Information Visualization (2004) 3, 36–48

www.palgrave-journals.com/ivs

Feel the information with VisPad: a large area vibrotactile device

Doanna Weissgerber[1]
Bruce Bridgeman[2]
Alex Pang[1]

[1]Computer Science, Computer Science Department, University of California, Santa Cruz, CA U.S.A.; [2]Psychology and Psychobiology, Psychology Department, University of California, Santa Cruz, CA, U.S.A.

Correspondence:
Alex Pang, Computer Science, University of California, Computer Science Department, Santa Cruz, CA 95064, U.S.A. Tel: +1 831 459 2712; Fax: +1 831 459 4829;
E-mail: pang@soe.ucsc.edu

Received: 13 July 2003
Revised: 15 October 2003
Accepted: 15 October 2003

Abstract

A new haptics design for visualizing data is constructed out of commodity massage pads and custom controllers and interfaces to a computer. It is an output device for information that can be transmitted to a user who sits on the pad. Two unique properties of the design are: (a) its large feedback area and (b) its passive nature, where unlike most current haptics devices, the user's hands are free to work on other things. To test how useful such a device is for visualizing data, we added the VisPad interface to our protein structure-alignment program (ProtAlign) and performed usability studies. The studies demonstrated that information could be perceived significantly faster utilizing our multi-modal presentation compared to vision-based graphical visualization alone.
Information Visualization (2004) **3,** 36–48. doi:10.1057/palgrave.ivs.9500060

Keywords: Haptics; usability study; bioinformatics; multi-modal; user interface

Introduction

For data sets up to 3D, visual display is often adequate. As data dimensionality increases, however, it becomes increasingly difficult to display images in an easily comprehensible form. What happens when you have exhausted the easily distinguishable colors, textures, and shapes possible in visual perception? What if you spend so much time interpreting an image that graphical visualization becomes a less effective tool?

Scientists have added new modalities to enhance our understanding of multi-dimensional data. Sonification has been shown to enhance the user's ability to understand data quickly.[1] But sound lacks a high-acuity spatial dimension, and when the reasonable (easily discerned) sound cues such as instruments, notes, tempo, etc. also become confusing, where does the scientist turn? We turn to the field of haptics for additional sensory input routes.[2]

Currently, most 'normal' haptic feedback devices tie up at least one hand or cover the hand with a glove (e.g. CyberGlove with CyberGrasp or CyberTouch by Immersion Inc., PHANToM by Sensable Technologies, or Haptic mice by Logitech). Gloves weigh almost 400 g and are tiring during prolonged use, and the PHANTOM takes up valuable desktop space.[3] Also, haptic feedback through the hand is not always appropriate. Because 'normal' applications use keyboards, it is intrusive to tie up a user's hand. If the user is typing, feedback from devices such as haptic mice is lost. Hand-oriented devices also make it impossible to just sit back and understand the visualization of multi-modal data.

There is research looking at haptic visualization using motors, air jets, and skin stretch. But again, most research is focused on the hand as the only haptic output channel. Hayward and Cruz-Hernandez[4] developed a version of the Pantograph, a force feedback device intended for direct manipulation with the hand. It was originally developed to allow visually handicapped users access to computer applications. Ogi and Hirose[5] developed a multi-sensory data environment. It is interesting to note that although they experimented with a wind sensation, they actually mounted small fans around the user's hands,[5] again restricting hand use.

To address the problem, we have developed VisPad, a non-intrusive large area force feedback device. The prototype is comprised of a massage chair pad with eight individually computer-controlled motors that simply attaches to the user's normal office chair. The prototype is able to control any variable voltage device. The device can be used with any software program. We demonstrated the prototype at the IEEE Information Visualization 2001 conference with the USGS earthquake data set. At that time, the device's usefulness had not been tested. The device has since been shown to be very effective during our recent usability study.

To test the effectiveness of VisPad, it was attached to our program ProtAlign,[6] a multi-dimensional tool useful in determining the structure of an unknown protein. VisPad enabled us to visualize haptically physical properties within a protein structure. We were able to determine that VisPad significantly decreased the time required to assess positions along a protein structure-sequence alignment.

VisPad can be used as the sole haptics device or can also be used in conjunction with other haptics devices if the user is so inclined.

Background

Haptic rendering translates computer information through the use of force feedback devices that convey movements or sensations felt by the user. Haptic feedback is a relatively new modality for human computer interaction.[3] In this emerging field, most of the research has focused on the use of haptics as an enhancement to virtual reality (VR) systems or as an alternative to visual displays for people who are visually impaired. Instead of considering haptics as a replacement for visual displays, we consider it as an additional mode for helping a user to understand information more quickly.

Multi-modal visualization techniques, along with enabling more channels for information, provide scientists with more appropriate models for information output. The hypothesis of modality appropriateness proposes that the sensory modalities are designed to process different types of information effectively. When there are bi-modal situations involving discrepant multi-modal inputs, the modality that is most reliable for the task is weighted most heavily.[7] The goal of scientific visualization is to improve the user's comprehension of information, which can be facilitated by the improvement of human–computer interaction. When we find better methods of information display, we increase the amount of information that a user can comprehend in a smaller amount of time.[8]

Multi-sensory information visualization gives the scientist more possible sensory input channels. With more input channels, natural mappings become easier to realize, and more information can be presented at once. Of course, with multi-sensory visualization comes the task of presenting the information in ways that do not confuse the user. Sound, touch, and the visual channels cannot compete for the user's attention, but should be complementary. The inappropriate use of these sensations can be counterproductive.[5]

Multi-sensory input mapping presents challenges. Not all input senses have the same input bandwidth. Vision has by far the highest information bandwidth, and sonification also has a fairly high bandwidth, while the capacity for comprehension of haptic stimuli is much lower.[8]

When considering haptic mappings, it is important to realize some of the limitations of mapping touch. Not all skin has the same haptic sensitivity. Sensitivity varies throughout the body. The two-point threshold for skin is the distance between two points where a subject is able to determine that two distinct points of contact are being made to the skin. Two-point thresholds have been directly correlated to skin sensitivity. Different people have different sensitivity levels. With sensitivity to light touches, females are more sensitive than males. Even temperature can cause skin sensitivity to vary. Cooled skin is less sensitive than warm skin. Age can also adversely affect sensitivity.[9]

There are several types of haptic stimulation. Pneumatic stimulation can be achieved using tiny air jets. Recent research has used air jets or air pockets that push the lining of a glove against the user's skin. Problems are caused by a numbing effect on the skin, which means that the ability to sense the air is temporarily disabled. This is a low bandwidth haptic method; the air jets must be well spaced for the user to sense each of them separately.[10]

Vibrotactile stimulation generates vibrations at different frequencies and amplitudes that are used to stimulate the user. A haptic device, employing vibrotactile stimulations, seems to be the best because it can be built very small and lightweight, and can achieve a high bandwidth. The 'Tacticon 1600' is an electrotactile system that is used to replace sound for the hearing impaired. It actually attaches small electrodes to the user's fingers and provides electrical pulses. These electrodes represent the frequency of sound waves. The range of available amplitudes is limited, however, because the intensity window between the sensory threshold and the pain threshold is rather narrow.

