

Interactive Mediated Reality


Raphael Grasset^α Laurence Boissieux, Jean-D. Gascuel^β Dieter Schmalstieg^γ

^αHITLAB NZ, University of Canterbury, Private Bag 4800, Christchurch, New Zealand

^βARTIS/GRAVIR, INRIA Rhone-Alpes, 655, Avenue de l'Europe, 38330 Montbonnot, France

^γ Interactive Media System Group, Vienna University of Technology, Favoritenstrasse 9-11/188/2, A-1040 Vienna, Austria

Email: Raphael.Grasset@hitlabnz.org Laurence.Boissieux, Jean-Dominique.Gascuel@imag.fr
 Dieter.Schmalstieg@ims.tuwien.ac.at



ated reality introduced by Mann (Mann 1994). Mediated reality describes the concept of filtering of our vision of the real world with and by virtual information realized with a linear or non-linear filter (in contrast augmented reality just uses an additive filter). Until now, AR systems alter the reality either by off-line placement of virtual elements in the real world, or use – like Steve Mann – hardware filtering for changing the vision of reality.




Figure 1: The interactive mediated reality concept: The user is tightly coupled with the visual filter, and controls it for altering the perception of reality.

Abstract

Mediated reality describes the concept of filtering our vision of reality, typically using a head-mounted video mixing display. We can redefine this idea in a more constructive context, applying dynamic changes to the appearance and geometry of objects in a real scene using computer graphics.

In this paper, we propose new tools for interactively mediated reality. After describing a new generic framework for achieving this goal, we present a prototype system for painting, grabbing and glueing together real and virtual elements. We also conducted an informal evaluation that provides an initial analysis of the level of interest, usability and current technical limitations of this approach.

Keywords: augmented reality, mediated reality, painting.

1 Introduction

During the last decade, augmented reality (AR) has demonstrated its potential in many application domains through a number of advantageous properties: easy collaboration, intuitive interaction, integration of digital information, and mobile computing. All AR systems, irrespective of whether they use a head mounted display (HMD) or other display system, share the goal of integrating virtual elements into the real world. A real image may be augmented with labels in screen space (Rekimoto & Nagao 1995), virtual objects placed in the world space (Feiner, MacIntyre & Seligmann 1993), or superposition of a virtual object superposed on or above a real object (Bajura, Fuchs & Ohbuchi 1992) (Poupyrev, Tan, Billingham, Kato, Regenbrecht & Tetsutani 2002). All these approaches can be subsumed into the concept of *medi-*

In this paper we explore the possibilities of “modifying reality” using computer graphics, by providing interactive tools for changing our vision of reality in real-time (figure 1). We have restricted this work to the modification of the geometric and photometric properties of real objects.

This topic is interesting because of the natural way in which computer graphics and the real world complement one another. Computer graphics can provide rapidly evolving representations of an object -with low cost in terms of energy, time, resources compared to the creation or modification of a real object - and can produce non-realistic effects (by animation or rendering). The real world allows intuitive manipulation of elements (interaction and navigation), haptic feedback, a natural environment for collaborative activities and perceptual possibilities that are not artificially limited by the capabilities of a computer. A system for modifying the real world can benefit from the advantages of both domains and therefore has enhanced expressive potential.

Mediated reality applications, such as those introduced by Raskar (Bandyopadhyay, Raskar & Fuchs 2001), are found in architecture (changing building appearance, sketching modifications), industry (adding of color/material to a scale mockup), art (visual assistance during realization of a sculpture), and also in rapid prototyping or packaging design (adding photometric properties). We will demonstrate some of these applications in the rest of this paper.

The main contribution of this work is the development of an interactive system enabling a user to paint, grab or glue both real and virtual objects and can be understood as an authoring tool for mediated

Copyright ©2005, Australian Computer Society, Inc. This paper appeared at the 6th Australasian User Interface Conference (AUIC2005), Newcastle. Conferences in Research and Practice in Information Technology, Vol. 40. M. Billingham and A. Cockburn, Eds. Reproduction for academic, not-for profit purposes permitted provided this text is included.

reality. The system is oriented towards the conceptual design domain, and consequently the main focus is on “sketching results”. We relax the physical lighting realism constraint and concentrate on aspects of the interactive system to obtain a “global idea” of the desired results. Also, we present a hierarchical description of this concept and its requirements, and an initial evaluation of different methods and solutions for modifying the reality.

