Activity Report 2013

Project-Team REVES

Rendering and virtual environments with sound

RESEARCH CENTER
Sophia Antipolis - Méditerranée

THEME
Interaction and visualization

Table of contents

Project-Team REVES

Keywords: Audio, Computer Graphics, Interaction, Visualization, Virtual Reality

Creation of the Project-Team: 2002 July 01.

1. Members

Research Scientists
George Drettakis [Team leader, Inria, Senior Researcher, HdR]
Adrien Bousseau [Inria, Researcher]

External Collaborators
George Koulieris [Technical University of Crete]
Ares Lagae [Gouvernement étranger, until Dec 2013]

Engineers
Rachid Guerchouche [Inria, FP7 VERVE project]
Clement Riant [Inria, Autodesk Inc, from Mar 2013]

PhD Students
Gaurav Chaurasia [MESR]
Emmanuelle Chapoulie [Inria, Conseil Régional PACA]
Sylvain Duchêne [Inria, FP7 VERVE project]
Emmanuel Iarussi [Inria, ANR DRAO project]
Rodrigo Ortiz Cayon [Inria, FP7 CR-PLAY project, from Nov 2013]

Post-Doctoral Fellows
Jorge Lopez Moreno [Inria, Autodesk Inc, until Jul 2013]
Stefan Popov [Inria, ANR ALTA project]
Christian Richardt [Inria]

Administrative Assistant
Sophie Honnorat [Inria]

Others
Kritarth Anand [Inria, intern, from May 2013 until Jul 2013]
Julia Chatain [Inria, intern, from Apr 2013 until Aug 2013]
Uditha Kasthuriarachchi [Inria, intern, Autodesk Inc, from Jun 2013 until Sep 2013]
Joan Sol Roo [Inria, intern, from May 2013 until Sep 2013]
Arunim Samat [Inria, intern, from Jun 2013 until Aug 2013]

2. Overall Objectives

2.1. General Presentation

Images, often accompanied by sound effects, have become increasingly present in our everyday lives; this has resulted in greater needs for content creation. Despite the fact that many traditional means exist, such as photography, artistic graphic design, audio mixing, they typically still remain the reserve of the expert, and require significant investment in time and expertise. Our main interest is computer image and sound synthesis, with an emphasis on automated methods. Our main goals include the simplification of the tasks required for the production of sound and images, as well as the development of new techniques for their generation.

The application domain is vast. It ranges from audiovisual production, which typically requires long, offline computation to obtain high quality results, all the way to real-time applications such as computer games or virtual reality, for which the main consideration is to guarantee 60 frames per second frame rates, or, in general the reduction of latency to user reaction. The process of generation of images and sound, generally called *rendering* is our primary interest; our second main interest are virtual environments (VE's) as well as augmented (AE's) or mixed environments (ME's), that is scenes containing both real objects (often digitized) as well as purely synthetic objects. We are interested in both the generation and the interaction with these environments. We use the term virtual environments for scenes with a certain degree of interactivity, potentially in a semi-immersive (stereo and tracking, workbench) or immersive (CAVE, RealityCenter) context.

2.2. Highlights of the Year

The past year was highly productive, with a large number of top-level publications. Most notably, 4 papers ([11], [14], [12], [18]) were presented at ACM SIGGRAPH 2013 in Los Angeles. These ACM Transactions on Graphics papers are the best publication in our field.

The continuing industrial interest in our work, both via bilateral contracts and EU initiatives with companies has grown in this year, on the topics of image-based relighting, image-based rendering and materials for vector art. This is a very promising trend for the future.

3. Research Program

3.1. Rendering

We consider plausible rendering to be a first promising research direction, both for images and for sound. Recent developments, such as point rendering, image-based modeling and rendering, and work on the simulation of aging indicate high potential for the development of techniques which render *plausible* rather than extremely accurate images. In particular, such approaches can result in more efficient renderings of very complex scenes (such as outdoors environments). This is true both for visual (image) and sound rendering. In the case of images, such techniques are naturally related to image- or point-based methods. It is important to note that these models are becoming more and more important in the context of network or heterogeneous rendering, where the traditional polygon-based approach is rapidly reaching its limits. Another research direction of interest is realistic rendering using simulation methods, both for images and sound. In some cases, research in these domains has reached a certain level of maturity, for example in the case of lighting and global illumination. For some of these domains, we investigate the possibility of technology transfer with appropriate partners. Nonetheless, certain aspects of these research domains, such as visibility or high-quality sound still have numerous and interesting remaining research challenges.

3.1.1. Plausible Rendering

3.1.1.1. Alternative representations for complex geometry

The key elements required to obtain visually rich simulations, are sufficient geometric detail, textures and lighting effects. A variety of algorithms exist to achieve these goals, for example displacement mapping, that is the displacement of a surface by a function or a series of functions, which are often generated stochastically. With such methods, it is possible to generate convincing representations of terrains or mountains, or of non-smooth objects such as rocks. Traditional approaches used to represent such objects require a very large number of polygons, resulting in slow rendering rates. Much more efficient rendering can be achieved by using point or image based rendering, where the number of elements used for display is view- or image resolution-dependent, resulting in a significant decrease in geometric complexity. Such approaches have very high potential. For example, if all object can be rendered by points, it could be possible to achieve much higher quality local illumination or shading, using more sophisticated and expensive algorithms, since geometric complexity will be reduced. Such novel techniques could lead to a complete replacement of polygon-based rendering for complex scenes. A number of significant technical challenges remain to achieve such a goal, including sampling techniques which adapt well to shading and shadowing algorithms, the development of algorithms and data structures which are both fast and compact, and which can allow interactive or real-time rendering. The type of rendering platforms used, varying from the high-performance graphics workstation all the way to the PDA or mobile phone, is an additional consideration in the development of these structures and algorithms. Such approaches are clearly a suitable choice for network rendering, for games or the modelling of certain natural object or phenomena (such as vegetation, e.g. Figure 1, or clouds). Other representations merit further research, such as image or video based rendering algorithms, or structures/algorithms such as the "render cache" [33], which we have developed in the past, or even volumetric methods. We will take into account considerations related to heterogeneous rendering platforms, network rendering, and the appropriate choices depending on bandwith or application. Point- or image-based representations can also lead to novel solutions for capturing and representing real objects. By combining real images, sampling techniques and borrowing techniques from other domains (e.g., computer vision, volumetric imaging, tomography etc.) we hope to develop representations of complex natural objects which will allow rapid rendering. Such approaches are closely related to texture synthesis and image-based modeling. We believe that such methods will not replace 3D (laser or range-finger) scans, but could be complementary, and represent a simpler and lower cost alternative for certain applications (architecture, archeology etc.). We are also investigating methods for adding "natural appearance" to synthetic objects. Such approaches include *weathering* or *aging* techniques, based on physical simulations [23], but also simpler methods such as accessibility maps [30]. The approaches we intend to investigate will attempt to both combine and simplify existing techniques, or develop novel approaches founded on generative models based on observation of the real world.