Another approach is functional neuromuscular stimulation (FMS), which provides stimulation directly to the neuromuscular system. This type of system places electrodes directly in muscles and in the nervous system. It is invasive and can be painful. Although very interesting from a theoretical standpoint, this method is definitely not appropriate for the standard user.[10]

Much of the haptic research is applied in virtual reality systems. Haptic output displays can be broadly grouped into two main categories: those used in immersive and non-immersive environments. The immersive systems tend to be for very specialized applications, such as surgical trainers, pilot trainers, etc. Immersive displays are useful for virtual reality, but are not really necessary or cost effective for the scientific interpretation of varied data sets.

Medical VR is interested in visualization of and interaction with anatomical data. It is currently used to train surgeons, and to plan, and perform surgeries and therapy. Other uses include evaluating MRIs, ultrasound scans, physiological images, etc. Increasingly, as a result of interactions, 'realizations' in modalities other than vision have started to surface: sound, force, touch, and smell. While this is very encouraging, most of the equipment is still quite bulky, specialized, and expensive.

We are interested in non-immersive haptic displays that are non-obtrusive. Computer users still expect to have a keyboard, a mouse, a monitor and a comfortable chair to sit in as they evaluate scientific data. Our sensory tactile display is currently an off-the-shelf massage chair pad that can be put on any existing chair. However, this display method is not limited to vibrating motors. Other possibilities include connecting the controllers to fans, lights, heating elements, etc.

The Haptic Interface Research Laboratory (HIRL) and the Human Design research group at the MIT Media Lab have been developing *perceptual user interfaces* (PUIs). PUIs enable an intelligent environment to sense the computer user and use machine perception and reasoning to react.[11] They have extensive research with 'smart' technologies. These technologies use sensors that allow a computer to see, hear, and interpret a user's actions. The computer tries to anticipate a user's need so that it can react intelligently.[12] Their labs have implemented *smart rooms, smart clothes, and the Smart Chair.*[13]

One of the most interesting PUI devices is the *Sensing Chair.*[14] This device senses the body posture of the occupant. The main purpose of their research is to transform ordinary chairs into perceptual and multi-modal human–computer interfaces. The ultimate goal is to have a haptic device that senses changes in body position and orientation to drive real-time applications.[15] Although the Sensing Chair is primarily an input device, recently the chair has incorporated feedback capabilities. This feedback focuses on the transmission of directional and simple geometric information on a user's back. This direct mapping is intuitive and can be presented in the local coordinates of the user's body.[16]

Recently, an office chair with a 3×3 array of vibrotactile units has been developed at the George Washington University (GWU). Initial experiments indicate that users are able to distinguish between different vibrational intensities. The users are also able to determine the location of the vibration.[17,18] In the most recent experiments, the goal was to select the correct letter in a field of letters. The chair used tactile cues to draw the user's attention to different areas on the screen. Visual cues consisted of briefly outlining the letter at the beginning of each experimental trial.[19]

VisPad was developed exclusively as an output device. Like the GWU haptic device, VisPad can draw the user's attention to certain areas on the screen. When rotating a protein molecule using ProtAlign, the change in the location of vibration can be useful in drawing the user's attention back to the amino-acid pair being evaluated. However, our primary goal is to convey multi-dimensional information to the user. VisPad uses vibrotactile output to represent one of the physical properties of an amino-acid position in a protein. ProtAlign determines the 'Exposure' of an amino-acid position to the surrounding substrate. VisPad transmits this 'Exposure' information to the user. VisPad used in conjunction with graphical visualization allows the user to visualize more information at the same time.

VisPad

The VisPad prototype was built using a common massage chair pad by Homedics. The control wires in the Homedics pad were cut and attached to an ACCES RDAG-128H by I/O Products, Inc. The RDAG-128H uses serial communication to set the power levels in eight digital to analog converters (DACS). Each of these DACS is capable of varying the voltage level on a wire from 0 to 10 V. The higher the voltage, the more intense the vibrations of the motor in the massage chair pad. We can independently control the intensity of vibration for eight motors. The DACS could also be used to control the speed of a fan, the amount of heat in a heating element, etc. VisPad can be controlled by any computer with the ability to write to the serial port. Currently, the system is being used with an 833 MHz Dell Inspiron 8000 (see Figure 1) running Windows 2000. Presently, serial communication with the RDAG-128H is implemented using Windows Messages. Unfortunately, this causes a noticeable delay before output voltage causes a motor to vibrate.

We have written several utility functions that allow efficient communication with the RDAG-128H. This allows for several interesting effects for each of the motors such as: pulsing, on/off, and waving (start slowly, build to a specified intensity and then return slowly to off).

Use of the VisPad

Our research has focused on the investigation of natural visualization techniques. Natural visualization can be a direct analogy or an easily understood mapping that requires little training. VisPad can be used for natural visualization techniques or it can be used for more abstract mappings. It can be used alone or in conjunction with other haptic devices. There are two main goals in haptics: motor control and perception.[20] Motor control allows the user to modify information. VisPad has been

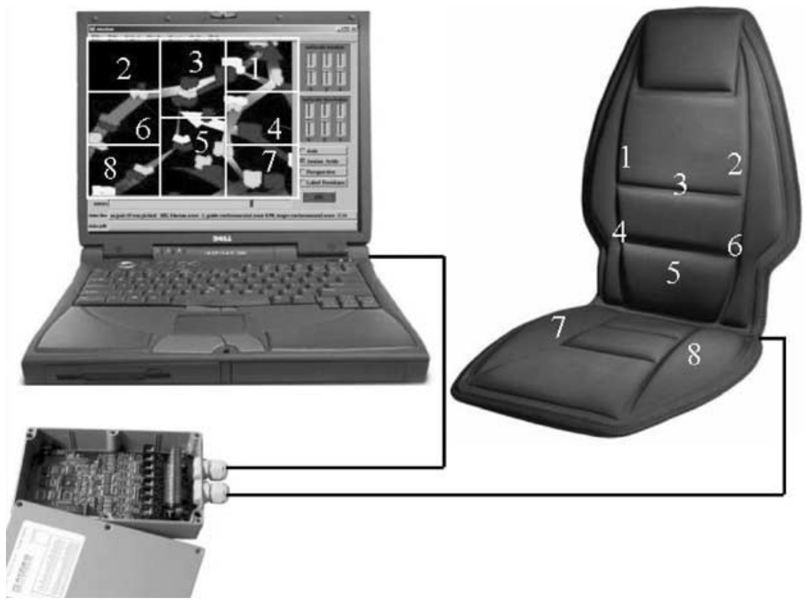

Figure 1 VisPad prototype and mapping the screen to the location of the eight motors.

designed for perceptual tasks, allowing the user to learn the properties of presented information.

During the IEEE Visualization 2001 conference, we demonstrated the capabilities of VisPad with a direct analogy to earthquake data. We used multi-modal techniques to visualize the USGS earthquake data set.[21] Earthquake date, time, location, magnitude, and depth were visualized.

A graphical slider visualized the dates and times of the quakes.

The coordinates of the states of California and Nevada were used to visualize a map graphically. The depth of the earthquake was represented by a colored dot at the location of the quake, while the size of the dot indicated the magnitude using the Richter scale. The magnitude of the earthquake and the location were mapped to the motors in the haptic device. Location was indicated by mapping the screen coordinates directly back to VisPad, while the magnitude of the earthquake was mapped to the motor intensity. Like the magnitude of a quake, the intensity was based on an exponential function.

$$\text{Voltage applied to Motor} = (e^x)/50. \qquad (1)$$

Our earthquake data set had magnitude values between 3.2 and 6.2. We wanted an exponential function that would map these values to the 0–10 V range we had available. We experimented with a few methods of mapping the voltages and found that the function in Eq. (1), where x is the Richter scale value, gave a feeling of earthquake magnitude that could be mapped to motor vibration. An earthquake of magnitude 4.0 causes the motors to vibrate with 1.09 V, while an earthquake of magnitude 4.1 causes them to vibrate with 1.21 V, for a difference of 0.12 V. This is compared to a 0.85 V difference between a magnitude 6.0 and 6.1 earthquake.