2 Related Work

In a very general context, acting on reality can include several approaches:

- Using a real object for interaction with virtual content – *tangible user interfaces* (Fitzmaurice, Ishii & Buxton 1995) (Ishii & Ullmer 1997) ;
- Enhancing real object properties in the environment – *amplified reality* introduced by (Falk, Redström & Björk 1999) ;
- Changing our vision of reality: This can be done off-line, by modifying video images with computer graphics effects (e.g., broadcast special effects). Alternatively, we can interact in real-time, which leads us to *mediated reality*.

The most interesting system for modifying the appearance of real objects is Raskar’s work on shader lamps (Bandyopadhyay et al. 2001). Based on a projection approach, he realized a prototype for painting on real objects, and also provided the illusion of animation of real static objects. Unfortunately, the system is limited to dull object surfaces (constrained by the dynamic color range of the projector), suffers from bad lighting conditions during work, occlusion of projection by the user and limited mixing of real and virtual object properties. In contrast, we chose a video-see through approach, that resolves some of these problems and introduces new possibilities such as painting on ordinary objects, adding 3D material, real-time texture acquisition from video, and a technique to avoid occlusion.

Lee (Lee, Hirota & State 2001) provides a system for digitizing real objects and deforming the resulting shape. However, the system uses polygon construction point by point, and samples are restricted to symmetrical object creation. Also the demonstration was limited to a single object. Piekarski (Piekarski & Thomas 2003) introduces an outdoor system for interactive modeling, at distance, real world object (geometry and reflectance). Even the system provides a posteriori editing tools they remain limited for accurate editing of real object properties.

In recent work, Fiorentino (Fiorentino, de Amicis, Monno & Stork 2002) described a system for sketching virtual objects in reality. The system is based on a tracked HMD, and uses CAD tools (extrusion, skinning, Coons patch) for very intuitive modeling of 3D elements. Although the user can add 3D virtual elements to a real object, the system – intended for engineers – does not allow a user to act upon a real object. The user’s options are limited to manual registration of real and virtual elements, and superposition of new virtual representations of real objects.

Several pieces of work are of particular interest to us because of the intuitive metaphors used for modifying the appearance of virtual objects. Art-Nova (Foskey, Otaduy & Lin 2002) provides an interactive system with haptic feedback for painting color and texture on objects, continuing previous work in (Agrawala, Beers & Levoy 1995)(Hanrahan & Haeberli 1990). Dab (Baxter, Scheib, Lin & Manocha 2001) provides more realistic tools that simulate movement of a real brush on a canvas (Yeh,

Lien & Ouhyoung 2002). (Curtis, Anderson, Seims, Fleischer & Salesin 1997) describes diffusion of paint on a virtual surface. Noteworthy projects in the virtual reality domain are CavePainting (Keefe, Feliz, Moscovich, Laidlaw & Jr. 2001) and Schkolne’s work (Schkolne, M. & P. 2001)(Schkolne, Ishii & Schrder 2002) on creating 3D surfaces and models in real space.

Relighting real-scene and inverse lighting from real images have been investigated by Fournier (Fournier 1994), Loscos (Loscos, Drettakis & Robert 2000), and Boivin (Boivin & Galalowicz 2001) using a radiosity approach, while Debevec (Debevec 1998), Gibson (Gibson & Murta 2000), and Sato (Sato, Sato & Ikeuchi 1999) take an image-based approach. However, none of these systems provide real-time interaction.

Our approach is in the same category as (Barrett & Cheney 2002), concentrating on new techniques for object-based editing of images, but in our case for editing a 3D model. We also want to provide tools in the spirit of SketchUp3D (Sketchup3D n.d.) dedicated to sketching designs, particularly for architectural models.

With the exception of Raskar, little research is going into the development of an *interactive system* for the modification of real 3D model properties. In the rest of the paper, we first describe the concepts of such an approach and present our prototype. In the following sections we demonstrate applications and initial results. We conclude with a discussion and intended future improvements of our system.

3 Interactive Mediated Reality Metaphor

3.1 Concept

Acting on reality with a computer graphics approach can be decomposed into four phases:

1. *Acquisition of information on reality.* We need to acquire a view of the world, the geometry and reflectance of real objects, environment and lighting conditions, as well as the pose of objects. Each of these elements can be acquired automatically (e.g., using 3D vision reconstruction of geometry), or manually (e.g., digitizing point by point).
2. *Modification of virtual or real elements.* We can modify elements automatically or interactively in the real environment, based on the availability of parameters acquired previously (such as the geometry of real objects).
3. *Registration of virtual and real elements.* Previously obtained or generated virtual information must be registered and calibrated with real tools, real objects, and the real environment.
4. *Display of new vision to user.* The result of the filtering is presented to the user, either with an HMD, a desktop monitor with video camera, or a special AR display such as the Virtual Showcase (Bimber, Gatesy, Witmer, Raskar & Encarnao 2002).

