3	*Tempo:* 3 clearly separate notes (C4 played by grand piano) with speed increasing in five stages from overall 1.5 to 0.3 seconds.
4	Combination of pitch and instrument (single notes).
5	Combination of pitch and tempo (played by grand piano).
6	Combination of tempo and instrument (playing a C4).
7	Combination of pitch, instrument and tempo.

The study was conducted on a standard SGI workstation (Octane or O2) equipped with small PC speakers and a software MIDI synthesizer (midisynth), which is part of the SGI digital media software package. At the beginning of each experiment, the subject was asked to complete a questionnaire about his or her age, gender, musical background, listening habits and professional background. The subject was guided through seven conditions. Each condition began with a *training phase* (Figure 1a), in which the subject was familiarized with the current meaning of the five audio signals. This was followed by a *recognition phase*, in which the subject was asked to identify a random series of 15 audio signals (which was different for each condition but the same for each subject), and indicate his or her confidence in the choice.

Figure 1b depicts the recognition phase, where the trained subject plays the audio signal (PLAY), chooses a number (1–5) that corresponds to the audio signal, and indicates his or her confidence (0–100). We also recorded the time spent for the recognition of each condition. The total time spent to complete the experiment varied between 20 and 40 minutes.

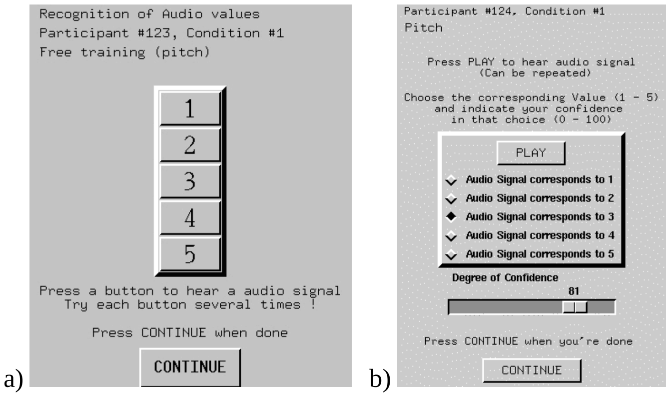

Figure 1. Example snapshots of the training phase (a) and of the recognition phase (b).

3.2 Results of the study

A total of 13 subjects (11 male, 2 female) were tested in this study. Their ages ranged from 21 to 55 (average 33.15). From the background questionnaire, it appears that the subjects are fairly evenly mixed in most categories. Three subjects have an undergraduate degree, five a Master's degree and five a Ph.D. degree. Four of the subjects are students, four are researchers and four are professionals. None of the subjects were music majors or professional musicians, although some like classical music and some are amateur musicians. However, the subjects' musicality had no significant impact on their performance. The results of three dependent variables were analyzed: success rate, time to recognize, and subjective confidence. To ensure statistical significance, Analysis of Variances tests (ANOVA) and planned comparison tests were conducted for all dependent variables.

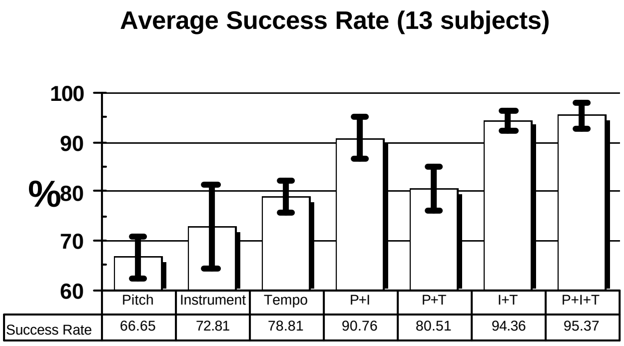

Figure 2. Average success rate across the seven conditions.

Figure 2 depicts the percentage of correct responses for all 13 subjects for each condition (with 15 audio signals per condition) along with the standard error bars. The standard error is a measure of variability in the sample and is defined as the standard deviation divided by the square root of the number of subjects. It appears, that of the three simple conditions (pitch, instrument, tempo) subjects are most successful in recognizing tempo, while least successful in recognizing pitch. This agrees with the generally accepted assumption that temporal variations are handled well by the human hearing system.