Generating haptic effects with VisPad

We have implemented library functions to allow easier access to VisPad's capabilities. At the layer closest to the hardware, there are five functions that allow access to the VisPad: *RunMotor*, *RunMotorAtLevel*, *Wave*, *WaveRow*, and *Pulse*.

The function *RunMotor*(*int motorNumber, int level*) allows the user to turn on any motor at any level from 0 to 10 V. The motor will continue to run at the assigned level until a new function changes the voltage. The RDAG-128H allows us to step up the voltage in 65535 increments (0000 to FFFF). This means there is a lot of flexibility in values. However realistically with the VisPad prototype, it is not possible to feel truly the differences between all 65536 levels. As a result, we have written functions that allow access to the motors in a more intuitive fashion.

The function *RunMotorAtLevel* (*int motorNumber, int level, int numberLevels*) scales the 65535 levels to *numberLevels* and sets the voltage to a *level* that should fall between 0 and *numberLevels*. For example, if the user chose to allow five levels, then level 1 would cause the motor to vibrate with 2 V, level 2 would be vibrate the motor with 4 V, level 3 with 6 V, level 4 with 8 V, and level 5 with 10 V. As in *RunMotor*, the motor will remain running until a call is made to set the level to 0.

The function *Pulse* (*int motor, int level, int numberLevels*) behaves like *RunMotorAtLevel*, except it turns the motor off after a small delay.

Wave (*int motor, int numberSteps, int level, int numberLevels*) allows the user to produce a triangle hat wave pattern for a specific motor. This function causes the motor to step up to the desired level in a specified number of steps and then step down in the same number of steps until the motor is turned off. For example, *Wave* (0, 1, 10, 10) will cause motor 0 to first vibrate at 5 V, then 10 V, 5 V, and then off. *Wave* (1, 2, 3, 5) will cause motor 1 to vibrate with 2 V, followed by 4 V, arriving at 6 V, then step back through 4 and 2 V before turning off.

The function *WaveRow* (*int row, int level, int numberLevels*, *char startSide* = *'L'*) causes a wave pattern to be drawn across the user's back or legs. The parameters level and *numberLevels* have the same meaning as in *RunMotorAtLevel*, *Wave* and *Pulse*. *WaveRow* builds a wave pattern across all of the motors in a single row (see Figure 2). The default direction is to start on these left and move across to the right. But the user can stipulate that a wave moves from right to left by setting *startSide* equal to '*R*'.

It is possible to combine the functions to get interesting patterns across multiple motors. For example, if we were to call the functions *WaveRow (0, 5, 5, 'L')*, *WaveRow (1, 4, 5, 'R')*, *WaveRow (2, 3, 5, 'L')*, *Pulse (0, 2, 5)*, we would get the effect of a wave of motion snaking across our back and legs and disappearing in the middle of our back.

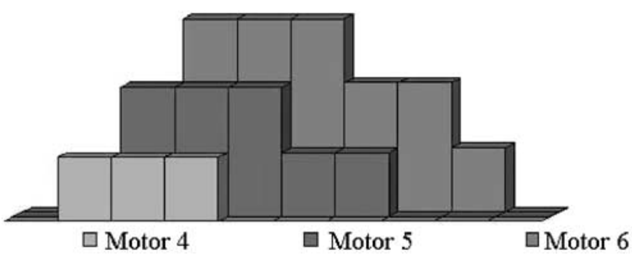

Figure 2 Wave effect of WaveRow as it travels across the back.

ProtAlign

We now give an overview of a visualization system developed for assessing the alignment between the 3D structure of a protein against another protein with a known structure.[6] We use this visualization system to run usability studies to see whether there is any improvement when a haptic interface is added.

Proteins are responsible for diverse tasks in nature. When scientists study the structure of a protein, they gain insight into its function. As genomes are sequenced, scientists have gained access to a huge number of potential proteins.[22] Currently, the human genome sequence is over 95% complete and the mouse genome sequence is over 90% complete.[23] Scientists today can determine the amino-acid sequence of a new protein, but they cannot guess its function without further analysis. One highly utilized analysis tool is the generation of a sequence alignment of the unknown protein to a protein of known structure.

ProtAlign (see Figure 3) was developed to address the problem of determining the structure of a protein from its amino-acid sequence. It is a 3D protein alignment assessment tool with multi-dimensional scoring algorithms developed to help predict the structure of unknown proteins.[24,6]

If a protein shares a 25% homology with a known protein, then it is very likely that both proteins fold in

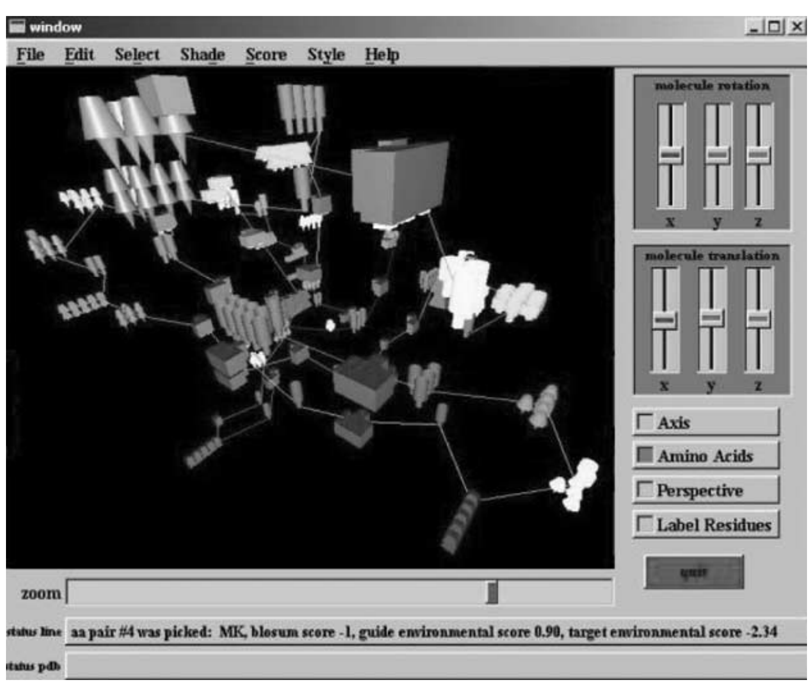

Figure 3 Screenshot of ProtAlign.

much the same way. Protein shape has much to do with the function of the protein. Scientists generate two-dimensional alignments using tools such as SAM-Sequence Alignment Modeling.[25] That is to say, the sequence of the unknown protein is lined up with the sequence of the known protein. After this alignment is made, ProtAlign allows the scientist to determine whether or not the generated alignment makes physical sense in three dimensions. ProtAlign also allows the scientist to edit the alignment in three dimensions and immediately visualize the results of the change.[6]

We have previously demonstrated methods for graphically visualizing the comparison data, but we found that the multi-dimensional data were difficult to understand quickly. Several screens were required to assess positions in the alignment fully.