3.1.1.2. Plausible audio rendering

Similar to image rendering, plausible approaches can be designed for audio rendering. For instance, the complexity of rendering high order reflections of sound waves makes current geometrical approaches inappropriate. However, such high order reflections drive our auditory perception of "reverberation" in a virtual environment and are thus a key aspect of a plausible audio rendering approach. In complex environments, such as cities, with a high geometrical complexity, hundreds or thousands of pedestrians and vehicles, the acoustic field is extremely rich. Here again, current geometrical approaches cannot be used due to the overwhelming number of sound sources to process. We study approaches for statistical modeling of sound scenes to efficiently deal with such complex environments. We also study perceptual approaches to audio rendering which can result in high efficiency rendering algorithms while preserving visual-auditory consistency if required.

3.1.2. High Quality Rendering Using Simulation

3.1.2.1. Non-diffuse lighting

A large body of global illumination research has concentrated on finite element methods for the simulation of the diffuse component and stochastic methods for the non-diffuse component. Mesh-based finite element approaches have a number of limitations, in terms of finding appropriate meshing strategies and form-factor calculations. Error analysis methodologies for finite element and stochastic methods have been very different in the past, and a unified approach would clearly be interesting. Efficient rendering, which is a major advantage of finite element approaches, remains an overall goal for all general global illumination research. For certain cases, stochastic methods can be efficient for all types of light transfers, in particular if we require a view-dependent solution. We are also interested both in *pure* stochastic methods, which do not use finite element techniques. Interesting future directions include filtering for improvement of final image quality as well as beam tracing type approaches [31] which have been recently developed for sound research.

3.1.2.2. Visibility and Shadows

Visibility calculations are central to all global illumination simulations, as well as for all rendering algorithms of images and sound. We have investigated various global visibility structures, and developed robust solutions for scenes typically used in computer graphics. Such analytical data structures [27], [26], [25] typically have robustness or memory consumption problems which make them difficult to apply to scenes of realistic size. Our solutions to date are based on general and flexible formalisms which describe all visibility event in terms of generators (vertices and edges); this approach has been published in the past [24]. Lazy evaluation, as well

Figure 1. Plausible rendering of an outdoors scene containing points, lines and polygons [22], representing a scene with trees, grass and flowers. We can achieve 7-8 frames per second compared to tens of seconds per image using standard polygonal rendering.

as hierarchical solutions, are clearly interesting avenues of research, although are probably quite application dependent.

3.1.2.3. Radiosity

For purely diffuse scenes, the radiosity algorithm remains one of the most well-adapted solutions. This area has reached a certain level of maturity, and many of the remaining problems are more technology-transfer oriented We are interested in interactive or real-time renderings of global illumination simulations for very complex scenes, the "cleanup" of input data, the use of application-dependent semantic information and mixed representations and their management. Hierarchical radiosity can also be applied to sound, and the ideas used in clustering methods for lighting can be applied to sound.

3.1.2.4. High-quality audio rendering

Our research on high quality audio rendering is focused on developing efficient algorithms for simulations of geometrical acoustics. It is necessary to develop techniques that can deal with complex scenes, introducing efficient algorithms and data structures (for instance, beam-trees [28] [31]), especially to model early reflections or diffractions from the objects in the environment. Validation of the algorithms is also a key aspect that is necessary in order to determine important acoustical phenomena, mandatory in order to obtain a high-quality result. Recent work by Nicolas Tsingos at Bell Labs [29] has shown that geometrical approaches can lead to high quality modeling of sound reflection and diffraction in a virtual environment (Figure 2). We will pursue this research further, for instance by dealing with more complex geometry (e.g., concert hall, entire building floors).

Finally, several signal processing issues remain in order to properly and efficiently restitute a 3D soundfield to the ears of the listener over a variety of systems (headphones, speakers). We would like to develop an open and general-purpose API for audio rendering applications. We already completed a preliminary version of a software library: AURELI [32].

4. Software and Platforms

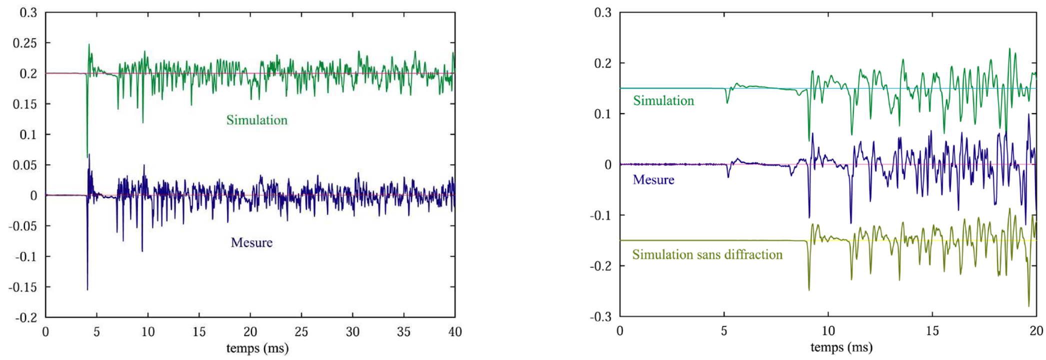

Figure 2. A comparison between a measurement (left) of the sound pressure in a given location of the "Bell Labs Box", a simple test environment built at Bell Laboratories, and a high-quality simulation based on a beam-tracing engine (right). Simulations include effects of reflections off the walls and diffraction off a panel introduced in the room.

