Multi-modal Exploration of Small Artifacts: An Exhibition at The Gold Museum in Bogota

Pablo Figueroa[*]
Universidad de los Andes
Colombia

Pierre Boulanger[†]
University of Alberta
Canada

Eduardo Londoño[‡]
Museo del Oro
Colombia

Flavio Prieto[§]
Universidad Nacional
Colombia

Mauricio Coral[¶]
Universidad de los Andes
Colombia

Juan Borda[‖]
Universidad de los Andes
Colombia

Felipe Vega[**]
Universidad de los Andes
Colombia

Diego Restrepo[††]
Universidad de los Andes
Colombia

Abstract

We present the iterative development and initial evaluation of a multi-modal platform for interacting with precious small artifacts from the Gold Museum in Bogota. By using a commercial haptic interface, loud speakers, and stereo displays, one can allow visitors to touch, hear, and observe in stereo those precious artifacts. We use this multi-modal interface in a novel way and in a novel context in order to provide virtual replicas that can be weighed, cleaned, and explored as if they were close to a visitor's hand. This platform is currently open to the public, and some of the lessons learned are reported in terms of usability in a real-world museum application.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual reality; I.3.8 [Computer Graphics]: Applications—; I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques; J.5 [Arts and Humanities]: Fine arts—;

Keywords: Virtual Heritage, Museo del Oro Colombia, Haptics, Multimodal Booth

1 Introduction

Museums have large collections of small pieces which should be both admired by current visitors and preserved for future generations. These competing goals force Museums to enclose small pieces behind glass windows with careful lighting, which preserve an artifact but also limit the ability for visitors to have real interaction with it. Moreover, some artifacts could have specific sound and material properties, but those features are usually hidden to visitors due the current standard presentation methods.

The purpose of this work is to innovate in the field of virtual heritage by using multi-modal technologies in order to manipulate virtual proxies. The developed system offers a multi-modal experience that allows touching, hearing, and seeing in stereo exact replicas of digitized artifacts. Such multi-modal interface allows museum visitors to experience virtual artifacts in a more compelling scenario than seeing real artifacts behind a glass display.

By integrating several commercially available hardware into one hardware setup, visitors can explore virtual artifacts that resemble their real counterparts. A combination of multiple commercial technologies offers several opportunities for the creation of a multimodal viewer to be used in a museum setting: 3D scanning technologies are becoming both reliable and affordable, high resolution cameras offer good quality data at commodity prices, scan data processing software allows the creation of 3D models that are suitable for real-time interaction, programmable graphic cards provide the necessary processing power to create interactive experiences on real scanned objects , commodity haptic interfaces are getting cheaper, and stereo displays are now easy to use by a large public.

This paper is organized as follows: First, we present some related work. Then, we describe in detail our platform and its features. Later, we present the evolution of our prototype and the evaluations performed. Finally, we give conclusions and directions of future work.

2 Related Work

The field of Virtual Heritage is wide and it is subject of several initiatives [Walczak and White 2003]. In particular, there have been several virtual heritage applications that have incorporated haptics in their interface development. In the CREATE project, Christou et al. [Christou et al. 2006] show an installation that uses a CAVE environment, a tracker, and two large haptic devices to create a realistic experience of manipulating artifacts from ancient Greece. Tecchia et al. [Tecchia et al. 2007] present a multi-modal exhibition with high-end haptics and stereo display, as well as a virtual gallery on the Internet with a selection of sculptures from several Museums involved in the project. A sophisticated haptic device is used in Bergamasco et al. [Bergamasco et al. 2002] in order to explore the shape of Museum artifacts, specially sculptures. Although these devices promise to create a very real haptic experience, it is probably too costly for most Museums. Laycock et al. [Laycock et al. 2006] present a system in which a simple haptic device (Phantom Omni) has been integrated to a high-quality pre-rendered environment, mostly for navigation.

In terms of usage, Bergamasco et al. [Bergamasco et al. 2002] describe two systems, one in front of a physical artifact and one in a virtual setup. In the first one, visitors are positioned in front of a sculpture where a haptic device allows them to feel its shape. An extra display shows more information, such as the current point of contact. In the second system, visitors are in a CAVE-like environment in which they can see in stereo a virtual replica of an artifact, and interact with it. Christou et al. [Christou et al. 2006] present also a CAVE-like scenario for the exploration of large scale archeological sites. Tecchia et al. [Tecchia et al. 2007] uses a virtual setup similar to the second system presented in Bergamasco et al. [Bergamasco et al. 2002]. Brewster [Brewster 2001] describes an exhibition in the Hunterian Museum at Glasgow University targeted to partially blind or blind visitors, in which it is possible to

[*]e-mail:pfiguero@acm.org
[†]e-mail:pierreb@cs.ualberta.ca
[‡]e-mail:elondola@banrep.gov.co
[§]e-mail:faprietoo@unal.edu.co
[¶]e-mail:ma.coral126@uniandes.edu.co
[‖]e-mail:jua-bord@uniandes.edu.co
[**]e-mail:fe-vega@uniandes.edu.co
[††]e-mail:dieg-res@uniandes.edu.co

feel the edges and differences in height in some carefully selected artifacts using a haptic mouse. Mclaughlin et al. [McLaughlin et al. 2000] developed a system for collaborative, remote haptic exploration of artifacts, based on heterogeneous setups composed of a Phantom and cybergrasp devices. Petridis et al. [Petridis et al. 2006] presents a multi-modal system that uses some novel devices, such as a Space-Mouse and artifact replicas manufactured by a 3D printer.