In this work, we restricted ourselves to analyzing the details of the second stage. We distinguish between the following types of parameters: content to be modified, temporal references used, and spatial references used. The modified content can be a real object, or additionally a newly introduced virtual object, which according to our definition becomes a “real object” once inserted into our environment. The temporal reference defines the synchronization frequency of perception to interaction: either no synchronization, interactive, or real-time. With respect to the

spatial reference, we distinguish the following possibilities along the arguments of (Feiner et al. 1993):

- *On image.* The final image perceived by the user can be modified. The 2D spatial reference mainly depends on the display technology, be it optical overlay, video composition or light projection.
- *On environment.* The user acts on global properties of the perceived real environment by affecting lighting, adding/suppressing both real and virtual content, and modifying rendering attributes.
- *On object.* Modify properties of a real object. Of particular interest are geometry and appearance, which can be modified on a global or local level with respect to the object. Global modifications of geometry include deformation on morphing, whilst local modifications can include the addition, removal, or deformation of parts of the scene. Local modifications of appearance can include changes to color and texture on a pixel/texel/vertex level, or alternatively on a structural part of the object, like a wall or door in a architectural mock-up. Global modifications to an object may involve rendering style and materials (non-photorealism, shininess, etc.).

An ideal system suiting the above description would work without an offline acquisition step and with arbitrary, unprepared real environments. As this is currently not technically feasible, we restrict the implementation presented in this study to objects with known geometry, acquired off-line with established methods such as laser scanning, 3D vision reconstruction, or manual digitalization. While it may be interesting to work in full-scale environments, we limited ourselves to smaller scale models similar to those used by architects in their daily work. We concentrated on modifying those object properties that are of greatest interest in conceptual design.

3.2 Our Approach: The Virtual Studio



Figure 2: A real studio (left image) and our virtual studio (right image). In the right image, we see a tool palette to the right, a tracked pencil in the center, a scratch area into the left, and a physical house model on the working surface towards the top.

We based our approach on real techniques and setups chosen by artists, painters, sculptors, and designers in their everyday activities. This guided our choice of setup, which aims to reproduce an artist's studio (figure 2). In a real studio, the main ingredients are a workspace, tools, materials, and the medium (object). We adopt these elements in our system by providing a working surface, tracked tools, a tool palette, and a scratch area for sketching and experimentation. The working surface is an area where the user places the content. Tools are used to modify the content, and create new elements. The palette easily permits

access to virtual elements (color, texture, options), using simple hierarchical menus. The scratch area is used to mix different content (e.g., real or virtual paint) before using it on the medium.

Designers often start with simple elements, sketches of new ideas, and proceed to combine and reuse many different ideas (Halper, Schlechtweg & Strothotte 2002). In our system we choose the same principle based on interactive modification of a real 3D models or scenes.

Our system aims to support this workstyle. We introduce a limited number of basic operations (paint, assemble, cut, glue), complemented by convenient computer editing operations (copy, paste, undo, multi-function tools, etc.). We propose three metaphors:

- *Painting tool.* The user interactively adds color, texture, material, video, real text, etc.
- *Grab tool.* The user selectively grabs real visual information from the live video of the real world.
- *Glue tool.* The user copies and pastes content. Content can be real or virtual, including 3D models, labels, textures, etc.



Figure 3: Snapshot of the system in action. Equipped with a tracked HMD, the user interactively paints on a real tracked object.

4 Setup

4.1 Hardware Setup

The user is equipped with a tracked video see-through HMD, and acts on content placed on a working surface (figure 3). The main interaction devices are a tracked brush and tool palette. The working surfaces can be individually arranged by placing tracked cardboard palettes as desired.

The Optotrak 3020 System provides high accuracy tracking of the user's head as well as hand-held tools, and is complemented by ARToolkit vision-based tracking for a number of palettes. Video acquisition is done with an EON PAL video camera, connected to an SGI Onyx and an i-glasses HMD. The video see-through HMD gives us full control over every pixel of the user's vision.

4.2 Software Setup

The application has been developed on top of the *Studierstube* framework (Schmalstieg, Fuhrmann, Hesina, Szalavari, Encarnao, Gervautz & Purgathofer 2002), a VR/AR platform that supports multiple users, devices, and applications. *Studierstube's* rapid

prototyping environment lends itself to the development of interactive editing functions, and provides the necessary handheld palette functions (Szalavari & Gervautz 1997).