Graphically, ProtAlign uses glyphs shaped like children's blocks to represent size and types of amino acids. The proportions of the blocks mimic the overall structure of the amino acid and the shape of the pegs reflects the amino acid type (see Figure 4). Each position represents a pair of amino acids in a protein structure-sequence alignment (see Figure 5). Cartoon shapes along the backbone show whether the amino acid positions are in a helix, a sheet, or a loop region (the secondary structure of the protein).[6,24]

Color is used to indicate the probability of an amino-acid substitution occurring naturally in nature as determined by the *Blosum62* matrix.[26] Color also indicates the 'Environment' and 'Exposure' scores as determined by the *Environments* program.[27] *Environments* looks at positions along the alignment, assesses the secondary structure, examines the neighboring amino acids, determines the secondary protein structure, and calculates the

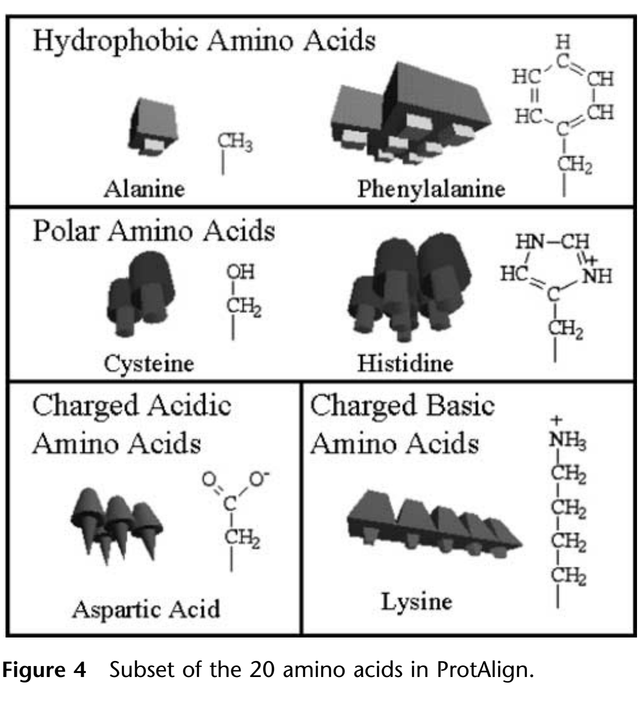

Figure 4 Subset of the 20 amino acids in ProtAlign.

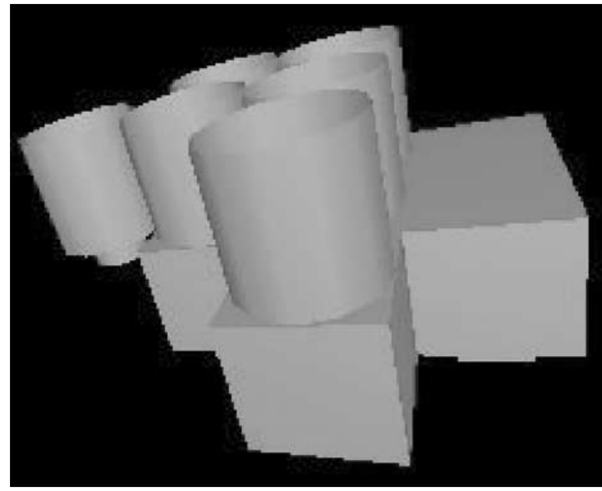

Figure 5 Pair of amino acids: histidine on top of phenylalanine.

exposure to substrate at each of the positions along the protein structure-sequence alignment. The 'Environment' metric is used to assess how much each of the amino acids in a position like the environment at that position. The 'Exposure' metric reveals the exposure to outside substrates at an amino-acid position.

Use of the VisPad with ProtAlign

The Exposure metric was also mapped to VisPad. When an amino-acid position is selected from the screen, the position is mapped back to the motors of VisPad (see Figure 1). The vibrational level represents the exposure level, with high vibration indicating a highly exposed amino acid position. ProtAlign allows full viewpoint control of the protein molecule. As the molecule is manipulated, the location of the vibration may move indicating the new position of a selected amino-acid pair. For example, if a user selects the amino-acid pair in the lower left hand corner of quadrant 8 in Figure 1, motor 8 (the lower left thigh) would indicate the exposure level. If the user then rotates the molecule moving the amino-acid position to quadrant 3, motor 3 (the upper middle back) would vibrate instead.

Because there are only eight motors in the prototype, we currently allow only two amino-acid positions to be haptically visualized at once. This greatly decreases the probability that the same motor will be needed to visualize more than one amino-acid position at the same time.

We also experimented with mapping the ProtAlign exposure metric to VisPad without the positional information. However, in this case only one amino-acid pair at a time could be haptically visualized. The exposure metric was mapped to the number of motors and the level of vibration. The higher the exposure score, the more an amino-acid position can move around, therefore we added more motors.

When an amino-acid position was buried, we applied 5 V to motor 3. This caused a mild vibration to the upper middle back. When a position was partially exposed, we applied 6 V to motors 3, 4, and 6. This caused a stronger vibrational effect on the middle and lower back. For exposed amino-acid positions, we expect more freedom of movement. Motors 1, 2, and 3 were vibrated with 10 V, while 6 V were applied to motors 7 and 8. Vibration of all of these motors indicated the freedom of movement of an amino-acid position.

Usability study

If all of the tasks were extremely easy to perform visually, there would be no need to pay attention to the haptics information. By attaching VisPad to the program ProtAlign,[6] we attempted to overload the visual system. For the usability study, we changed the user interface to remove any features that would not be used during the experiment (see Figure 6). Users still had the ability to rotate the protein molecule. However, the names of the metrics were chosen so as to reduce confusion during the experiment. The naming was as much for people who understood protein biology as those people who had absolutely no understanding of a protein molecule.

The *Blosum* metric is familiar to people knowledgeable about protein, but because people not familiar with protein would have no idea what a *Blosum* score indicates, the name of the metric was changed to reflect exactly what it is. We called it the 'Mutation Probability Index'. When each test began, this metric was mapped to the color of the backbone of the protein. The backbone was represented by the ribbon-like structure on the screen.

For the 'Environment' metric, the user was asked to compare the colors of two amino acid glyphs in a position along the alignment. This metric measures how much each amino-acid in the pair likes the environment. When each test began, this metric was mapped to the top and bottom glyphs in the amino-acid pair. The user could then compare the top glyph to the bottom glyph in the pair at a position. ProtAlign can distinguish five levels of environmental preference: much better, a little better, the same, a little worse, or much worse, although for this experiment the metric was condensed to three levels: more comfortable, the same comfort level, and less comfortable.