In terms of usability studies, Burke et al. [Burke et al. 2006] have shown that additional modalities to visual feedback such as touch and sound improve performance. Butler et al. [Butler and Neave 2008] suggest that visitors might take more time looking at an exhibition when their interface involve haptic devices. Asano et al. [Asano et al. 2005] suggest high expectations on exhibition planners and visitors about haptic enabled displays of artifacts in remote museum applications.

Finally, there have been more advanced systems to capture more of the physical characteristics of an artifact, such as the work by Pai et al. [Pai et al. 2001] who measures mechanical properties at different positions, or by Corbett et al. [Corbett et al. 2007] who measures sound. Although we are not using all these complex apparatus, our work is inspired by Pai and Corbett as we try to measure more than just one modality per artifact.

3 System Description

The current version of our multimodal setup offers visitors of the Gold Museum in Bogota a unique experience, in which they can touch, hear, and see in stereo six selected pieces. It consists of a stereo display, a haptic device (Phantom Omni or Novint's Falcon), and consumer-level speakers, all mounted in a way that co-locates haptic manipulation with stereo visualization. Figure 1 shows a visitor interacting in our current setup. The aluminum frame allows to co-locate the stereo image with the space for haptic manipulation, so users see and touch the virtual artifact in the same space, although no detailed calibration of spaces have been performed between the real haptic pen and its virtual counterpart. The active stereo display requires shutter glasses, which at first was considered uncomfortable for users. However, due to its higher output quality, we preferred an active display than an autostereoscopic one.

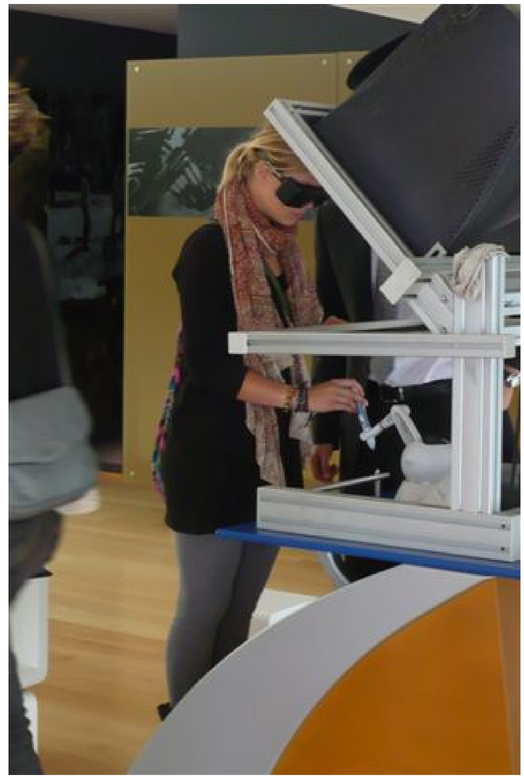

Figure 1: *The Haptic Installation of the Blue Gold Project.*

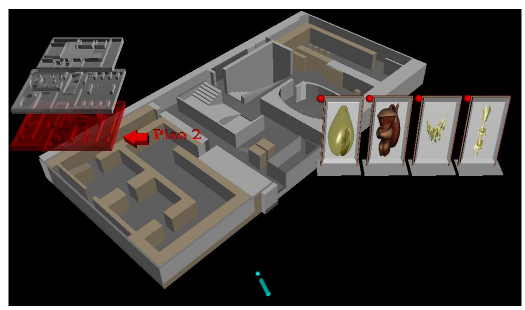

Figure 2: *Entry Screen on the Haptic Installation.*

Figure 2 shows the screen when a user enters the system. It shows small translucent versions of each floor of the Museum with the selected one in red, a big copy of the current floor, and a set of door-like widgets with the virtual replicas that could be visited. The small cylinder with a sphere in its tip represents the haptic pencil. By moving such a pencil and touching virtual objects, visitors can change the current floor, so they can see which other objects are available in this virtual exhibition, or they can push one door to see details of a particular object. A door is activated by pushing it after a certain angle, a condition that is shown by changing the color of the door's tip from red to green. Haptic feedback is active in all virtual objects, so in particular, it is possible to feel the steps in the stairs in the background representation of the floor plan. Although the main purpose of haptic feedback is to create a multimodal experience with Museum's pieces, we believe that adding such feature to all graphic elements in our interface create a rich, uniform experience that allow visitors to familiarize themselves with this new modality faster and easier.