4.3 Tracking and Calibration

We need to calibrate the different elements of our workspace: Optotrack, brush, and HMD tracker are calibrated using (Tuceryan, Greer, Whitaker, Breen, Crampton, Rose & Ahlers 1995), while the camera parameters are calculated via Heikkila's algorithm (Heikkila 2000). We obtain a typical error in the order of 3-5 mm for the head, 1-3 mm for the pen and 2 mm for the camera (camera distortion is not treated). The system calibration is followed by the calibration of real objects, both static and dynamic. We based our method on (Grasset, Decoret & Gascuel 2001), and provide two interactive methods for dynamic registration and placement of real objects.

5 The Prototype System

5.1 Object Modification: Pixel Level

Users can change the appearance of a real or virtual object by painting on it with the tracked brush. This approach requires knowledge of the geometry of the real object. The painting algorithm is based on the ArtNova approach (Foskey et al. 2002) which includes: hierarchical collision detection between brush and object geometry, extraction of polygons in brush volume, modification of the object's properties based on brush properties and paint type. This approach works in texture space, avoiding the typical distortions introduced by systems working in screen space (see (Low 2001) for more details).

The painting methods available to the user when modifying particular object depend on how much information the system has concerning this object and more specifically the *type* of information.

- When the *geometry* of an object is available we use a transparency texture, combined with an occluder phantom of the real object, which is necessary to avoid seeing the interior of the object. The occluder object can be defined by an erosion operation on the real object's geometry. Depending on the painting mode, the applied color is mixed with the transparency texture to create the new appearance of the real object (figure 4).
- When the *reflectance* properties of the object are acquired on the fly texture is extracted from the video image using the current viewpoint (figure 7). We compute a projection of the brush to the image and extract the corresponding pixel color¹.



Figure 4: Painting with only geometry information. Paint is simply blended with the current color on the transparent texture of the object. Note that there is no mix with the appearance of the real object.

- When the *geometry and reflectance* properties of an object are available we define the reflectance of the object using texture and Phong material

parameters. In this case the mixing is done with the current brush and the real texture of the real object (figure 5). We observe a difference between real and virtual appearance due to the difference of real and virtual lighting. This problem can be compensated manually or automatically by adjusting virtual lights (Debevec 1998).




Figure 5: Painting with geometry and texture information. Note the improved integration with the real texture, but also the problem with different lighting conditions.




Figure 6: Painting with a "lighten" painting mode on 2D paper. We manually adjust the different lighting conditions. Note the effect of the lighten mode that can remove information with high color value.




Figure 7: Painting with texture acquisition on the fly. Shown is interaction in a pixel copy mode.

The system allows the user to paint with a variety of brush sizes, pencils, airbrushes, etc. In each case the user can not only paint with virtual color, but also with virtual texture. Texture orientation can be fixed before painting, or interactively modified by a plane dragging widget similar to (Agrawala et al. 1995). The user can grab a real patch on the fly by selecting four points on a real surface. We extract this patch texture from the user's view, unwarp

¹The user can minimize perspective error by looking perpendicularly at the grabbing area.




Figure 8: Upper left: palette with color/texture. Upper right: color and texture painting. Lower left: painting with real texture. Lower right: results of different brush types.

it and use it for the current brush (a similar approach have been proposed in (Ryokai, Marti & Ishii 2004)). This approach allows the user to interactively create a real texture by real painting, or use a texture from an alternative source such as a book or magazine. If texture information is also available for the destination surface, the user can choose from a variety of ways in which the paint should be blended with the surface: copy, add, subtract, lighten.

5.2 Object Modification: Patch Level

If manipulation of individual pixels is too ineffective, the user can select a specific patch area of an object. A patch can be predefined or created interactively by selecting individual triangles or painting a mask. Once a patch is available, operations can be efficiently applied to it as a whole, for example introducing changes to color, texture or material properties such as specularities. Patches give the user a high level of control over the object. Another powerful tool is the placement of textured labels (similar to (Kalnins, Markosian, Meier, Kowalski, Lee, Davidson, Webb, Hughes & Finkelstein 2002)). The label is a textured patch which is projected onto the texture of the destination object. We use classical methods involving planar, cylindrical or spherical projection (figure 9). While basic virtual manipulators for configuring the projection are available, we did not attempt to correct distortion. A user can also grab a label from a real object and glue it onto another object. This feature permits the incorporation of real text, photos, and hand painted images.



Figure 9: Interactive placement of labels: Left and middle – placement of a virtual door on a real house; right - placement of a window with the “label and copy” tool.