During an earlier usability study, we found that people took a long time to distinguish between 'much better' and 'a little better'. A user had to look at the color mapping and determine what the colors meant when

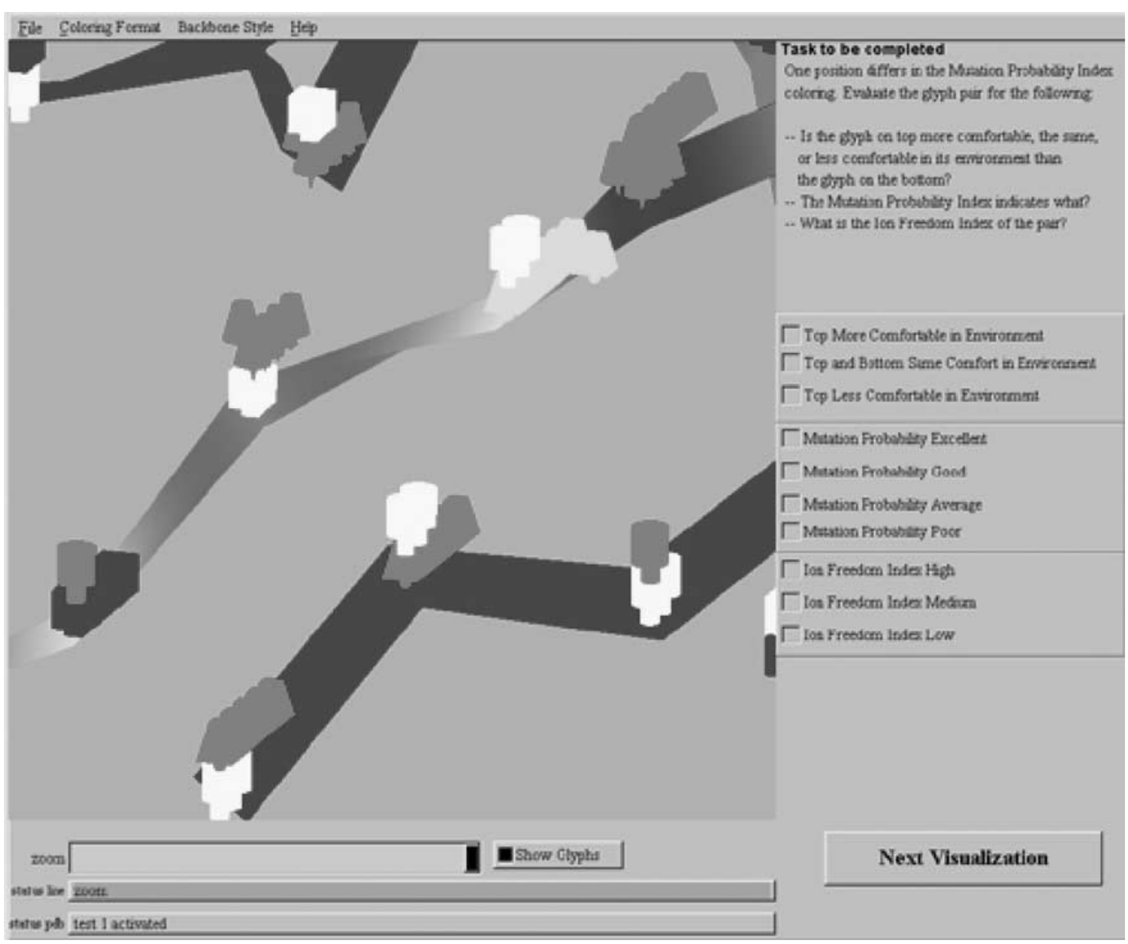

Figure 6 Screenshot of the usability test.

applied to each of the amino acids, then had to decide whether the top was more comfortable or not. Finally, they would have to make the decision on how much better or how much worse. We found that some of the color combinations were more prone to confusion. While there were very few errors determining whether the top was more comfortable or less comfortable, more errors occurred in estimating the magnitude of the difference, that is, 'much better' *vs* 'a little better'. These errors decreased when a user was trained for longer periods of time with the color mapping. Because we were limited in our experimental time, and therefore our training time, we limited the choices to 'better', 'the same', or 'worse'.

The 'Ion Freedom Index' of an amino-acid pair tells us how exposed a position is to the surrounding water molecules. There are three levels defined by the program *Environments*: exposed, partially buried, and buried. We mapped these three levels to red for exposed, green for partially buried, and blue for buried. The mapping indicates that exposed positions are relatively free to move around. Things that move tend to generate heat, thus the color red. Buried positions tend not to have much freedom. Things that do not move around get cold (blue). We needed an intermediate color to indicate partially buried positions. We chose green to keep the rainbow analogy that is pervasive to the color mapping in ProtAlign. This metric required the user to select the metric from a pull-down menu.

The intensity of haptic vibration was also mapped to the 'Ion Freedom Index' (exposure level) of the amino acid. As before, the screen coordinates were mapped to the motors in the VisPad (see Figure 1). The user used a mouse to click on the amino-acid pair that was to be evaluated in the protein alignment. The position of the selected amino-acid pair was then mapped to the motors of the haptic device, the vibrational intensity of the motor indicating the 'Ion Freedom Index'. The more the vibration, the more exposed the amino-acid pair is to the surrounding water molecules.

Participants

There were a total of 49 volunteers for the experimental process, 48 students, and one UCSC professor. No one was tested for color blindness, although none of the volunteers indicated that deciphering colors was a problem. The 48 student volunteers were all given extra credit in their class for participating in the study. All knew they were to be given the same amount of extra credit in the class for their participation regardless of their performance on the tasks.

Process

Refinining the experimental method Because the original experiment had too many tasks, they were paired down to a reasonable workload during pilot runs in Dr. Bridgeman's laboratory. In the initial experimental setup, memory of all the mappings was being tested rather than the usefulness of VisPad.[28]

The usability study was fine-tuned before running the final experiment. During the first few days, 12 of the student volunteers were used to fine-tune the experiment, so that the entire session, including debriefing questionnaire, took an hour. Pilot subjects were asked what tasks were confusing, what questions seemed ambiguous, etc. The questions and the tasks were adjusted to avoid confusion. The number of separate experiments was changed so that all of the users could finish in the allotted time of 1 h.

The final experimental method There were 36 participants in the final experiment. None of these participants had any prior exposure to ProtAlign or VisPad. They were randomly assigned to either the experimental group or the control group. There were 19 participants in the experimental group and 16 participants in the control group. There was also an equipment failure that invalidated the results from one user.

Five experimental phases

Briefing Each participant was given a written sheet to read before the test.

The briefing re-iterated that no matter how they performed, they would earn the extra credit. They were given a little information on ProtAlign and what the testing entailed (tour, training, testing, and debriefing). They were not told to perform tasks quickly, but rather they were asked to perform correctly to the best of their ability. Users were informed that during the tour and training phases, we would answer any questions they had regarding the protein visualization and the assessment methods. They were also told that during testing, we would no longer answer their questions. They were allotted 5 min for this phase.

Self guided tour A self-paced guided tour was given that showed all of the possible ways to manipulate the protein molecule. It interactively showed the user how to use all of the scoring metrics to assess a position along the alignment. Most of the tour was identical for the experimental and control groups. The only difference was that the experimental group was given a tour of the haptic mappings. The guided tour was allotted 15 min.

Training Because the control group was not exposed to the haptic vibration during the tour, they were told that the motors would give them a good massage during the training and testing phases so that the vibrations would not surprise them. The control group was not given any indication that there was information in the vibrations.

The users were given nine different training alignments. When the training alignment was presented, the Mutation Probability was mapped to the backbone and the Environmental Comfort of each of the amino acids in the pair was mapped to the glyphs.

We needed a way to draw attention to the pair that was to be assessed. We opted to generate alignments that had only one position that differed in Mutation Probability from all of the other positions. The backbone ribbon differed in color under one pair of glyphs (see Figures 7a and d). How it differed was changed in each test alignment.

Once the user determined which amino-acid pair needed to be assessed, they were allowed to use any of the three metrics to color the amino-acid pairs and the backbone. Additionally, the experimental group could click on the amino-acid pair and receive haptic feedback.