Although all of the more complex conditions (conditions 4–7) showed at least marginally higher success rates than any of the three basic conditions, only conditions 6 and 7 resulted in a significantly better performance with little variability among the subjects. This agrees with our assumption that these more complex combinations should be easier to recognize and therefore result in higher success rates. Although subjects had lower success rates in condition 2 (instrument only) than in condition 3 (tempo only), the instrument property seems to positively affect performance when used within a combination. Of the conditions that combine two properties, the change in the instrument property seems to make the recognition more successful. For example, condition 4 (Pitch + Instrument) and condition 6 (Instrument + Tempo) are more successful than condition 5 (Pitch + Tempo). Although condition 7 (Pitch + Instrument + Tempo) seems to result in higher success rates, it is not much different from condition 6 (Instrument + Tempo). This may indicate that pitch does not contribute much to the success rate within a combined condition.

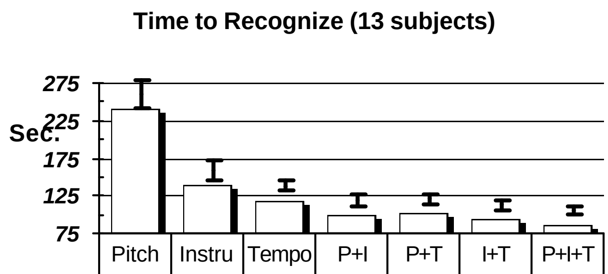

	Pitch	Instru	Tempo	P+I	P+T	I+T	P+I+T
Time	241.3	139.15	118.53	98.53	101.01	92.69	85.46

Figure 3. Average total time needed to recognize a series of 15 audio signals.

Figure 3 depicts the average time to complete the recognition phase (in seconds) for all subjects for each condition with the standard error bars. Note how much faster the subjects are able to recognize the signals in conditions after the first one (nearly twice as fast!). This time still improves in conditions 2 and 3 and reaches a plateau for the more complex conditions (4 to 7). This improvement may indicate a very strong leaming curve, were learning is essentially complete after the subject has become familiar with three different sound properties (after condition 3).

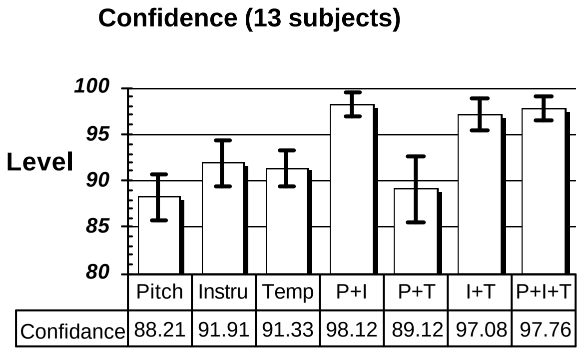

Figure 4. Average subjective confidence levels.

A measure of the subject's confidence about his or her choice was recorded on a scale of 0 to 100. Figure 4 depicts the average confidence level for all 13 subjects for each condition; along with the standard error bars.Of the four more complex conditions, subjects reported the least confidence in condition 5 (Pitch + Tempo), with only slight differences among the other three conditions (4, 6 and 7). This unexpected finding is interesting because condition 5 is the only complex condition that lacks a change in the instrument. A similar trend was also observed in the success rate. This suggests that the combination of changing instrument and either pitch or tempo may enhance both success in recognizing the audio signal as well as confidence in the ability to correctly recognize the audio signal.

The study demonstrated that any subject, musical or not, can be trained to differentiate between five different audio signals and connect those audio signals with numbers 1 to 5 (a so-called ballpark setup). After running though several different conditions (scenarios), each divided into a short training phase and a recognition phase, the subjects achieved a 95% success rate for a recognition scheme that simultaneously varies pitch, tempo and instrument. This result indicates that it is possible to use the auditory channel for applications that need to convey a secondary data stream and training can be achieved in a short period of time (20 to 40 minutes).

4. Multi-sensory investigation of geoscientific data with GDIS

We conducted a case study of GDIS's multi-sensory abilities by investigating surface-mesh based geoscientific data. A high-resolution elevation model of a seafloor (derived from multi-beam bathymetry data) was texture-mapped with the residual mantle-boguer-anomaly (RMBA gravity map, see Figure 6). The gravity is visually mapped with a Blue-Green-Red color map, a map of the age of the oceanic crust (calculated from magnetics) was mapped into the audio domain (sound map) and the change of slope of the bathymetry was expressed in a "friction map".