4.1. Multi-View Image-Based Rendering and Relighting Suite

Participants: Clement Riant, Sylvain Duchêne, Pierre-Yves Laffont, Adrien Bousseau, George Drettakis.

We have designed and implemented a set of libraries for handling multi-view image-based rendering and relighting algorithms. These constitute the basis for the software developed for the EU projects VERVE and CR-PLAY.

4.1.1. RID: Rich Intrinsic-image Decomposer

We developed a software platform to perform rich intrinsic decomposition methods from photographs of outdoor scenes, as described in [13]. It includes main scripts and functions in Matlab for treatment of the input data, interfaces to software for multi-view reconstruction (Bundler, PMVS) and meshing from point clouds (method developed by Julie Digne, a postdoc in the GEOMETRICA project team). We then interface software for image matting using the Matting Laplacian, and User-Assisted Intrinsic Images. The system also includes an interface with Adobe Photoshop, for visualization and demonstration of our results in end-user image editing software. The method performs the computation of sun, sky and indirect lighting received at 3D points of an automatically reconstructed scene, using a modified version of the PBRT stochastic raytracer. Finally, there is a scene calibration module and an OpenGL viewer.

4.1.2. ROSSE: Relighting Outdoor Scenes with Shadow Editing

This software package includes a set of modules for processing point clouds and meshes produced by automatic multi-view stereo computer vision solutions. It includes all file management, point cloud and mesh handling, as well as ray-tracing using the Intel Embree ray tracer to computer illumination properties on the mesh. An interactive viewer is also included. A new intrinsic image approach is included as well as a module for relighting and shadow movement, based on an image-driven approach to moving cast shadows.

4.1.3. SWARPI: Superpixel Warp for Image-based rendering

Depth Synthesis and Warped-Based Superpixel Image-Based Rendering. This software package is the implementation of the publication [12]. The main software consists of two components: the depth synthesis step and the image-based rendering step:

a) The depth synthesis step is a Matlab package that reads 3D points coming from an automated 3D reconstruction pipeline, together with images and calibrated cameras, and produces the superpixel decomposition and the depth synthesis algorithm. The current version uses the open source packages bundler and PMVS (GPL v3 license), but other 3D reconstruction approaches could be used instead. b) The rendering step is in C++, and takes the result of the first step as input to allow interactive navigation. The code uses multi-pass deferred shading with geometry shaders (OpenGL 4.0 or above) to perform the rendering.

In addition to the implementation of [12], we have developed a Matlab interface for manual depth correction ("depth painting") An APP (Agency for the Protection of Programs) registration of this software is pending.

4.2. APF: state-of-the-art 3D audio library

Participants: Adrien David, George Drettakis.

This work was performed in collaboration with Jean-Christophe Lombardo of the DREAM group (i.e., the research support development group of our Inria center). REVES has several audio research publications over the last 10 years, which correspond to a class of functionalities such as clustering, masking, progressive processing etc.. The first component is the masking or culling algorithm, which aims at removing all the inaudible audio sources from a virtual scene based on perceptual metrics. The second component, called clustering, aims at grouping audio sources that are spatially close to each other and premix them to a representative cluster source, so that all spatialization related processing can be applied only on the representative premixed source [9]. Other audio topics were also considered and developed, like progressive and scalable frequency domain mixing, sound propagation, scalable reverberation, modal sound synthesis and contact sounds generation [1].

In order to maintain all the knowledge in the group and re-use these technologies in the Immersive Space, a previous young engineer (David Grelaud) wrote a fully documented audio library (APF) which gathers about 10 audio publications and 1 US patent. APF is a cross-platform, object oriented C++ API available on GForge. All the code has been re-implemented and a completely new software architecture resulted in a twofold increase in the speed of our algorithms. APF runs in the Immersive Space and uses the tracking system to spatialize virtual audio sources around the listener. It can also exploit personal Head Related Transfer Functions (HRTF).

We have implemented a network communications layer to create an audio rendering server on a separate machine, and the library is fully integrated into the osgVR platform.

APF has been critical in establishing collaborations in the context of various grant proposals (EU and national).

4.3. GaborNoise Software

Participants: Ares Lagae, George Drettakis.

We proposed a new procedural noise function last year, Gabor noise [6]. In the context of this project, we have developed a software package, which includes a CPU reference implementation of the 2D noise, and a complete GPU implementation of the 2D noise, surface noise, and 3D noise. This software package has been filed for APP protection and is in the process of being transferred to industrial partners.

This work is a collaboration with Sylvain Lefebvre, former member of the team, now in the ALICE project-team, Inria Nancy - Grand Est.

4.4. Gabor Noise By Example

Participant: George Drettakis.

In collaboration with B. Galerne, S. Lefebvre and A. Lagae (KU Leuven) we have released to code for the 2012 SIGGRAPH paper Gabor Noise By Example (see http://www-sop.inria.fr/reves/Basilic/2012/GLLD12/). This includes a matlab code for the analysis and C++/cuda code for the synthesis.

5. New Results

5.1. Plausible and Realistic Image Rendering

5.1.1. Depth Synthesis and Local Warps for Interactive Image-based Navigation

Participants: Gaurav Chaurasia, Sylvain Duchene, George Drettakis.

Figure 3. Novel views generated by the image-based rendering approach of [12] along with a visualization of novel camera position relative to the 3D scene and input cameras. This approach is among the first to handle very complex urban scenes such as those shown here and provide a stable solution for viewpoints that are far from the input cameras.

Modern camera calibration and multi-view stereo techniques enable users to smoothly navigate between different views of a scene captured using standard cameras. The underlying automatic 3D reconstruction methods work well for buildings and regular structures but often fail on vegetation, vehicles and other complex geometry present in everyday urban scenes. Consequently, missing depth information makes image-based rendering for such scenes very challenging. This paper introduces a new image-based rendering algorithm that is robust to missing or unreliable geometry, providing plausible novel views even in regions quite far from the input camera positions. The approach first oversegments the input images, creating superpixels of homogeneous color content which preserve depth discontinuities. It then introduces a *depth synthesis* step for poorly reconstructed regions. It defines a graph on the superpixels and uses *shortest walk* traversals to fill unreconstructed regions with approximate depth from regions that are well-reconstructed and similar in visual content. The superpixels augmented with synthesized depth allow a local shape-preserving warp which warps each superpixel of the input image to the novel view without incurring distortions and preserving the local visual content within the superpixel. This allows the approach to effectively compensate for missing photoconsistent depth, the lack of which is known to cause rendering artifacts. The final rendering algorithm blends the warped images, using heuristics to avoid ghosting artifacts. The results demonstrate novel view synthesis in real time for multiple challenging scenes with significant depth complexity (see Figure 3), providing a convincing immersive navigation experience. The paper presents comparisons with three of the state of the art image-based rendering techniques and demonstrate clear advantages.