Once a visitor enters a door, the system shows a screen such as the one illustrated at Figure 3, that permits observation of an artifact from a user selected point of view. The sticky white semi-sphere on the left is a widget for rotating the artifact. Once the virtual pen is inside, it uses the distance from its center to the pen's tip to define a direction of rotation. Visitors can stop rotation by moving the virtual pen to the center of this widget or by leaving its space. This widget allows us to avoid different states in the interface that could appear with the use of buttons in more common techniques such as dragging, so we keep the interface manipulation as simple as possible. It also allows visualization of artifacts without occlusion from the virtual pen. At this level visitors can touch the artifact and feel its surface and shape. Some selected artifacts can also play sound when hit with the virtual stylus. Options at this level allow visitors to weigh or clean an artifact, and they are represented as doors at the top, the current level with a darker background and other levels with a lighter color The system returns to the entry screen when there is no activity from visitors. In this way, we encourage a visitor to abandon the experience so other visitors can also enjoy the interaction.

Weighing (Figure 4) shows the selected artifact at one side of a simple scale. By pressing the other side with the haptic pencil, visitors can feel a relative weight of this artifact. We took the weights of all

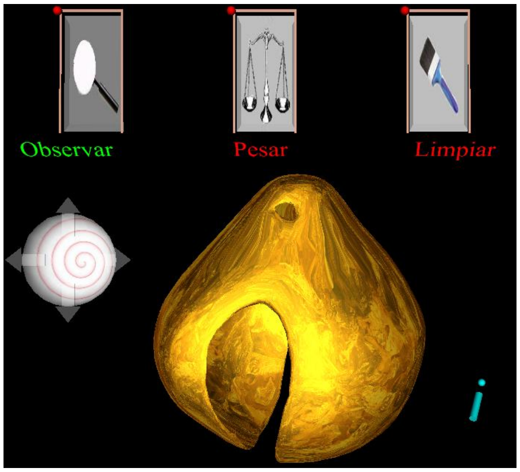

Figure 3: *Observing an Object in the Haptic Installation.*

objects in our installation and we scaled those values to the range available in the device. Although there is evidence that it is not easy for users to distinguish more than 3 feedback values in simpler haptic devices [Cholewiak et al. 2008], expert users could feel some differences in this setup beyond three levels. Artifacts can also be touched in this level as in the observing mode.

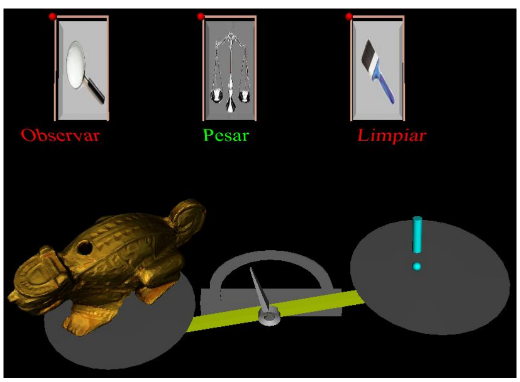

Figure 4: *Weighing an Object in the Haptic Installation.*

Finally, the cleaning mode is designed to ask visitors to look around an artifact and learn more about the cleaning process that curators perform. Depending on the artifact, we designed two cleaning procedures. In the first one, visitors can see the dirt as stained sections on the artifact. As they touch such an artifact with the haptic pencil, it gets cleaned. The second procedure shows a recipient full of cleansing liquid below an object, and controls that allow visitors to move this object up and down. The section of the object that enters the acid becomes clean. Figure 5 shows these two cleaning experiences.

3.1 Software Implementation

We chose H3D [Sensegraphics 2009], an open source API that facilitates the description of a scene, show stereo output, and facilitate interaction with haptic devices such as the Sensable's Omni and the Novint's Falcon. H3D is designed as an extension of X3D, with extra nodes for handling haptics and behavior. Our application uses a module in H3D that allows us to define the interface as a state machine, depicted in Figure 6.

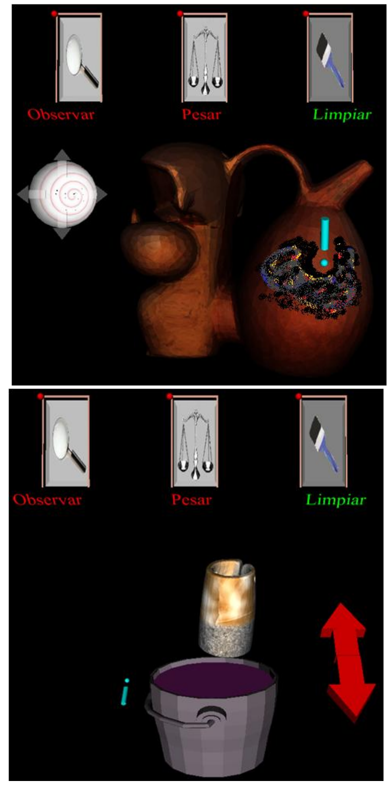

Figure 5: *Two Ways to Clean Artifacts.*

Floor selection at the main menu is implemented with the *Switch* X3D node: each time a visitor touches with the virtual pointer any of the mini maps, the system activates a Python script that identifies the selected object, clears the current visualization, and updates the scene graph with the new model. The functionality of pressing a door uses the *DynamicTransform* node . This node allows you to define torque, rotational axis, and an event that triggers a movement of each door, in this case when making contact with the haptics device.