As a result of this study, the sonification in GDIS uses both a relative and an absolute form of sonification. Only a single tone, defined by pitch, tempo and instrument are played. The data value of the sound map is used to synthesize the current tone with a MIDI synthesizer. The pitch and the tempo use a simple linear mapping: low data values map to a low pitched, slow sequence of notes whereas high data values map to a series of high pitched, fast notes. Most users of GDIS seemed to intuitively understand this form of mapping. The data represented by the sound map can be divided into five “bins”, which represent, for example, a progression from “very low” and “low” over “medium” to “high” and “very high”. These bins and the concepts connected to them can be heard via the instrument property. By assigning a distinct instrument to each bin (for example a tuba for the lowest data values and a piccolo flute for the highest data values), the trained user can hear (via pitch and tempo) not only how the data are changing, but also in what general part of the data (ballpark) it currently falls.

When the PHANToM's “virtual fingertip” (represented graphically by a cone) touches the surface, the force feedback allows the user to feel the sometimes-delicate surface features and get an idea of the age of the surface at this point by listening to the pitch, instrument and duration of the currently played notes. The user is also able to perceive inflection points by feeling an increase in friction. The user navigates by “grabbing” the surface with the stylus. Holding down the stylus button attaches the surface to the stylus. With the surface attached to the user’s hand, the user can easily look at the surface from different angles and distances and, in general, investigate the change of surface curvature. The application offers a special surface coloring, which shows not only the magnitude of the slope, but also its direction as a change of color (Figure 7). By changing the position of the global light source (virtual sun) and by using a virtual flashlight attached to the end of the virtual fingertip, the user can enhance the appearance of geologically interesting surface structures.

Besides the aforementioned settings, GDIS can be configured to freely explore other combinations of visual (texture) maps, sound maps and friction maps by allowing the user to designate any surface variable as either a visual, sound or friction map. Although the pure exploration of a data set by multi-sensory means is one important aspect of our research, we also augmented the commonly employed, interactive task of digitizing line segments onto a surface and extracting surface data from digitized polygons. We were particularly interested in the PHANToM’s ability to interact with surface morphology while simultaneously receiving input about several surface attributes via visual, acoustic and haptic means. Hitting the space key on the keyboard digitizes line segments–a new point will be dropped exactly where the virtual fingertip touches the surface. The system then drapes a new three-dimensional line segment on the surface. Line segments can be closed to form polygons from which internal points can be extracted. Figure 6 depicts the seafloor data set after digitizing an important geological feature, a major fault line (bright line left).

Using GDIS, we have investigated a tectonically interesting area on the Mid-Atlantic Ridge, where a new oceanic crust is created in water depth of 3000 m to 6000 m. By precisely digitizing fault structures while simultaneously accessing other surface properties, we are working to improve models of recently discovered dome-like structures called mega-mullions, which promise to grant insight into the deeper structures of the ocean floor. Figure 7 depicts a structural model of fault planes directly the on Mid-Ocean ridge. Note the “virtual fingertip”, represented as a cone and its flashlight-effect (lower left corner). In the future, we plan to apply this system to other surface data with multiple overlapping properties. For example, in the remote sensing domain, GDIS may offer a way to aid in the interactive planning of a pipeline by giving access to multiple constraints at once.

5. User-feedback and Discussion

GDIS was presented on several occasions to a number of geoscientists at the Annual Conference of the Society of Exploration Geophysicists (SEG 2000) in Calgary, Canada in August 2000. At the conclusion of the twenty minute demonstration of GDIS' functionality, the audience was invited to experience the multi-sensory system first-hand by "test-riding" GDIS. About 50 geoscientists with different backgrounds used the system for different periods of time (5 min to 60 min), during which they were brought into contact with all major points of the system: three-dimensional visualization, haptic feedback, and sonification. All users were geoscientists, from a variety of geoscientific backgrounds including seismic interpretation, visualization, remote-sensing/GIS, structural geology, seismic acquisition, and geophysical data processing. Some of the users (27) recorded their impressions on a short questionnaire. This section reports on the results of these impressions. The results may be indicative of general trends and provide feedback with regard to a deployment in other parts of the geosciences. The users spent between 5 and 60 minutes with the demonstrator (13.15 minutes on average). Of the 27 geoscientists, 18 use some form of digitizing as part of their work. The user's ratings for the system's three components and the overall system (Figure 9) were very high: on the average the graphics part was rated 8.52; haptics was rated 7.59; sound was rated 6.52 and the overall system 7.93.