This work was in collaboration with Olga Sorkine-Hornung at ETH Zurich. It has been published in ACM Transactions on Graphics 2013 [12] and presented at SIGGRAPH.

5.1.2. Megastereo: Constructing High-Resolution Stereo Panoramas

Participant: Christian Richardt.

There is currently a strong consumer interest in a more immersive experience of content, such as 3D photographs, television and cinema. A great way of capturing environmental content are panoramas (see Figure 4). We present a solution for generating high-quality stereo panoramas at megapixel resolutions. While previous approaches introduced the basic principles, we show that those techniques do not generalise well to today's high image resolutions and lead to disturbing visual artefacts. We describe the necessary correction steps and a compact representation for the input images in order to achieve a highly accurate approximation to the required ray space. In addition, we introduce a flow-based upsampling of the available input rays which

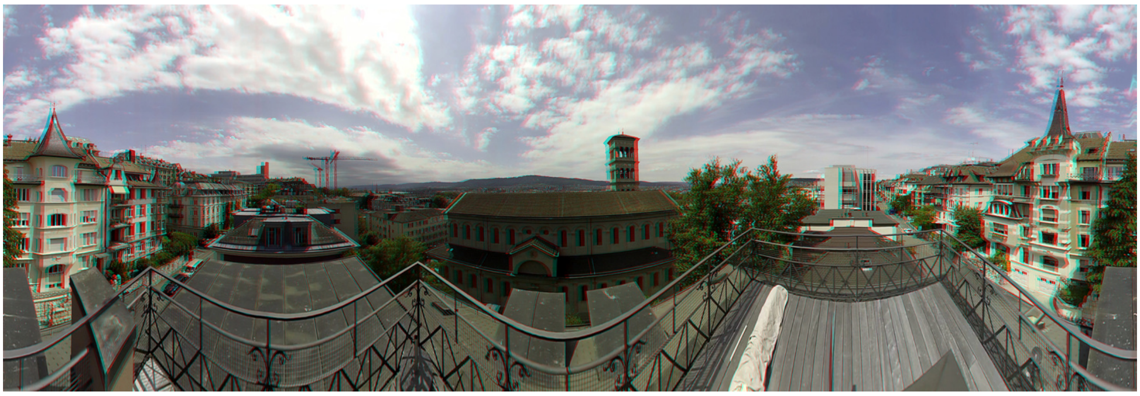

Figure 4. A stereoscopic panorama corrected and stitched using our techniques, shown as red-cyan anaglyph image.

effectively resolves known aliasing issues like stitching artefacts. The required rays are generated on the fly to perfectly match the desired output resolution, even for small numbers of input images. This upsampling is real-time and enables direct interactive control over the desired stereoscopic depth effect. In combination, our contributions allow the generation of stereoscopic panoramas at high output resolutions that are virtually free of artefacts such as seams, stereo discontinuities, vertical parallax and other mono-/stereoscopic shape distortions.

This work was carried out in collaboration with Yael Pritch, Henning Zimmer and Alexander Sorkine-Hornung at Disney Research Zurich. The paper has been published as an oral presentation at CVPR 2013 [20].

5.1.3. Probabilistic Connection Path Tracing

Participants: Stefan Popov, George Drettakis.

We propose an unbiased generalization of bi-directional path tracing (BPT) that significantly improves its rendering efficiency. Our main insight is that the set of paths traced by BPT contains a significant amount of statistical information, that is not exploited.

BPT repeatedly builds an eye and a light sub-paths, connects them, estimates the contribution to the corresponding pixel and then throws the path away. Instead, we propose to first trace all eye and light sub-paths, and then probabilistically connect each eye sub-path to one or more light sub-paths. From a Monte-Carlo perspective, this will connect each light to each eye sub-path, substantially increasing the number of paths used to estimate the solution. As a result, the convergence will be significantly increased as well.

This work is a collaboration with Frédo Durand from the Massachusetts Institute of Technology, Cambridge and Ravi Ramamoorthi from University of California, Berkeley in the context of the CRISP Associated Team.

5.1.4. Parallelization Strategies for Associative Image Processing Operators

Participants: Gaurav Chaurasia, George Drettakis.

Basic image processing operations have been optimized on a case-by-case basis such as prefix sums and recursive filters. Moreover, these optimized algorithms are very complicated to program because parallelization involves non-trivial splitting of the input domain of the operator. The target of this is to generalize the optimization heuristics of a generic class of associative image processing operators by developing an algebraic understanding of the operator and parallelization options. The algebra can transform associative operations such as box filters, summed area table, recursive filters etc. by splitting their domain to smaller subsets of the input image that can be executed in parallel and recombine the intermediate result later. The ultimate target is to develop a compiler front-end based on the Halide language that implements this algebra and is capable of parallelizing associative operators of arbitrary footprints by a few lines of code, thereby relieving the programmer of the tedious task for programming the parallelized algorithms. Such a compiler would allow programmers to easily experiment with a plethora of parallelization strategies in a systematic manner.

This work is in collaboration with Jonathan Ragan-Kelley and Fredo Durand of MIT and Sylvain Paris (Adobe Research).

5.1.5. Lightfield Editing

Participant: Adrien Bousseau.

Lightfields capture multiple nearby views of a scene and are consolidating themselves as the successors of conventional photographs. As the field grows and evolves, the need for tools to process and manipulate lightfields arises. However, traditional image manipulation software such as Adobe Photoshop are designed to handle single views and their interfaces cannot cope with multiple views coherently. In this work we evaluate different user interface designs for lightfield editing. Our interfaces differ mainly in the way depth is presented to the user and build upon different depth perception cues.