Figure 7 shows some possible evaluation windows used to assess an amino-acid position during the usability training and testing. When the window first appears, the user must determine which amino-acid position to evaluate (see Figure 7a). Suppose the backbone is red in one amino-acid position, while the rest of the backbone is yellow. Red indicates that the Mutation Probability is poor. The Environmental Comfort requires us to compare the colors of the top and bottom glyphs. Users learn that green indicates that an amino acid is comfortable in the current environment, while blue indicates that the amino acid is extremely comfortable in the current environment. If the amino acid on top is blue and the bottom is green, the user would be able to determine that the top amino acid is more comfortable in the environment.

Once the Mutation Probability and Environmental Comfort are determined, the Ion Freedom Index can be chosen to color both the backbone and the glyph pair (see Figure 7b). If the glyph pair and the backbone were blue, the user would be able to determine that the Ion Freedom Index is low.

If the user wants to re-evaluate Environmental Comfort, it is possible to color the glyph pair and have the

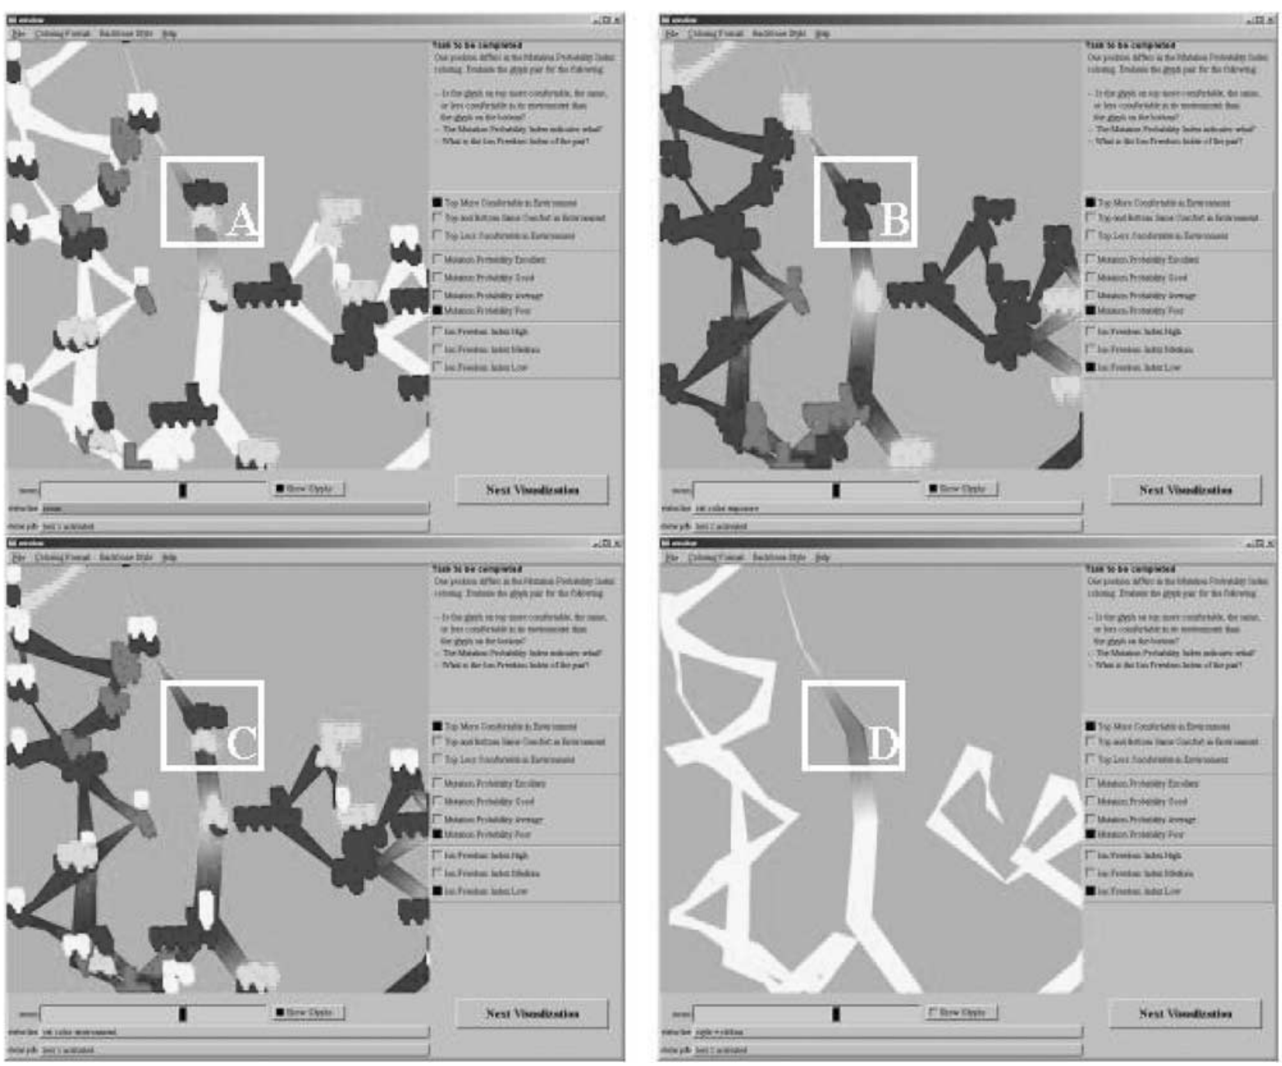

Figure 7 Evaluating an amino-acid position during usability study.

backbone represent the Ion Freedom Index (see Figure 7c). If the user had difficulty finding the position to evaluate at the beginning of the test, it is possible to turn off the glyphs highlighting the Mutation Probability Index (see Figure 7d).

The Environmental Comfort Level and the Ion Freedom Index were determined using the Environments program. The Mutation Probability was determined using the Blosum62 matrix. This metric colors the backbone of the protein by default. However, users could choose to color the amino-acid pair using the Blosum62 matrix, using a menu option. When an amino-acid pair is selected within the protein, the exposure level of the position is mapped to the motors in the chair.

Ion Freedom Index was also mapped to the intensity of motor vibration. There were three intensities for this experiment. If the amino-acid pair position was exposed to the outside solution, the motor vibrated intensely. If the position was partially exposed, a medium vibration was given, and if the position was buried in the core of the protein, very little vibration was used.

Both the experimental and control groups were asked three questions about a specified amino-acid pair in a set of protein structure-sequence alignments Users answered the questions by clicking on the desired radio buttons. They were told that the questions were always the same (see Figure 8). Question 1 was 'Is the glyph on top more comfortable, the same, or less comfortable in its environment than the glyph on the bottom?' Question 2 asked the user to determine the mutation probability. Question 3 asked the user to assess the Ion Freedom Index.

During the training when users pressed the 'Next Visualization button' a second group of radio buttons appeared with the correct answers. They were able to compare their answers and ask any questions about why they answered a question incorrectly.

The training was allotted 15 min. Before testing began, the user was asked if they were ready for the test. If they indicated they were not, we allowed them to repeat the training session.