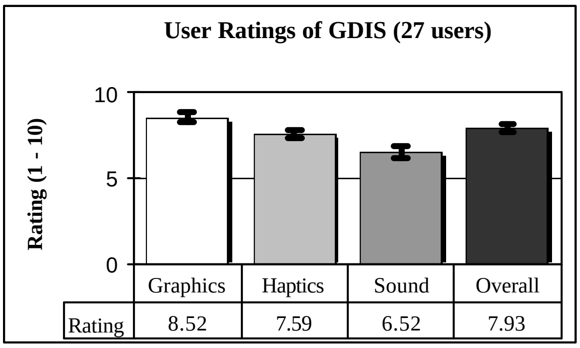

Figure 9. GDIS User Ratings with standard error bars.

Among the users, 16 could see a system like GDIS as useful in their current job and that it would add value to it. When asked about what they liked most about GDIS, the users mentioned the ability to perceive information through more than the visual channel, the ability to navigate and manipulate in three-dimensional with the PHANToM and the ability to switch between the surface's attributes. The following list contains some of the users responses:

- The ability to realize unseen information
- Haptics provides additional sensory input for investigation surfaces
- Ability to follow faults in the "inverted" surface
- PHANToM is a natural tool for navigation and manipulating three-dimensional objects
- Feeling the surface to "visualize" the structural framework
- Interesting application of sound
- Tactile feedback
- Navigation and the spotlight
- Liked the sound idea (hearing your data)
- Graphics (rotating, illumination, texture)
- The way to hear "steps" between instruments
- Overall system integration
- Interactive feedback
- Multiple parameters sensed simultaneously
- The uniqueness of the touch/hear system
- The option, ease to change between attributes
- Integration of senses - ability to query multiple data sets at once

When asked what they liked least, the users mentioned the lack of available customization, problem in distinguishing between different levels of friction and the limitation to perceive the volume of areas with the stylus tip. Specifically:

- Sound is somewhat distracting but could be useful for listening to additional datasets
- Needs user interface to customize sound/haptics for parameters
- Needs to be personalized for each user
- Hard to reliably distinguish the friction component
- Did not feel the surface friction very clearly
- Hard to scan large surfaces or volume with single point

When asked, if the users could imagine an area for a “killer application” they mentioned seismic interpretation with multiple attributes, aiding GIS-related planning tasks and help for the visually challenged:

- Digitizing faults on coherence time slices
- Planning of logistics and routes

- Shaping the surface to volume data
- Medical (surgical) application
- Adding 3D graphics to GIS
- Color blind, blind interpreters, interactive visualization, learning tool
- Attribute interpretation based on engineering and geophysical data
- Assisting seismic interpretation by tracking amplitude and hearing other attributes simultaneously

GDIS was generally very well received and praised for its novel approach. Most users were exploration geophysicists with a seismic interpretation or visualization background from the oil and gas or minerals domain. Although the demonstrated task, the structural exploration of high-resolution bathymetry data with gravity data and age data mapped onto it, was not familiar to many users, most of them could see a deployment of the system for a similar task closer to their own work. Users specifically liked the ability to receive information about “invisible” data via friction or sound while exploring the surface and the ability to explore and to interact with the surface in true 3D.

The demonstrations were expected to stimulate the users about the potential deployment of multi-sensory systems in their specific line of work. Several users voiced suggestions for applications in other areas, for example seismic interpretation and planning tasks based on remote sensing/GIS type data. In the seismic interpretation area, it was suggested to be used for modeling of faults or other structural elements on subsurface horizons with multi-sensory access to several different seismic attributes such as amplitude or coherence or geostatistically derived properties such as rock facies, permeability, or porosity. For the remote sensing area, users suggested possible applications related to planning of structures such as pipelines or roads on topographic or bathymetric surfaces within a three-dimensional GIS system, with haptics conveying the slope and sound warning about unseen problems.

6. Summary

In this paper, we have presented a system that integrates a novel sonification method and haptic force feedback with three-dimensional visualization to form a multi-sensory user interface for exploring geoscientific data. The system has been developed for the specific geoscientific task of investigating geological structures on the seafloor. However, it is generic enough to be used in the exploration and modeling of other geoscientific surfaces. The system has been demonstrated to geoscientists and was very well received. Many geoscientists acknowledged the system’s potential for a future deployment in other parts of the geosciences and provided valuable feedback.