This work is a collaboration with Adrian Jarabo, Belen Masia and Diego Gutierrez from Universidad de Zaragoza and Fabio Pellacini from Sapienza Universita di Roma.

5.2. Perception for Plausible Rendering

5.2.1. Perception of Perspective Distortions in Image-Based Rendering

Participants: Peter Vangorp, Christian Richardt, Gaurav Chaurasia, George Drettakis.

Figure 5. This interactive navigation tool shows an inset (in the lower left) that predicts comfort ratings for all possible camera orientations as seen from the blue camera's viewpoint. The application also restricts the user's motion to regions with acceptable predicted quality (in blue and yellow).

Image-based rendering (IBR) creates realistic images by enriching simple geometries with photographs, for example by mapping the photograph of a building façade onto a plane. However, as soon as the viewer moves away from the correct viewpoint, the image in the retina becomes distorted, sometimes leading to gross misperceptions of the original geometry. Two hypotheses from vision science state how viewers perceive such image distortions, one claiming that they can compensate for them (and therefore perceive scene geometry reasonably correctly), and one claiming that they cannot compensate (and therefore can perceive rather significant distortions). We modified the latter hypothesis so that it extends to street-level IBR. We then conducted a rigorous experiment that measured the magnitude of perceptual distortions that occur with IBR for façade viewing. We also conducted a rating experiment that assessed the acceptability of the distortions. The results of the two experiments were consistent with one another. They showed that viewers' percepts are indeed distorted, but not as severely as predicted by the modified vision science hypothesis. From our experimental results, we develop a predictive model of distortion for street-level IBR, which we use to provide guidelines for acceptability of virtual views and for capture camera density. We perform a confirmatory study to validate our predictions, and illustrate their use with an application that guides users in IBR navigation to stay in regions where virtual views yield acceptable perceptual distortions (see Figure 5).

This work is a collaboration with Emily Cooper and Marty Banks at UC Berkeley, within the associate team CRISP. The paper was accepted as a SIGGRAPH 2013 paper and published in the ACM Transactions on Graphics journal [18].

5.2.2. Gloss Perception in Painterly and Cartoon Rendering

Participant: Adrien Bousseau.

Depictions with traditional media such as painting and drawing represent scene content in a stylized manner. It is unclear however how well stylized images depict scene properties like shape, material and lighting. In this project, we use non photorealistic rendering algorithms to evaluate how stylization alters the perception of gloss (see Figure 6). Our study reveals a compression of the range of representable gloss in stylized images so that shiny materials appear more diffuse in painterly rendering, while diffuse materials appear shinier in cartoon images.

From our measurements we estimate the function that maps realistic gloss parameters to their perception in a stylized rendering. This mapping allows users of NPR algorithms to predict the perception of gloss in their images. The inverse of this function exaggerates gloss properties to make the contrast between materials in a stylized image more faithful. We have conducted our experiment both in a lab and on a crowdsourcing website. While crowdsourcing allows us to quickly design our pilot study, a lab experiment provides more control on how subjects perform the task. We provide a detailed comparison of the results obtained with the two approaches and discuss their advantages and drawbacks for studies similar to ours.

This work is a collaboration with James O'Shea, Ravi Ramamoorthi and Maneesh Agrawala from UC Berkeley in the context of the Associate Team CRISP (see also Section 1) and Frédo Durand from MIT. It has been published in ACM Transactions on Graphics 2013 [11] and presented at SIGGRAPH.

5.2.3. A High-Level Visual Attention Model

Participant: George Drettakis.

The goal of this project is to develop a high-level attention model based on memory schemas and singleton theory in visual perception. We have developed an approach extending a Bayesian approach to attention, which incorporates these high level features and can be directly used in a game engine to improve scene design.

This project is in collaboration with the Tech. University of Crete in the context of the Ph.D. of George Koulieris, supervised by Prof. Katerina Mania and BTU Cottburg (D. Cunningham).

5.3. Interaction and Design for Virtual Environments

5.3.1. Diffusion Curves: A Vector Representation for Smooth-Shaded Images

Participant: Adrien Bousseau.

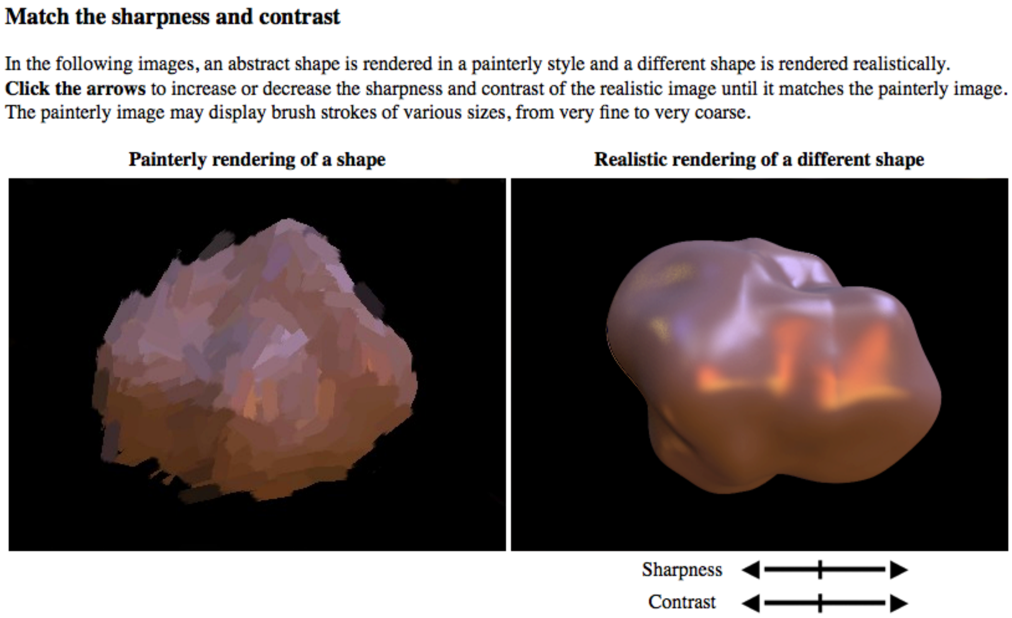

Figure 6. The experimental task used for studying gloss perception in stylized images.