By pushing any of the doors, the state machine changes to the Observe state. The rotation widget in this state is a scene that contains a half textured sphere and a *ProximitySensor* node that activates a *ForceField* that keeps the pencil close by. A *SpringEffect* node is added to help locate the haptic device inside the widget. Each time the proximity sensor detects movement within its field, it computes a delta of orientation for the 3D model. The three doors in this scene

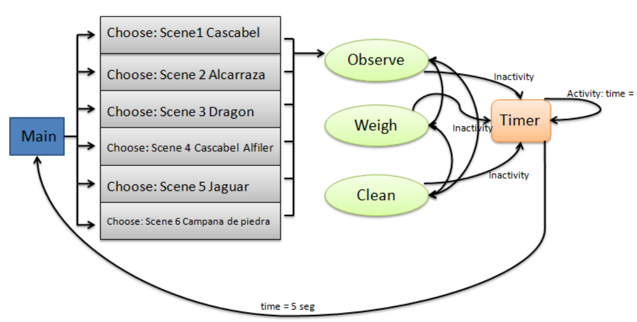

Figure 6: *State Machine for Our Application.*

are governed by the same behavior that the doors of the main scene. Objects are X3D scenes that include a *SmoothSurface* node in order to allow H3D to automatically handle haptic feedback and registration. We manually tuned the parameters of this node in order to create an effect similar to reality.

The weigh state shows a balance that was created in both Blender and X3D code, in order to facilitate the implementation of the balance movement. Each base receives an event once the haptic device makes contact with it, and this event activates a *DynamicTransform* that produces opposite force proportional to the mass defined for each object. We used $RealMass/100$ as the virtual mass for most of our objects in order to use the maximum of the available feedback range, except for the Jaguar and the Ceramic which used $RealMass/1000$, since their range of weights is way too high in comparison to other objects.

The clean state uses a custom node called *PaintableTexture* and objects with two textures: the appearance of the clean object and another one for dirt with full opacity. When users touch an object, this node receives the texture coordinates and changes the opacity level in the dirt texture to transparent. so the real appearance can be seen. The second method of cleaning in Figure 5 uses a similar custom node, *AcidTexture*, that changes opacity according to the level of the object.

We monitor user's activity with a *TimeSensor* node and a Python script that counts 5 seconds of inactivity before switching to the main state.

3.2 Interactive Content Creation

In terms of content, we captured shape and images from six objects of the Museum's collection, and sound for musical objects from this set. We used a high precision laser scanner that produced models of about 8Mb and 400.000 polygons for objects of about 7cm, but we had to reduce them to about 400Kb and 12.000 polygons, in order to allow interactive frame rates for both visual and haptics rendering. Since the laser scanner data did not include textures and we had not completed the automatic texturing process yet, we had to rely on hand texturing and fake coloring to finalize the virtual artifact models in this exhibition. In the case of smooth gold pieces it was possible to create a compelling material that looks good, such as the gold bell illustrated at Figure 7.

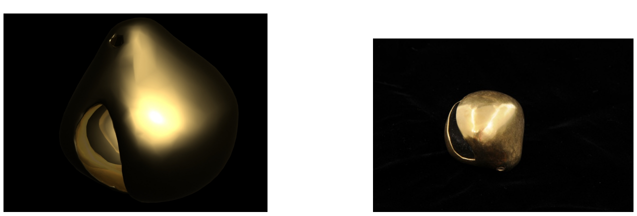

Figure 7: Comparison between a bell's replica with faked gold material and its picture.

In the case of artifacts with complex textures, it was necessary to manually stitch photographs from the artifacts as textures on the polygons. This keeps the polygon count low and allows us to show enough detail of an object. We took two sets of 36 high resolution images around two main axes of each artifact, and we used such photographs in order to create the texture map of each object. Although this process is common in the game industry, it has two main problems that need to be overcomed in the future: textures may not perfectly blend due to perspective distortions, and low polygon shapes and standard rendering techniques may show artifacts not present in real life objects, such as the apparent cracks of the ceramic in Figure 8.

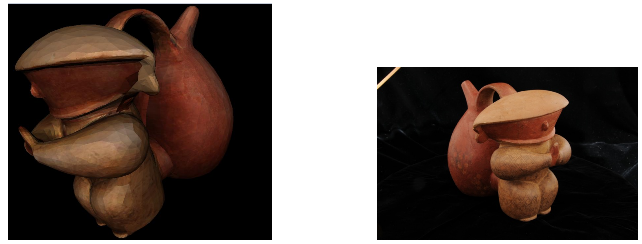

Figure 8: A model of a ceramic for interaction and its picture.

The process of capturing data can be performed by technicians with knowledge about the particular scanner and other technologies for sound and high resolution pictures, while the processes of reducing polygon count and stiching textures should be performed by personnel with knowledge in 3D design and 3D tools, such as Blender or RapidWorks. Although we have not yet created high level editors for easily adding new content to our application, we managed to create modules in H3D that can be instantiated by a XSLT-based code generation tool to the particular content we want to add[1]. We believe this also can be done by a technician with a proper training, in order to facilitate the entire content creation process.