Using a combination of graphics, haptics and sound to solve a particular geoscientific task is still a very new field that needs to be explored in a wider context. The fields of haptics and sonification are still at an early stage of their development––perhaps comparable to computer graphics 10 or 15 years ago. Only time will tell if these fields will experience a similar explosion of technological possibilities over the next years and become an integral part of our life as well. The way that our system makes use of the three senses is clearly just a first step and a glimpse of what could be possible in the future.

References

1. Aviles, W.A. and J.F. Ranta. Haptic Interaction with Geoscientific Data. In: Fourth PHANToMS User Group (PUG) Meeting. 1999. Cambridge, MA.
2. Barrass, S. and B. Zehner. Responsive Sonification of Well-logs. In ICAD 2000. 2000. Atlanta, GA: ICAD.
3. Fröhlich, B., et al. Exploring Geo-Scientific Data in Virtual Environments. In: IEEE Visualization. 1999.
4. Harding, C., Multi-Sensory Investigation of Geoscientific Data: Adding Touch and Sound to 3D Visualization and Modeling, PhD Thesis, Department of Geosciences. 2000, University of Houston: Houston. p. 170.
5. Harding, C., R.B. Loftin, and A. Anderson, Visualization and modeling of geoscientific data on the Interactive Workbench. The Leading Edge, 2000. 19(5): p. 506-511.
6. Kramer, G., An Introduction to Auditory Display, in Auditory Display: Sonification, Audification and Auditory Interfaces, G. Kramer, Editor. 1994, Addison Wesley: Reading, MA. p. 1-58.
7. Moore, B.C.J., An introduction to the psychology of hearing. 4th ed. 1997, Orlando, FL: Academic Press.

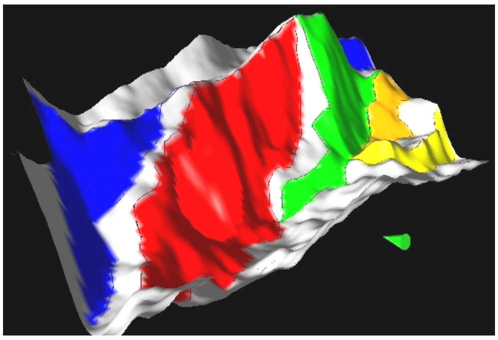

Figure 5: GDIS uses stereographics, haptic force feedback and sonification.

Figure 8: Structural Model created with GDIS

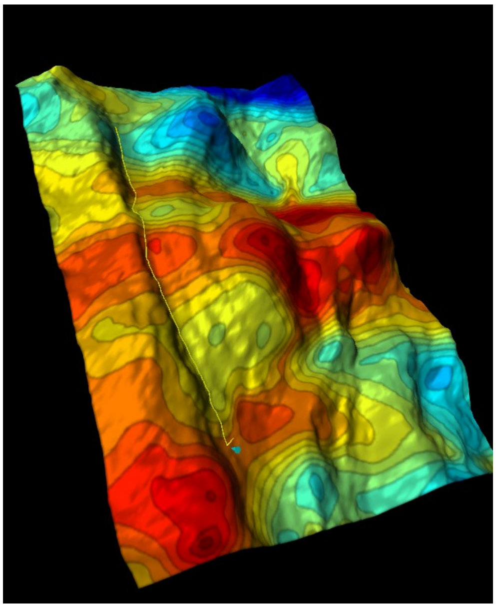

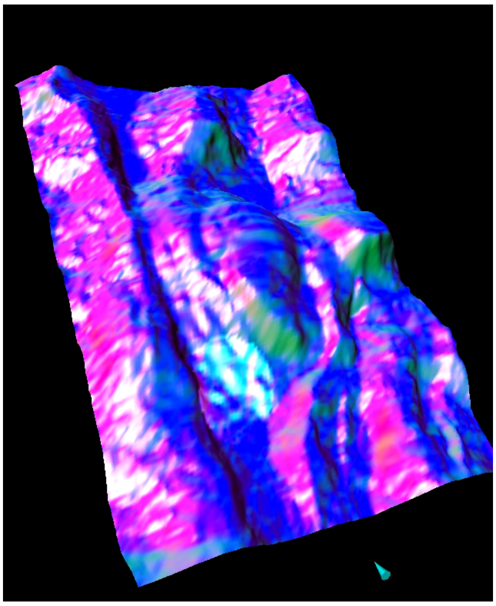

Figure 6. Depiction of gravity data texture-mapped onto the seafloor.

Figure 7. Depiction of the seafloor data with curvature coloration