This paper was selected for presentation in the Communications of the ACM, as an important graphics research result of interest to the entire Computer Science community. We describe a new vector-based primitive for creating smooth-shaded images, called the diffusion curve. A diffusion curve partitions the space through which it is drawn, defining different colors on either side. These colors may vary smoothly along the curve. In addition, the sharpness of the color transition from one side of the curve to the other can be controlled. Given a set of diffusion curves, the final image is constructed by solving a Poisson equation whose constraints are specified by the set of gradients across all diffusion curves (Figure 7). Like all vector-based primitives, diffusion curves conveniently support a variety of operations, including geometry-based editing, keyframe animation, and ready stylization. Moreover, their representation is compact and inherently resolution independent. We describe a GPU-based implementation for rendering images defined by a set of diffusion curves in real time. We then demonstrate an interactive drawing system allowing artists to create artwork using diffusion curves, either by drawing the curves in a freehand style, or by tracing existing imagery. Furthermore, we describe a completely automatic conversion process for taking an image and turning it into a set of diffusion curves that closely approximate the original image content.

This work is a collaboration with Alexandrina Orzan, Pascal Barla (Inria / Manao), Holger Winnemöller (Adobe Systems), Joëlle Thollot (Inria / Maverick) and David Salesin (Adobe Systems). This work was originally published in ACM Transactions on Graphics (Proceeding of SIGGRAPH 2008) and was selected for publication in Communications of the ACM July 2013 [15].

5.3.2. Natural Gesture-based Interaction for Complex Tasks in an Immersive Cube

Participants: Emmanuelle Chapoulie, George Drettakis.

We present a solution for natural gesture interaction in an immersive cube in which users can manipulate objects with fingers of both hands in a close-to-natural manner for moderately complex, general purpose tasks. Our solution uses finger tracking coupled with a real-time physics engine, combined with a comprehensive approach for hand gestures, which is robust to tracker noise and simulation instabilities. To determine if our

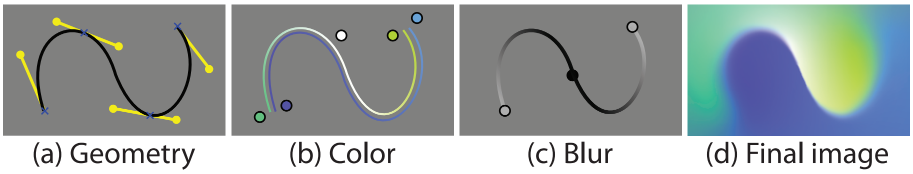

Figure 7. A diffusion curve consists of a Bézier curve (a) enriched with color (b) and blur (c) control points. The final image (d) is obtained by diffusing the colors in the image domain.

natural gestures are a feasible interface in an immersive cube, we perform an exploratory study for tasks involving the user walking in the cube while performing complex manipulations such as balancing objects. We compare gestures to a traditional 6-DOF Wand, and we also compare both gestures and Wand with the same task, faithfully reproduced in the real world. Users are also asked to perform a free task, allowing us to observe their perceived level of presence in the scene. Our results show that our robust approach provides a feasible natural gesture interface for immersive cube-like environments and is perceived by users as being closer to the real experience compared to the Wand.

This work is a collaboration with Jean-Christophe Lombardo of SED, with Evanthia Dimara and Maria Roussou from the University of Athens and with Maud Marchal from IRISA-INSA/Inria Rennes - Bretagne Atlantique. The work is under review in the journal Virtual Reality.

5.3.3. Evaluation of Direct Manipulation in an Immersive Cube: a Controlled Study

Participants: Emmanuelle Chapoulie, George Drettakis.

We are pursuing a study for interaction using finger tracking and traditional 6 degrees of freedom (DOF) flysticks in a virtual reality immersive cube. Our study aims at identifying which factors make one interface better than the other and which are the tradeoffs for the design of experiments, thus decomposing the movements into restricted DOF.

5.3.4. The Drawing Assistant: Automated Drawing Guidance and Feedback from Photographs

Participants: Emmanuel Iarussi, Adrien Bousseau.

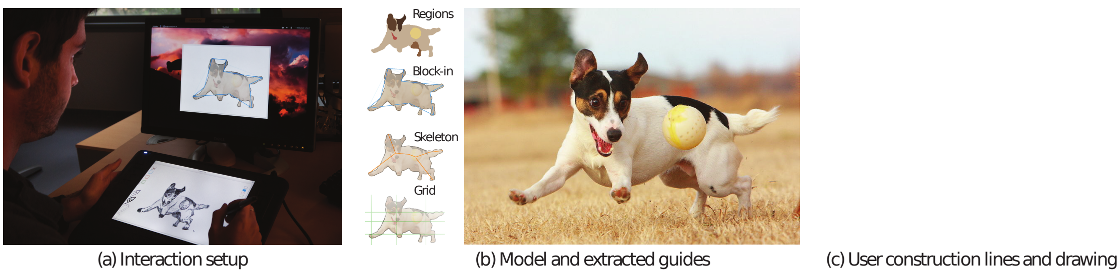

Figure 8. Our drawing assistant provides guidance and feedback over a model photograph that the user reproduces on a virtual canvas (a). We use computer vision algorithms to extract visual guides that enhance the geometric structures in the image (b). In this example, the user first sketched the block-in construction lines (c, blue) before drawing the regions and adding details. This guidance helps users produce more accurate drawings.

Drawing is the earliest form of visual depiction and continues to enjoy great popularity with paint systems. However, drawing requires artistic skills that many people feel out of reach. We developed an interactive drawing tool that provides automated guidance over model photographs to help people practice traditional drawing-by-observation techniques. The drawing literature describes a number of techniques to help people gain consciousness of the shapes in a scene and their relationships. We compile these techniques and derive a set of construction lines that we automatically extract from a model photograph (see Figure 8). We then display these lines over the model to guide its manual reproduction by the user on the drawing canvas. Our pen-based interface also allows users to navigate between the techniques they wish to practice and to draw construction lines in dedicated layers. We use shape-matching to register the user's sketch with the model guides. We use this registration to provide corrective feedback to the user. We conducted two user studies to inform the design of our tool and evaluate our approach with a total of 20 users. Participants produced better drawings using the drawing assistant, with more accurate proportions and alignments. They also perceived that guidance and corrective feedback helped them better understand how to draw. Finally, some participants spontaneously applied the techniques when asked to draw without our tool after using it for about 30 minutes.