4 User Evaluation

During development, we performed several user studies in order to understand better our interface, how visitors would react to it, and improve the interface by an iterative process. User evaluation is key in any user centered development process such as this one, since it helps to narrow the design space and the options during development by means of subject's feedback. The following subsections show the main results of each study.

4.1 Cycle 1: Small Expert Group

In this early development stage, 6 experts from our lab and the Museum were asked to perform an expert evaluation on the interface. We showed our prototype to them in two main hardware setups (with active stereo display and with auto-stereoscopic display) and they were asked to give their opinions. In general, they found the interface very compelling, they liked the weighing experience, and they were amazed by the possibilities of auto-stereoscopic displays, although they clearly took more time to adjust their position in front of the auto-stereo display and therefore it was decided to use active stereo for the public version. An issue in this first version was the lack of detail in certain pieces, since at that time no artifact had textures. We selected the two pieces with more details and we manually added textures to them.

4.2 Cycle 2: Large Expert Group

This study was previous to the opening at the Museum, and it was aimed to find significant differences in the hardware setups we were planning to use for real.

A total of 123 subjects were given a short introduction to the interface and then they were left to freely interact. After experiencing

[1] We tried to implement such modules as X3D's protos, but the current implementation of H3D prevented us from instantiating more than one copy of a Proto containing PythonScript nodes.

the interface for about 2 minutes they were asked to fill a questionnaire inspired in [Shneiderman et al. 2009], with questions about gender, age, previous experience using similar systems (yes or no), and with subjective questions in the form of 9-level Likert scales in the following issues[2]: general reactions (4.1: frustrating - excellent, 4.2: boring - exciting, 4.3: difficult - easy, 4.4: rigid - flexible), image quality (5.1: blurry - sharp, 5.2: confusing - intuitive, 5.3: inadequate - adequate, 5.4: too small - too large)[3], 3D display (6.1: inadequate - adequate), sound (6.7: inadequate - adequate), metaphor (7.1: inconsistent - consistent, 7.3: confusing - clear), feedback (7.4: not enough - enough), and learning process (8.1: difficult - easy). They could also give open comments at the end of the session.

All tested hardware setups used the Phantom Omni from Sensable as an input device and haptic rendering, and we tried several stereo displays and configurations. The conditions were the following:

1. Active-stereo monitor with shutter glasses.

2. Auto-stereoscopic screen with an inter-ocular distance of 0.01m[4].

3. Auto-stereoscopic screen with an inter-ocular distance of 0.02m.

Previously to the test, we tried other inter-ocular distance values, and we chose these ones, since our team found them significant and different enough.

In the case of condition 3, we took samples in two different places, at a conference and at our lab. For this reason we did first a between subjects analysis in both samples, in order to find significant differences between them. Figure 9 shows the averages and standard deviations for questions in the two samples.

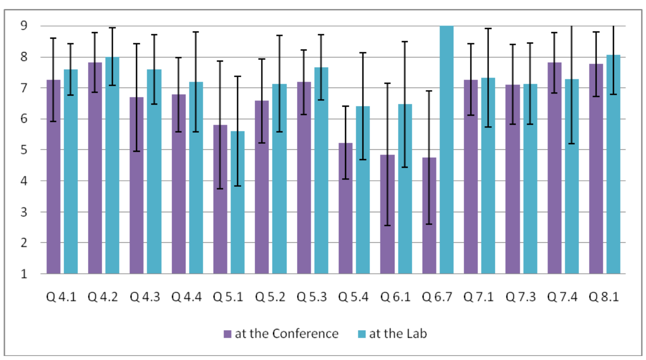

Figure 9: *Two Samples for Autostereoscopic with 0.02m, for Questions of Sec. 4.2.*

We had 27 subjects at the conference and 15 at our lab for condition 3. In order to find out differences between these two samples and due to the nature of the data, we performed a non-parametric Mann-Whitney U test. Although there are some significant differences in some answers (size in image quality and how adequate is sound), this test failed to find other significant differences. In the case of size and image quality, people at our lab considered interface elements bigger (M=6.4) than the people at the conference (M=5.23), and sound was considered better at the lab scenario (M=9) than in the conference (M=4.75). We believe these differences are due to the controlled light and sound environment we had at the lab, versus the normal light and sound at an installation during the conference, which may affect visitors' sight and hearing.

After this first analysis, we performed a between subjects Kruskal-Wallis[Wikipedia 2009] analysis on four samples: condition 1 with 15 subjects, condition 2 with 64, condition 3 at the conference with 27, and condition 3 at our lab with 15. Our hypothesis in this case is that samples are not equal, specially between conditions 1 and 2 or 2 and 3. Figure 10 shows the averages and standard deviations of these samples.