5.3 Object Modification: Material Level

The user can also modify the geometry of a real object. As shown in figure 10, the geometric additions can be chosen from a browser on the palette, then manually aligned with the destination.




Figure 10: Adding predefined elements - a cardboard cuboid is augmented with a virtual roof, porch, and a sketch label.

Material modification is complemented with sculpture tools for interactively adding material to a real object. Available shapes include cubes, ribbons, extrusions, tubes, sphere spray, etc. The material is deposited interactively at the brush location. The new matter becomes part of the scene and can be painted or modified in the same way as other scene objects. Because no full model of the background environment was available, we did not attempt material removal.

5.4 Global Modification: Adding Object

The overall scene can be extended with new objects unassociated with any existing real object. Again, additions can be predefined (polygonal models, billboards and image-based rendering models) or interactively created through material deposition. We can also add real elements to a real scene, registering them with the approach presented in section 4.3.

6 Results and Applications

In this section, we present results from our system. In all cases, measurements and screenshots are realized in real-time on an SGI Onyx with a single active processor. Our system has been tested with different shapes: polygonal, curved surfaces (like a cup), and complex surfaces (scanned mockup).

The fine-grained control over the deposit of paint mainly depends on technical limitations, namely accuracy of tracking and texture resolution. In practice, tracking permits the placement of a brush with an accuracy in order of 4-5 mm. Paint diffusion looks realistic and operates in real time. We typically need a 256x256 pixels texture for an element with a maximum size of 10cm x 10cm. Painting in larger area requires excessive computational power and often involves texture memory swapping, which adversely affects the efficiency of the method.

Little previous work, with the exception of Dynamic Shader Lamps, can be directly compared with our approach in terms of results. We noted the following:

- Advantages: It is capable of painting on objects with arbitrary surface properties, not only dull surfaces. The method is usable under any lighting condition. The user’s viewpoint and preferences are incorporated for grabbing and retrieving information. It is possible to mix 2D surface augmentation with 3D virtual content, which cannot be achieved with a projection system. Finally, our system supports a larger set

of design tools, which mostly draw from the rich legacy of 2D painting systems.

- **Disadvantages:** The video head-mounted display is intrusive and lacks stability due to head tracking deficiencies. Occlusion of the hand is less natural than with a projection system. Moreover, the user's perception is limited to the resolution of the video camera.

We now describe opportunities and limitations in several potential applications areas.




Figure 11: Results with identified applications : architecture, design and cosmetic.

Architecture: The user can paint on a real mockup to experiment with different materials for a future renovation of a building. We used a scanned model for testing our algorithm for painting on a complex object (see figure 6). Unfortunately, the limited accuracy of reconstruction sometimes yields aesthetically unsatisfactory results. We can easily paint on a door or window or any other planar surface, but the system is inefficient in dealing with very small details.

Packaging design: A designer can easily modify an existing product, or paint on a rapid prototyping model. In this case, dull objects are sufficient, and the main problem seems to be finding efficient manipulation techniques for easy placement of labels on the real surface. Our approach to sketching a real label and then gluing it onto a surface can nicely combine the advantages of traditional pencil-and-paper sketching with 3D modeling.

Cosmetic: A user can easily test a new look or style, as shown in figure 6. With a virtual representation of one's head, a user can play with make-up and accessories like glasses or jewellery. The idea of scanning and modifying a representation of oneself has been explored extensively in the context of e-commerce, but a mediated reality interface has the potential to combine unlimited digital manipulation with convenient and established interaction techniques like make-up painting.

Collaborative games: Children can easily use simple real models for a new kind of game. Some toy vendors have already shown the potential of combined real and virtual construction games (e.g., Lego Mindstorm robots). Mediated reality can bring simple and easy visual customization to such games. The constructed objects and scenes can subsequently be used for continued play; here the remaining question is what kind of "simulation engine" (e.g. car traffic in a city model) to use for play. Another challenge is to make the overall system sufficiently robust for everyday playing.

7 Evaluation and Discussion

We conducted an informal evaluation for analysis of the level of interest, usability and users' perception of our new concept with the previously described prototype. We developed a formative evaluation based on subjective user comments (opinion survey), interviews and direct observations.



Figure 12: Some user results with free experimentation.