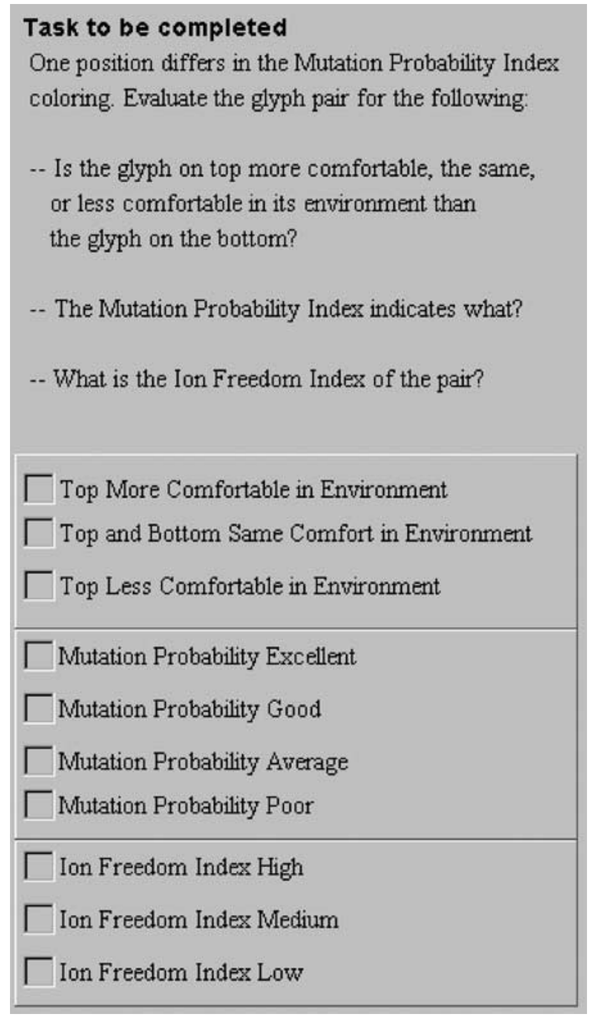

Figure 8 Usability test questions.

Testing The actual testing was allotted 20 min. There were 21 different protein alignments given in random order. The order of the experiments was determined using a random number generator that was seeded with unique numbers. Because we knew that people would become more comfortable with the testing process as it continued, the tests were given in random order to remove the effects of extra learning during the testing process.

The time it took to assess each alignment fully was recorded. The answers users gave were recorded, as were their menu selections. This enabled us to determine whether the answers were correct, how the user was assessing the amino-acid position, and to know how long it took to perform all of the tasks for that position.

The only difference between the control and the experimental groups is that the experimental group understood the meaning of the vibrational motors. Both groups were asked what they thought the motors meant on the debriefing sheet. This was to ensure that none of the control group subjects accidentally figured out the haptic mapping. Some of them figured out the positional information, but none of them discovered the link to exposure.

Debriefing The final phase was the debriefing, which was allotted 5 min. Participants were asked to answer questions about themselves and the test. We asked them if the vibration from the haptic chair pad device distracted them from the task. We also asked them about their declared major, what they considered easy, and what they considered confusing. Lastly, they were asked to rate the difficulty and enjoyability of the test.

Experimental results

Overall For each test, we verified that answers were correct, and if the answers were all correct, we recorded the time. For each test, using a two-sample *t*-test, we

Table 1 **Average time (in seconds) for users to complete all tasks for each of the 21 experiments**

Alignment test #	*Experimental group*	*Control group*
1	17.80	32.99
2	18.67	40.85
3	25.06	39.07
4	21.50	40.43
5	19.94	39.71
6	19.25	32.92
7	21.90	33.62
8	15.05	44.86
9	14.50	33.40
10	17.90	31.00
11	16.22	27.73
12	18.29	34.88
13	17.41	34.88
14	15.65	26.57
15	16.06	31.53
16	15.39	24.33
17	16.94	32.58
18	16.00	24.43
19	20.11	32.79
20	16.06	25.69
21	14.89	25.79

compared the average times of the experimental and control groups. We were able to determine that the difference between the mean times for the experimental and control groups was significant for each of the tests ($P<0.03$).

We were expecting to find that the group with haptic feedback performed their tasks more quickly. We did not anticipate how much faster they were able to complete the tasks (see Table 1). The overall average time to perform all tasks in a test for the experimental group was 17.30 s, *vs* the 31.09 s for the control group. In all of the tests, the group with haptic information performed at least 1.5 times faster than the group without haptic information.

Error rates We were concerned that the control group might find the motors distracting to the point that it would cause extra errors. This turned out not to be the case. When we compared the error rates of the experimental and control groups, the experimental group had an error rate of 11.03% compared to the control group's error rate of 11.31%. Analysis showed that the difference was not statistically significant.

Results of the debriefing questionnaire During the debriefing, we asked the users their declared major and whether they were familiar with protein biology. Because the bulk of the test pool was taken from a basic computer-engineering course, we did not expect expert computer users. Only two of the users were computer-engineering majors, and only one of them rated the test as easy. Both of these users were in the experimental group. The only biology major was in the control group. Of the experimental group, 16% had a little knowledge about protein biology as compared to 25% of the control group. Test subjects reported varied majors: Business Management, Economics, Sociology/Psychology, History, Film and Music.

We looked at what the users found confusing, easy, and/or difficult. Both the control and experimental groups considered color mapping to be the most difficult task. Many of the control group considered determining the Ion Freedom Index the most confusing task, while none of the experimental group considered it the most confusing task. However, 16% of the experimental group reported that the most difficult task was differentiating between low and medium Ion Freedom Indices. The Ion Freedom Index was considered the easiest task by 38% of the experimental group compared to 17% of the control. Comparing the debriefing answers of the control and experimental groups, we observe that VisPad appears to make the evaluation of the Ion Freedom Index less confusing.

We asked the users whether they considered the vibrational motors distracting. Of the control group, the motors at some point distracted 39% during the test compared to 43% of the experimental group, a difference that was not statistically significant.

Haptics as a possible duplicate information channel We were able to determine whether a user in the experimental group decided to combine haptic and graphical information to assess the exposure of a position in the protein alignment. They were taught that this was possible during the tour and the training. During training, the users were taught all of the possible methods of assessing a position in the alignment. However, they were not told they had to use a specific method. In fact, we told the experimental group that it was possible to get all of the information visually. We were interested in whether people would choose to use haptic feedback when given the option of determining the information visually.

Most of the experimental group chose not to visualize the exposure information graphically, but for some of the experiments, a few of the users preferred bi-modal visualization. They chose to both feel and see the exposure information to assess a position. Presumably, haptics was used to reinforce the visual information.

Using the *t*-test, we compared the times of the control group to the experimental group users who preferred bi-modal visualization. Although the results look promising ($P<0.06$), only two people consistently chose to both visualize and feel the information. Three others used bi-modal visualization sporadically throughout the experiments. This haptic device may be useful both as its own information channel and as a duplicate information channel. Further experimentation would be required to verify this.

Conclusions

We presented a prototype of a large area passive haptic device using off-the-shelf commodity massage pads. Using this prototype, we modified an existing application – ProtAlign – in order to run some usability studies. VisPad was used to visualize haptically the exposure of an amino-acid position to the surrounding substrate. The experiments demonstrate that the use of our VisPad haptics unit during the visualization of complex data greatly decreased the time it took to understand all of the visualized information. Also interesting was the finding that VisPad decreased the time with no error rate increase. Whether the user chose to use the haptic information as its own information channel or whether they chose to use haptics as a duplicate information channel, VisPad increased their performance.

Why should use of an additional modality with a very low bandwidth improve performance on a principally visual task? One answer comes from a comparison of the amounts of time that are necessary for switching modalities in attention. Psychophysical work shows that switching attention from one modality to another requires about 50 ms, while moving visual gaze requires 200 ms for saccade preparation and execution.[29–31] Although small, the latency difference would occur frequently, leading to significant advantage for cross-modal information sources.