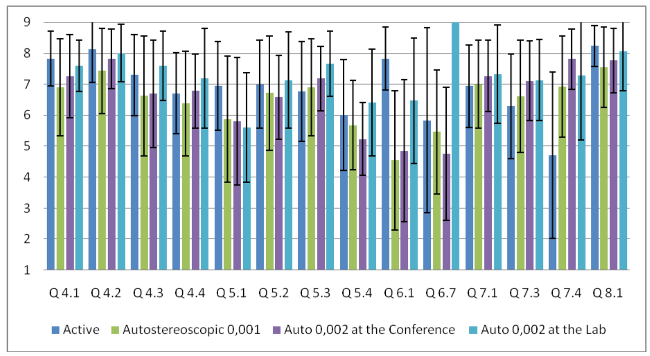

Figure 10: *Comparing Data from 4 Samples and 3 Conditions, for Questions of Sec. 4.2.*

Although the analysis apparently gives us significant differences for an alpha level of 0.05 in terms of the adequacy of the 3D display (Q. 6.1) ($p = .000$), sound (Q. 6.7) ($p = .032$), and amount of feedback (Q. 7.4) ($p = .000$), which means that subjects considered better the active than the autostereoscopic display, sound was perceived worse in the autostereoscopic setup in the lab, and they perceived less feedback in the active stereo setup than in the other ones, the shapes of the distributions of corresponding samples seem different, which invalidates these results (Figure 11). We also found from analyzing the data that although sound adequacy suggests a significant difference, this was due to a big number of lost data in condition 1, which resulted in a high mark with no standard deviation.

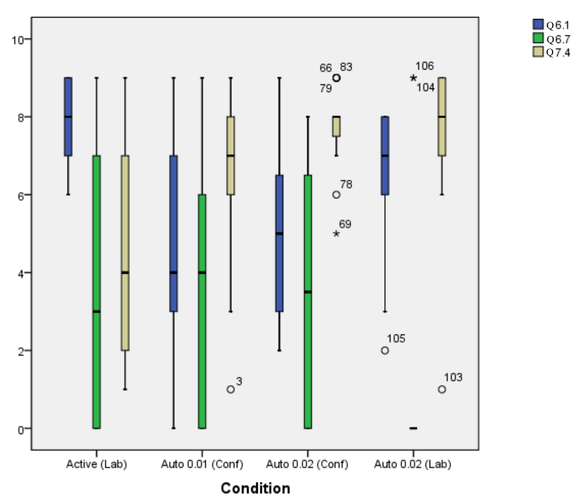

Figure 11: *Shapes of Samples from Selected Questions in Cycle 1.*

[2] Numbers associated with each scale here are used in the following figures.

[3] Readers should be aware of the difference of scale that this question has, since benign answers have a mark of 5 instead of 9 as other questions. It will be changed in future surveys.

[4] This information comes from the parameters in H3D.

In terms of average per question in all subjects, all setups received good marks. There is a tendency to mark blurrier the visualization (Q. 5.1) in an autostereoscopic display (conditions 2 and 3) than in the active display ([1: M=6.94] [2: M=5.68] [3 Conf: M=5.8] [3 Lab M=5.6]), which is interesting and support a significant result above, but our test failed to show significant differences ($p = .135$).

4.3 Cycle 3: Observation

Our multimodal installation is open to the public in the last floor of the Museum, in a place dedicated to interactive experiences and videos. It is expected that visitors will see first all real pieces and later on they will find our exhibition with virtual replicas. As a way to control the experience, we have a guide that invites people to see the exhibition, shows the way it works, and solves small technical problems if they appear. After the opening, we have performed several observation sessions at different hours, over a period of 20 days. We wanted to see how the system was used by visitors, and how useful our guide help is in the interaction between our system and visitors.

We have found that although our guide plays more as a director than as a coach, it has been necessary for her to play an active role in the interactive experience. Visitors lack in general exposure to haptic technologies, and they are either afraid or too enthusiastic about expected feedback. Adult visitors rarely approach the site by themselves, so our guide encourage them to interact, and show them what to do, what to expect, and what are the limits of the system. Kids usually go in groups and they are usually eager to try our setup, so in this case it is our guide's duty to organize them and control the use of the haptic device, since there is only one installation. Once visitors try, they are usually amazed by the technology and the virtual exhibition, although they do not express any comment related to the relationship between the real pieces and the virtual replicas. Ambient light is an issue for 3D viewing, and although we have added black covers around the environment, there are pieces that are hardly visible when ambient light is high.

4.4 Cycle 4: Hardware Changes

The purpose of the last study was to evaluate the differences of use in our latest software version in 4 different conditions, depending on the stereo displays and haptic interface available, in order to explore different form factors. We performed a between subjects analysis in which we followed a similar procedure as in Section 4.2, with the addition of two questions related to an observation task we asked subjects to perform. In this case, the conditions and number of subjects were the following:

1. Active stereo display in normal position with the Omni (10 subjects).
2. Active stereo display in normal position with the Falcon (10 subjects).
3. Auto-stereoscopic 3D display in normal position with the Omni (8 subjects).
4. Active stereo display in co-location with the Omni (11 subjects).

All conditions were performed at our lab except condition 4, which was taken at the Museum installation. Our hypothesis in this case is that conditions are different. Figure 12 shows the averages and standard deviations of the data, and Figure 13 shows the pictures of these conditions.