7.1 Pilot Study

We asked several users to evaluate our system by creating a scene from a simple cardboard box. They sequentially evaluated the different possibilities of our interactive system with the aid of an assistant explaining the use of the available tools. We briefly describe their comments, resulting from the interview analysis:

- Users have the sensation that paint is really applied to the object, but with a little gap. They really appreciated this new concept and suggested interesting new applications.
- They had no problems manipulating the palette, and chose options and elements in the hierarchical menus.
- They liked the possibility of choosing real elements.
- The stylus was found to be too heavy, and the tracking to be insufficiently stable.
- The tools on tablet were found to be too small, and interpretation was difficult without assistance.
- Sometimes it is difficult to perceive depth, and the users did not intuitively paint on the real object, or sense where the palette is placed². Users proposed a snapping method or grid alignment to compensate the lack of perception.
- Misregistration of virtual and real elements sometimes caused problems in painting.

7.2 Main Study

Based on the feedback from the pilot study, we attempted to improve our system. The rigid tracker on the pencil was replaced with a more lightweight and robust system, icons and interface were simplified and registration was improved. We suppressed rotation movement of the brush and compensated by constraining placement techniques. We then conducted

²This problem seems to be connected to the principle of monoscopic video see-through and insufficient resolution, and needs to be addressed with improved hardware

two more evaluations, one with non-professionals and the other with professional designers.

Casual users. In the first study, we conducted an evaluation with 18 staff members of our institute, from diverse backgrounds. We replaced the cardboard box with a real mockup (see figure 13), and conducted the evaluation based on an interior design scenario.

As previously, different tools were presented to a user, and the user was then free to experiment. At the end a questionnaire presented, which divided into five parts: experienced sensations, setup acceptance, tool set functionality, level of interest in application, and requests for improvements. Each session took approximately 20 minutes.

We analyzed 18 questionnaires (about half of them from managers, the remainder divided between technical staff and students). Almost all users found the application interesting (16x “very interesting”, 1x “interesting”), simple to use (4x “very simple”, 10x “simple”), but few found the current application efficient for modifying the real content (1x “very efficient”, 7x “efficient” and 5x “moderately efficient”).

Overall, users appreciated the idea of virtual painting in the real world, naturally moving the head around the mockup, and the freedom of object manipulation (mockup or pencil tool).

We found that HMD ergonomics remain problematic (6x “comfortable”, 7x “moderately comfortable”, 6x “not comfortable”). The main issue appears to be a misregistration problem. The working environment (workspace, tool palette and brush) has been largely appreciated, although users mentioned a lack of space and discomfort with the wires of the tracking system.




Figure 13: The setup experiment (in left image) and a typical result (in right image): the architectural mockup and interactive device, with the result obtained by a user.

Users rated the provided tools as very interesting (72%). The painting tool was the most appreciated (88%) while the tool for adding of real elements the least appreciated (60%). Some users noticed low efficiency in comparison of real or traditional digital tools for a part of presented tools. They noticed also a lack of established options for painting systems, such as command history, lasso selection mode, zooming, etc. Although we set out to recreate a physical painting/sculpting experience, users stated that the application is more similar to computer-based painting and modeling tools.

Only 55% of users found that the painting appears to be really on the mockup surface when they turn their head around, or move the mockup. In contrast, the perception of paint diffusion seemed to be effective. At the beginning of a session, new users generally have some problems with occlusion between pencil and virtual painting, but seem to rapidly become accustomed to this within a few minutes.

We asked users to rate the perceived potential for our envisioned application areas, and received the following scores (“has potential”): architecture (77% yes), games (50% yes), art (38% yes), cosmetic (33% yes), packaging (33% yes). Some users also propose




Figure 14: Two results during the evaluation by casual users.

other interesting domains like interior design, product presentation or virtual prototyping.

Professional users. Complementary to the evaluation with casual users, we conducted another evaluation with architectural students in their final year of college (future potential users of this system). 20 students and 3 professors of Lyon Architecture School were invited for testing and discussing the system (figure 15). They largely appreciated the concept (85% rating it “very good” or “good”). In their domain, the main use would be for interactive modification, real time preview of designs, or client presentations. The main perceived limitations were scale (as mockup was small), precision, and efficiency. They suggested the system could be complementary to more traditional tools and 3D modeling computer tools.



Figure 15: An architectural student experiments with the system.

Our observations confirmed that users can generally accommodate to the HMD, easily master the different tools, and rapidly develop their imagination despite the limitations imposed by the available technology.

It was encouraging to learn that users have generally expressed encouraging interest in this concept, and found the application tools very useful. In comparison to others tools or solutions, efficiency and usability are limited by the prototypical characteristics of our system. Improvements can guide us to realize a formal evaluation (like with a comparison with projective approach). Several users judged the system in its current state to be more a curiosity rather than a professional tool. However, we believe that technological limitations can be overcome with sufficient time and effort.