Future work

Ideally, our VisPad haptics prototype could be turned into a more robust output device. Because of the limitations of the RDAG-128H, only eight motors could be controlled. Also, the chair pad itself has motors designed for massage. What this means is that the motors have a very large overlapping vibrational pattern. It would be interesting to see VisPad evolve so that it uses a larger array of smaller motors. It would be nice to see them radio controlled rather than by direct connection using wires.

Users took significantly less time to interpret information when VisPad was added. However, there is also evidence that the users with the VisPad learned the tasks more quickly. During training they spent less time with the help menus. It would be interesting to see a usability study to determine whether a task could be learned more quickly when haptics was used in conjunction with graphics and sonification.

Acknowledgments

We thank Craig Abbott for his help in developing the VisPad prototype, particularly with getting the RDAG-128H to work, and his help with the experimental design. His expertise in clinical psychology helped minimize the users' confusion during the usability study. Thanks to Marc Hansen for his initial efforts with ProtAlign. Initial funding for this work came from NSF NPACI ACI-9619020.

References

1 Hansen M, Charp E, Lodha S, Weissgerber-Meads D, Pang A. *PROMUSE: a prototype system for multi-media visualization of protein structural alignments*. In: Altman R, Dunker A, Hunter L, Klein T, Lauderdale K (Eds). Pacific Symposium on BioComputing. World Scientific Publishing: London, 1999; 367–379.

2 van Veen H, van Erp J. *Tactile information presentation in the cockpit*. Haptic Human–Computer Interaction 2000 (Glasgow, UK) Springer-Verlag: Heidelberg, 2000; 174–181.

3 Burdea G. *Haptics issues in virtual environments. Computer Graphics International 2000* (Geneva, Switzerland), IEEE Computer Society Press: Los Alamitos, CA, 2000; 295–302.

4 Hayward V, Cruz-Hernandez J. *Tactile display device using distributed lateral skin stretch*. In: Proceedings of the Haptic Interfaces for Virtual Environment and Teleoperator Systems Symposium (Orlando, Florida), ASME IMECE2000, 2000; DSC-69-2: 1309–1314.

5 Ogi T, Hirose M. *Multisensory data sensualization based on human perception*. IEEE Virtual Reality Annual International Symposium 1996 (Santa Clara, U.S.A.), IEEE Computer Society Press: Los Alamitos, CA, 1996; 66–70.

6 Weissgerber-Meads D, Hansen M, Pang A. *ProtAlign: a 3-dimensional protein alignment assessment tool*. In: Altman R, Dunker A, Hunter L, Klein T, Lauderdale K (Eds). Pacific Symposium on BioComputing. World Scientific Publishing: London, 1999; 354–367.

7 Lederman S, Martin A, Tong C, Klatzky R. *Relative performance using haptic and/or touch produced auditory cues in a remote absolute texture identification task*. 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems 2003 (Los Angeles, U.S.A.) IEEE Computer Society Press: Los Alamitos, CA, 2003; 151–158.

8 Fritz J. *Haptic representation of scientific data for visually impaired or blind persons*. Ph.D. thesis, University of Delaware, 1996.

9 Sekuler R, Blake R. *Perception, 3rd edn*. McGraw-Hill: New York, 1994; 572pp.

10 Edwards G. *Performance and Usability of Force Feedback and Auditory Substitutions in a Virtual Environment Manipulation Task*. Master's thesis, 2000; Virginia Polytechnic Institute and State University.

11 Turk M, Robertson G. *Perceptual user interfaces (introduction)*. Communications of the ACM 2000; **3:** 32–34.

12 Pentland A. *Perceptual user interfaces: perceptual intelligence*. Communications of the ACM 2000; **3:** 35–44.

13 Lu I. *Augmented smart chair*. http://www-white.media.mit.edu/vismod/demos/smartchair.

14 TAN H. *Perceptual user interfaces: haptic interfaces*. Communications of the ACM 2000; **3:** 40–41.

15 Tan H, Slivovsky L, Pentland A. *A sensing chair using pressure distribution sensors*. IEEE/ASME Transactions on Mechatronics 2001; **3:** 261–268.

16 Tan H, Pentland A. *Tactual displays for wearable computing*. Personal Technologies 2001; **1:** 225–230.

17 Lindeman R, Cutler J. *Controller design for a wearable, near-field haptic display*. 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems 2003 (Los Angeles, U.S.A.) IEEE Computer Society Press: Los Alamitos, CA, 2003; 397–403.

18 Lindeman R, Yanagida Y. *Empirical studies for effective near-field haptics in virtual environments*. IEEE Virtual Reality 2003 (Los Angeles, U.S.A.), IEEE Computer Society Press: Los Alamitos, CA, 2003; 287–288.

19 Lindeman R, Yanagida Y, Sibert J, Lavine R. *Effective vibrotactile cueing in a visual search task*. Ninth IFIP TC13 International Conference on Human–Computer Interaction (Zurich, Switzerland), to appear in INTERACT 2003, 2003.

20 Kirkpatrick A, Douglas S. *Application based evaluation of haptic interfaces*. 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems 2002 (Orlando, U.S.A.), IEEE Computer Society Press: Los Alamitos, CA, 2002; 32–39.

21 NCEDC, *Northern California earthquake data center* [WWW document] http://quake.geo.berkeley.edu (accessed 1 August 2001).

22 Karolchik D, Baertsch R, Diekhans M, Furey T, Hinrichs A, Lu Y, Roskin K, Schwartz M, Sugnet C, Thomas D, Weber R, Haussler D, Kent J. *The UCSC genome browser database*. Nucleic Acids Research 2003; **1:** 51–54.

23 Haussler D. *UCSC genome bioinformatics [WWW document] http://genome.ucsc.edu/* (accessed 10 April 2003).
24 Hansen M, Weissgerber-Meads D, Pang A. *Comparative visualization of protein structure-sequence alignments*. IEEE Symposium on Information Visualization 1998 (Research Triangle Park, USA), IEEE Computer Society Press: Los Alamitos, CA, 1998; 106–110.
25 Hughey R, Karplus K, Krogh A. *SAM sequence alignment and modeling software system*. Technical Report UCSC-CRL-99-11, 2001.
26 Henikoff S, Henikoff J. *Amino acid substitution matrices from protein blocks*. Proceedings of the National Academy of Sciences 1992; **89:** 10915–10919.
27 Bowie J, Luthy R, Eisenberg D. *A method to identify protein sequences that fold into a know three–dimensional structure*. Science 1991; **253:** 164–170.
28 Huey B, Wickens C. *Workload factor*. In: Huey B, Wickens C (Eds). Workload Transition: Implications for Individual and Team Performance. National Academy Press: Washington, DC, 1993; 54–93.
29 Bridgeman B, Mayer M. *Failure to integrate visual information from successive fixations*. Bulletin of the Psychonomic Society 1983; **21:** 285–286.
30 Turatto M, Benso F, Galfano G, Umilta C. *Nonspatial attentional shifts between audition and vision*. Journal of Experimental Psychology: Human Perception and Performance 2002; **28:** 628–639.
31 Turatto M, Galfano G, Bridgeman B, Umilta C. *Space-independent modality-driven attentional capture in auditory, tactile and visual systems*. Visual Cognition 2003, in press.