We performed a non-parametric, Kruskal-Wallis test on this data, with the hyphotesis that samples are different. Again, although we might think there are significant differentes for questions 5.1 (blurry-sharp image) ($p = .016$), 5.4 (too small or too large image) ($p = .012$), 6.7 (sound adequacy) ($p = .025$), 9.1 (test question 1) ($p = .000$), and 9.2 (test question 2) ($p = .016$), the shapes of their samples (Figure 14) seem different enough to invalidate these results[5].

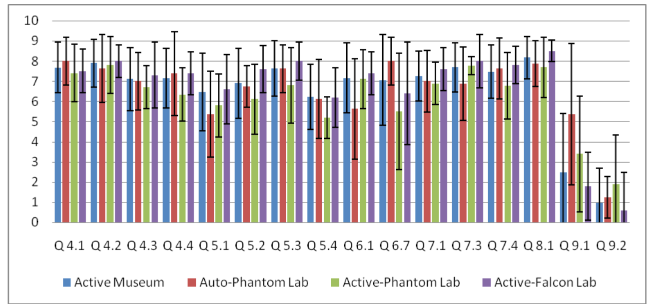

Figure 12: *Data from this Study, for Questions of Sec. 4.2.*

Figure 13: *Hardware Setups in this Study*

5 Lessons Learned

We are not aware of previous experiences in the use of commercial and relatively affordable haptic devices for interactive tasks in a public setup in the field of Virtual Heritage. We believe these devices create an opportunity for multimodal installations in Museums and similar venues, in order to complement the real encounter with pieces and offer visitors thrilling experiences. This experience showed us that multimodality and specially haptic technologies are very new to the public in general, but with proper guidance they easily overcome the learning curve and enjoy this interactive experience.

During development we also became aware of differences between peoples and venues. In terms of users, researchers attending VR conferences are very aware of available technologies and although they may be amused by novel applications, they know how to use such technologies and their limitations. In the case of visitors of a Museum, they are not used to novel technologies and they usually have to be invited and guided in order to enjoy the experience. However, once they pass the first learning barrier, the interaction technology becomes interesting and exciting, sometimes more that the actual content of the application that uses it. In particular, the

[5] In the X axis of the figure, P denotes Phantom and F denotes Falcon.

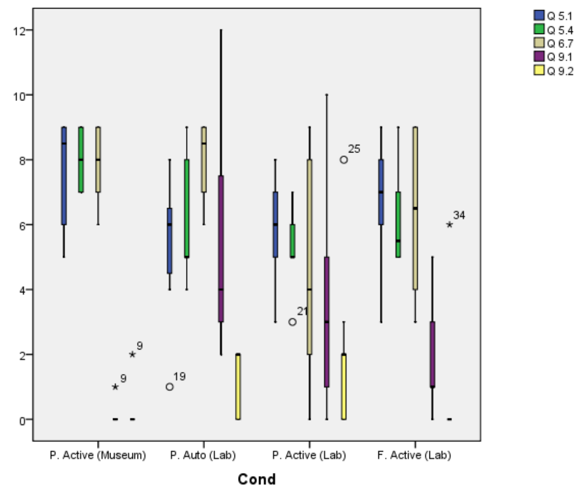

Figure 14: *Shapes of Samples from Selected Questions in Cycle 4.*

weighing experience was very interesting to all people, and they liked the effect even though the simulated weight is scaled due to limitations of the device. We still have to improve our methods of evaluation in order to measure both the impact that our application has in the general public and the sense of weight that they perceive.

The venue and the context of use are also very important in these type of experiences. Current 3D visualization technologies require controlled light environments, sound effects can be better appreciated if ambient sound is low, and the effect of small haptic devices has to be explained to visitors in order to make it more evident and enjoyable. In this direction, we found necessary to count with a guide person that mediated the experience, and facilitated the use of haptic devices and our interface. This person was also very important in the process of explaining limitations of the interface, which may be not evident to visitors. For example, force feedback of small haptic devices is limited, so users have to limit the amount of force they use during the experience. If visitors apply too much force to the device, it gets loose and it may break, so it is important to train them regarding this issue. It was also important to create a booth that allowed access to visitors of all ages and heights. Our interactive environment lies on a table that offers stairs for kids in a non-obstructive manner for adults.

Finally, user studies gave us a better understanding of how our system works, and how users approach to it. Although statistical analysis of these first studies were not statistically significant, our process of observation of these experiences and the suggestions that the data gave us were very valuable during the development and improvement of our interface. More work has to be done in order to find more effective ways to analyze user's data.

6 Conclusions and Future Work

We have shown that the use of a multimodal platform in the field of virtual heritage, in which commercial devices are integrated in order to allow visitors to see in stereo, hear, and touch replicas of small objects. Some evaluation tasks have been performed during development and after the opening at the Museum, which allow us to confirm the interest of using this technology in real-world setups, although there are open issues related to the use of these technologies by the general public. We plan to perform more studies about how visitors understand haptic feedback and how they evaluate different haptic feedback responses. We also plan to incorporate author systems in order to facilitate the inclusion of new artifacts, such as the one presented in [Wojciechowski et al. 2004]. It is also future work to find out ways to relate better the virtual replicas with originals, in order to complement in a better way real expositions.