This work is a collaboration with Theophanis Tsandilas from the InSitu project team - Inria Saclay, in the context of the ANR DRAO project. It has been published at proceedings of UIST 2013 the 26th annual ACM symposium on User interface software and technology [19].

5.3.5. Shape-Aware Sketch Editing with Covariant-Minimizing Cross Fields

Participants: Emmanuel Iarussi, Adrien Bousseau.

Free-hand sketches are extensively used in product design for their ability to convey 3D surfaces with a handful of pen strokes. Skillful artists capture all surface information by strategically positioning strokes so that they depict the feature lines and curvature directions of surface patches. Viewers envision the intended 3D surface by mentally interpolating these lines to form a dense network representative of the curvature of the shape. Our goal is to mimic this interpolation process to estimate at each pixel of a sketch the projection of the two principal directions of the surface, or their extrapolation over umbilic regions. While the information we recover is purely 2D, it provides a vivid sense of the intended 3D surface and allows various shape-aware sketch editing applications, including normal estimation for shading, cross-hatching rendering and surface parameterization for texture mapping.

This work is a collaboration with David Bommes from the Titane project team, Sophia-Antipolis.

5.3.6. Depicting Stylized Materials with Vector Shade Trees

Participants: Jorge Lopez-Moreno, Stefan Popov, Adrien Bousseau, George Drettakis.

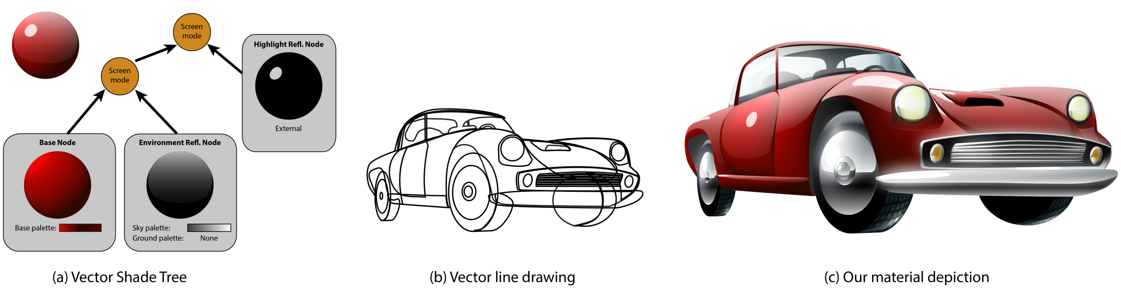

Figure 9. Depicting Stylized Materials with Vector Shade Trees.

Vector graphics represent images with compact, editable and scalable primitives. Skillful vector artists employ these primitives to produce vivid depictions of material appearance and lighting. However, such stylized imagery often requires building complex multi-layered combinations of colored fills and gradient meshes. We facilitate this task by introducing vector shade trees that bring to vector graphics the flexibility of modular shading representations as known in the 3D rendering community. In contrast to traditional shade trees that combine pixel and vertex shaders, our shade nodes encapsulate the creation and blending of vector primitives that vector artists routinely use. We propose a set of basic shade nodes that we design to respect the traditional guidelines on material depiction described in drawing books and tutorials. We integrate our representation as an Adobe Illustrator plug-in that allows even inexperienced users to take a line drawing, apply a few clicks and obtain a fully colored illustration. More experienced artists can easily refine the illustration, adding more details and visual features, while using all the vector drawing tools they are already familiar with. We demonstrate the power of our representation by quickly generating illustrations of complex objects and materials.

Figure 9 illustrates how our algorithm works. We use a combination of basic shade nodes composed of vector graphics primitives to describe Vector Shade Trees that represent stylized materials (a). Combining these nodes allows the depiction of a variety of materials while preserving traditional vector drawing style and practice. We integrate our vector shade trees in a vector drawing tool that allows users to apply stylized shading effects on vector line drawings (b,c).

This work is a collaboration with Maneesh Agrawala from University of California, Berkeley in the context of the CRISP Associated Team. The work was accepted as a SIGGRAPH 2013 paper and published in ACM Transactions on Graphics, volume 32, issue 4 [14].

5.3.7. Auditory-Visual Aversive Stimuli Modulate the Conscious Experience of Fear

Participants: Rachid Guerchouche, George Drettakis.

In a natural environment, affective information is perceived via multiple senses, mostly audition and vision. However, the impact of multisensory information on affect remains relatively undiscovered. In this study, we investigated whether the auditory-visual presentation of aversive stimuli influences the experience of fear. We used the advantages of virtual reality to manipulate multisensory presentation and to display potentially fearful dog stimuli embedded in a natural context. We manipulated the affective reactions evoked by the dog stimuli by recruiting two groups of participants: dog-fearful and non-fearful participants. The sensitivity to dog fear was assessed psychometrically by a questionnaire and also at behavioral and subjective levels using a Behavioral Avoidance Test (BAT). Participants navigated in virtual environments, in which they encountered virtual dog stimuli presented through the auditory channel, the visual channel or both. They were asked to report their fear using Subjective Units of Distress. We compared the fear for unimodal (visual or auditory) and bimodal (auditory-visual) dog stimuli. Dog-fearful participants as well as non-fearful participants reported more fear in response to bimodal audiovisual compared to unimodal presentation of dog stimuli. These results suggest that fear is more intense when the affective information is processed via multiple sensory pathways, which might be due to a cross-modal potentiation. Our findings have implications for the field of virtual reality-based therapy of phobias. Therapies could be refined and improved by implicating and manipulating the multisensory presentation of the feared situations.

This work is a collaboration with Marine Taffou and Isabelle Viaud-Delmon from CNRS-IRCAM, in the context of the European project VERVE. The work was published in the Multisensory Research journal 2013 [17].

5.3.8. Memory Motivation Virtual Experience

Participants: Emmanuelle Chapoulie, Rachid Guerchouche, George Drettakis.