We think that a fundamental drawback of our approach is mainly related to the insufficient visual properties of the video see-through system which reduces the user’s ability to perceive the world at video definition. We hope that in the future optical see-

through displays will support a high dynamic color range and be comfortable enough for casual use. Another major obstacle lies in poor registration destroying the illusion of painting with haptic feedback. One solution for improved registration may be a vision-based tracking system with a model-based approach.

8 Conclusion and Future Work

We presented a new system for interactively modifying our vision of reality. The system, based on an artist's metaphor, allows a user to paint, glue and paste elements on a real object. Users have access to a large palette of tools and can completely modify reality. We describe a framework for interacting with reality, and have built and informally evaluated a first working prototype.

In the future, we plan to perform a more in-depth evaluation with several domain specialists, such as architects, beauticians and children. Another interesting challenge is the possibility of painting on objects with unknown geometry, allowing a user to paint and simultaneously create a virtual model on the fly. Another untreated aspect is the removing of real matter from an object or removing elements from a scene. We also have plans to provide more tangible tools, indirect painting tools (spray painting), hardware supported painting with programmable shaders, and other additional features.

In conclusion, we speculate that mediated reality may be a nascent form of art.

Acknowledgements. This work was sponsored by the Austrian Science Fund FWF under contract no. Y193. We would like to thank Joe Newman for proof-reading and suggestions. Many thanks to the subjects who have tested the system.

References

- Agrawala, M., Beers, A. C. & Levoy, M. (1995), 3d painting on scanned surfaces, *in* 'Symposium on Interactive 3D Graphics'.
- Bajura, M., Fuchs, H. & Ohbuchi, R. (1992), 'Merging virtual objects with the real world: Seeing ultrasound imagery within the patient', *Computer Graphics* **26**(2), 203–210.
- Bandyopadhyay, D., Raskar, R. & Fuchs, H. (2001), Dynamic shader lamps: Painting on real objects, *in* 'The Second IEEE and ACM International Symposium on Augmented Reality (ISAR'01)'.
- Barrett, W. A. & Cheney, A. S. (2002), 'Object-based image editing', *ACM Transaction Graphics* **21**(3), 777–784.
- Baxter, W., Scheib, V., Lin, M. & Manocha, D. (2001), Dab: Interactive haptic painting with 3d virtual brushes, *in* 'Proceedings of ACM SIGGRAPH 01', pp. 461–468.
- Bimber, O., Gatesy, Witmer, Raskar, R. & Encarnao, M. (2002), 'Merging fossil specimens with computer-generated information', *IEEE Computer* **35**(9), 45–50.
- Boivin, S. & Gagalowicz, A. (2001), Image-based rendering of diffuse, specular and glossy surfaces from a single image, *in* 'Proceedings of ACM SIGGRAPH 2001'.
- Curtis, C. J., Anderson, S. E., Seims, J. E., Fleischer, K. W. & Salesin, D. H. (1997), 'Computer-generated watercolor', *Computer Graphics* **31**(Annual Conference Series), 421–430.
- Debevec, P. (1998), 'Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography', *Computer Graphics* **32**(Annual Conference Series), 189–198.
- Falk, J., Redström, J. & Björk, S. (1999), 'Amplifying reality', *Lecture Notes in Computer Science* **1707**, 274–282.
- Feiner, S., MacIntyre, B. & Seligmann, D. (1993), 'Knowledge-based augmented reality', *Communication of the ACM* **36**(7), 52–62.
- Fiorentino, M., de Amicis, R., Monno, G. & Stork, A. (2002), Spacedesign: A mixed reality workspace for aesthetic industrial design, *in* 'International Symposium on Mixed and Augmented Reality (ISMAR'02)'.
- Fitzmaurice, G. W., Ishii, H. & Buxton, W. (1995), Bricks: Laying the foundations for graspable user interfaces, *in* 'Proceedings for ACM CHI 1995', pp. 442–449.
- Foskey, M., Otaduy, M. A. & Lin, M. C. (2002), Artnova: Touch-enabled 3d model design, *in* 'IEEE Virtual Reality Conference 2002'.
- Fournier, A. (1994), Illumination problems in computer augmented reality, TR 95-35, Department of Computer Science at the University of British Columbia.
- Gibson, S. & Murta, A. (2000), Interactive rendering with real-world illumination, *in* '11th Eurographics Workshop on Rendering'.
- Grasset, R., Decoret, X. & Gascuel, J.-D. (2001), Augmented reality collaborative environment: calibration and interactive scene editing, *in* 'VRIC 2001'. Laval Virtual.
- Halper, N., Schlechtweg, S. & Strothotte, T. (2002), Creating non-photorealistic images the designer's way, *in* 'International Symposium on Non Photorealistic Animation and Rendering (NPAR2002)', pp. 97–104.
- Hanrahan, P. & Haerberli, P. (1990), Direct wysiwyg painting and texturing on 3d shapes, *in* 'Proceedings of ACM SIGGRAPH 90', pp. 215–223.
- Heikkila, J. (2000), 'Geometric camera calibration using circular control points', *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(10), 1066–1077.
- Ishii, H. & Ullmer, B. (1997), Tangible bits: Towards seamless interfaces between people, bits and atoms, *in* 'Proceedings of ACM CHI 1997', pp. 234–241.
- Kalnins, R. D., Markosian, L., Meier, B. J., Kowalski, M. A., Lee, J. C., Davidson, P. L., Webb, M., Hughes, J. F. & Finkelstein, A. (2002), Wysiwyg npr: drawing strokes directly on 3d models, *in* 'Proceedings of ACM SIGGRAPH 2002', pp. 755–762.
- Keefe, D. F., Feliz, D. A., Moscovich, T., Laidlaw, D. H. & Jr., J. J. L. (2001), Cavepainting: a fully immersive 3d artistic medium and interactive experience, *in* 'Symposium on Interactive 3D Graphics', pp. 85–93.
- Lee, J., Hirota, G. & State, A. (2001), 'Modeling real objects using video see-through augmented reality', *Presence - Teleoperators and Virtual Environments* **11**(2), 144–157.