Acknowledgments

This research was funded by Colciencias and Renata in Colombia (December 2007). We thank the AMMI Lab at the University of Alberta for lending us a high-precision scanner for this project, and to the Gold Museum for granting us access to selected pieces of their collection.

References

ASANO, T., ISHIBASHI, Y., MINEZAWA, S., AND FUJIMOTO, M. 2005. Surveys of exhibition planners and visitors about a distributed haptic museum. In *ACE '05: Proceedings of the 2005 ACM SIGCHI International Conference on Advances in computer entertainment technology*, ACM, New York, NY, USA, 246–249.

BERGAMASCO, M., FRISOLI, A., AND BARBAGLI, F. 2002. Haptics technologies and cultural heritage applications. In *CA '02: Proceedings of the Computer Animation*, IEEE Computer Society, Washington, DC, USA, 25.

BREWSTER, S. 2001. The impact of haptic 'touching' technology on cultural applications. In *In proceedings of EVA'01*, 1–14.

BURKE, J. L., PREWETT, M. S., GRAY, A. A., YANG, L., STILSON, F. R. B., COOVERT, M. D., ELLIOT, L. R., AND REDDEN, E. 2006. Comparing the effects of visual-auditory and visual-tactile feedback on user performance: a meta-analysis. In *ICMI '06: Proceedings of the 8th international conference on Multimodal interfaces*, ACM, New York, NY, USA, 108–117.

BUTLER, M., AND NEAVE, P. 2008. Object appreciation through haptic interaction. In *Hello! Where are you in the landscape of educational technology? Proceedings ascilite Melbourne 2008*, Ascilite, 133–141. http://www.ascilite.org.au/conferences/melbourne08/procs/butler-m.pdf.

CHOLEWIAK, S. A., TAN, H. Z., AND EBERT, D. S. 2008. Haptic identification of stiffness and force magnitude. In *HAPTICS '08: Proceedings of the 2008 Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, IEEE Computer Society, Washington, DC, USA, 87–91.

CHRISTOU, C., ANGUS, C., LOSCOS, C., DETTORI, A., AND ROUSSOU, M. 2006. A versatile large-scale multimodal vr system for cultural heritage visualization. In *VRST '06: Proceedings of the ACM symposium on Virtual reality software and technology*, ACM, New York, NY, USA, 133–140.

CORBETT, R., VAN DEN DOEL, K., LLOYD, J. E., AND HEIDRICH, W. 2007. Timbrefields: 3d interactive sound models for real-time audio. *Presence: Teleoper. Virtual Environ. 16*, 6, 643–654.

LAYCOCK, R. G., LAYCOCK, S. D., AND DAY, A. M. 2006. Haptic Navigation and Exploration of High Quality Pre-rendered Environments . Eurographics Association, Nicosia, Cyprus, M. Ioannides, D. Arnold, F. Niccolucci, and K. Mania, Eds., 17–24.

MCLAUGHLIN, M. L., SUKHATME, G., SHAHABI, C., HESPANHA, J., ORTEGA, A., AND MEDIONI, G. 2000. The haptic museum. In *Proceedings of the EVA 2000 Conference on Electronic Imaging and the Visual Arts*.

PAI, D. K., DOEL, K. V. D., JAMES, D. L., LANG, J., LLOYD, J. E., RICHMOND, J. L., AND YAU, S. H. 2001. Scanning physical interaction behavior of 3d objects. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, 87–96.

PETRIDIS, P., PLETINCKX, D., AND WHITE, K. M. M. 2006. The epoch multimodal interface for interacting with digital heritage artefacts. In *Interactive Technologies and Sociotechnical Systems*, Springer Berlin / Heidelberg, Berlin, 408–417.

SENSEGRAPHICS, 2009. H3d.org open source haptics. http://www.h3dapi.org.

SHNEIDERMAN, B., PLAISANT, C., COHEN, M., AND JACOBS, S. 2009. *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, 5th ed. Addison Wesley.

TECCHIA, F., RUFFALDI, E., FRISOLI, A., BERGAMASCO, M., AND CARROZZINO, M. 2007. Multimodal interaction for the web. In *Museums and the Web 2007: Proceedings*, J. Trant and D. Bearman, Eds., Archives & Museum Informatics. http://www.archimuse.com/mw2007/papers/ tecchia/tecchia.html.

WALCZAK, K., AND WHITE, M. 2003. Cultural heritage applications of virtual reality. In *Web3D '03: Proceedings of the eighth international conference on 3D Web technology*, ACM, New York, NY, USA, 182–183.

WIKIPEDIA, 2009. Kruskal-wallis one-way analysis of variance. http://en.wikipedia.org/wiki/Kruskal-Wallis_one-way_analysis_of_variance.

WOJCIECHOWSKI, R., WALCZAK, K., WHITE, M., AND CELLARY, W. 2004. Building virtual and augmented reality museum exhibitions. In *Web3D '04: Proceedings of the ninth international conference on 3D Web technology*, ACM, New York, NY, USA, 135–144.