- Loscos, C., Drettakis, G. & Robert, L. (2000), 'Interactive virtual relighting of real scenes', *IEEE Transactions on Visualization and Computer Graphics* **6**(3).
- Low, K.-L. (2001), Simulated 3d painting, TR 01-02, Department of Computer Science, University of North Carolina at Chapel Hill.
- Mann, S. (1994), 'mediated reality', TR 260, M.I.T. Media Lab Perceptual Computing Section.
- Piekarski, W. & Thomas, B. H. (2003), Interactive augmented reality techniques for construction at a distance of 3d geometry, in 'Proceedings of International Immersive Projection Technology Workshop and Eurographics Workshop on Virtual Environments', pp. 461-468.
- Poupyrev, I., Tan, D. S., Billingham, M., Kato, H., Regenbrecht, H. & Tetsutani, N. (2002), 'Developing a generic augmented-reality interface', *Computer Graphics* **35**(3), 44-50.
- Rekimoto, J. & Nagao, K. (1995), The world through the computer: Computer augmented interaction with real world environments, in 'Symposium on User Interface Software and Technology', pp. 29-36.
- Ryokai, K., Marti, S. & Ishii, H. (2004), I/o brush: Drawing with everyday objects as ink, in 'Proceedings of ACM CHI 2004', pp. 234-241.
- Sato, I., Sato, Y. & Ikeuchi, K. (1999), 'Acquiring a radiance distribution to superimpose virtual objects onto a real scene', *IEEE Transactions on Visualization and Computer Graphics* **5**(1), 1-12.
- Schkolne, S., Ishii, H. & Schrder, P. (2002), Tangible + virtual a flexible 3d interface for spatial construction applied to dna, Tr, Caltech Multi-Res Modeling Group.
- Schkolne, S., M., P. & P., S. (2001), Surface drawing: Creating organic 3d shapes with the hand and tangible tools, in 'Proceedings of CHI 2001'.
- Schmalstieg, D., Fuhrmann, A., Hesina, G., Szalavari, Z., Encarnao, L. M., Gervautz, M. & Purgathofer, W. (2002), 'The studierstube augmented reality project', *Presence - Teleoperators and Virtual Environments* **11**(1), 33-54.
- Sketchup3D (n.d.), 'www.sketchup.com/'.
- Szalavari, Z. & Gervautz, M. (1997), 'The personal interaction panel - a two-handed interface for augmented reality', *Computer Graphics Forum* **16**(3), 335-346.
- Tuceryan, M., Greer, D. S., Whitaker, R. T., Breen, D. E., Crampton, C., Rose, E. & Ahlers, K. H. (1995), 'Calibration requirements and procedures for a monitor-based augmented reality system', *IEEE Transactions on Visualization and Computer Graphics* **1**(3), 255-273.
- Yeh, J.-S., Lien, T.-Y. & Ouhyoung, M. (2002), On the effects of haptic display in brush and ink simulation for chinese painting and calligraphy, in 'Proceedings of Pacific Graphics 2002 (PG2002)', pp. 439-441.