Memory complaints are known to be one the first stages of Alzheimer's disease, for which -up to now, there is no known chemical treatment. In the context of the European project VERVE, and in collboartion with the Resources and Research Memory Centre of Nice Hospital (CM2R), we performed a study on the feasibility of treating memory complaints using realistic immersive virtual environments. Such environments are created

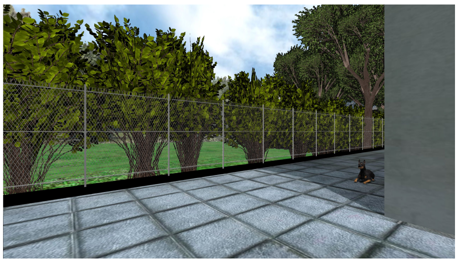 

Figure 10. Pictures of the auditory-visual VEs used to measure the participants' fear when encountering virtual dogs. On the left, the outdoor garden scene and on the right, the indoor hangar scene.

using Image-Based Rendering technique developed by REVES. It is possible to easily provide, realistic 3D environments of places familiar to the participants using only a few photograph, and investigate whether IBR virtual environments can convey familiarity.

This work is a collaboration with Pierre-David Petit and Pr. Philippe Robert from CM2R. The work will be presented in IEEE Virtual Reality conference 2014 and will be published in the conference proceedings.

5.3.9. Layered Image Vectorization

Participants: Christian Richardt, Adrien Bousseau, George Drettakis.

Vector graphics enjoy great popularity among graphic designers for their compactness, scalability and editability. The goal of *vectorization* algorithms is to facilitate the creation of vector graphics by converting bitmap images into vector primitives. However, while a vectorization algorithm should faithfully reproduce the appearance of a bitmap image, it should also generate vector primitives that are easily editable – a goal that existing methods have largely overlooked. We investigate layered representations which are more compact and editable, and hence better preserve the strengths of vector graphics. This work is in collaboration with Maneesh Agrawala in the context of the CRISP Associated Team and Jorge Lopez-Moreno, now a postdoc at the University of Madrid.

5.3.10. True2Form: Automatic 3D Concept Modeling from Design Sketches

Participants: Adrien Bousseau.

We developed a method to estimate smooth 3D shapes from design sketches. We do this by hypothesizing and perceptually validating a set of local geometric relationships between the curves in sketches. We then algorithmically reconstruct 3D curves from a single sketch by detecting their local geometric relationships and reconciling them globally across the 3D curve network.

This work is a collaboration with James McCrae and Karan Singh from the University of Toronto and Xu Baoxuan, Will Chang and Alla Sheffer from the University of British Columbia.

6. Bilateral Contracts and Grants with Industry

6.1. Bilateral Contracts with Industry

6.1.1. Autodesk

Participants: Adrien Bousseau, George Drettakis, Clement Riant, Sylvain Duchene.

We extended our technology transfer agreement with Autodesk concerning the RID technology on single-lighting condition intrinsic images. We transfered a first version of the software on Autodesk servers.

6.2. Bilateral Grants with Industry

6.2.1. Autodesk

Participants: Adrien Bousseau, George Drettakis, Clement Riant, Sylvain Duchene.

Autodesk has offered a significant research donation to REVES in support of our work on intrinsic images. Autodesk has also donated several licenses of Maya, 3DS Max and SketchBookPro.

6.2.2. Adobe

Participants: George Drettakis, Gaurav Chaurasia.

Adobe has offered a small donation in the context of the Halide project (Sec. 5.1.4). Adobe has also signed an evaluation license for the Vector Shade Trees software developed in the context of [14].

7. Partnerships and Cooperations

7.1. National Initiatives

7.1.1. ANR ALTA

Participants: Emmanuelle Chapoulie, Stefan Popov, George Drettakis.

The ANR ALTA project started in October 2011, and focuses on the development of novel algorithms for realistic and efficient global illumination. The project is coordinated by the Grenoble Inria group ARTIS (N.Holzschuch), and the Bordeaux Inria group MANAO (X. Granier) is also a partner. Our participation is the study of error bounds for these algorithms and the development of interactive global illumination, and the development of the new global illumination algorithm described in Sec. 5.1.3.

7.1.2. ANR DRAO

Participants: Emmanuel Iarussi, Adrien Bousseau.

https://www-sop.inria.fr/members/Adrien.Bousseau/drao/

The ANR DRAO is a young researcher project coordinated by Adrien Bousseau, in collaboration with the InSitu project team at Inria Saclay - Ile de France (W. Mackay and T. Tsandilas) and the MANAO project team (P. Barla and G. Guennebaud) and POTIOC project team (M. Hachet) at Inria Bordeaux - Sud Ouest. The goal of this collaboration is to develop novel drawing tools for amateurs as well as for expert designers and illustrators, combining expertise in Computer Graphics (REVES and MANAO) and Human-Computer Interaction (InSitu, POTIOC). This ANR project funds the PhD of Emmanuel Iarussi.

The first part of the project will be to observe how people draw with existing tools. To do so we will conduct observational studies where we will interview designers and illustrators and collect data by videotaping drawing sessions and by recording drawings with digital pens. In the second part of the project we will deduce from our observations new user interfaces and rendering algorithms that automate part of the drawing process and enrich 2D drawings with realistic rendering capabilities. We will combine computer vision and computer graphics techniques to estimate geometric information from sketches. We will then use this information to guide rendering algorithms that generate plausible depictions of material and lighting over the drawing. In the third part of the project, we plan to develop computer-assisted drawing lessons to teach amateurs how to draw from photographs and 3D models. We will apply image analysis algorithms to estimate the structure of a photograph and use that structure as guidance for drawing. To summarize, the goal of the ANR DRAO project is to make amateurs more confident in their drawing skills and to allow expert designers to produce complex illustrations more effectively.

The ANR DRAO has resulted in two publications this year on assisting drawing from photographs [19] and vector drawing of stylized materials [14].

7.1.3. ANR SEMAPOLIS

Participant: George Drettakis.

This ANR project started in October 2013. The goal is to use semantic information to improve urban reconstruction and rendering. The consortium is led by ENPC (R. Marlet) and includes the Inria Willow team and the GREY-C laboratory on image processing. Our contribution will be in the rendering part.

7.2. European Initiatives

7.2.1. FP7 Projects

7.2.1.1. VERVE

Title: VERVE

Type: COOPERATION (ICT)

Defi: Services to promote E-inclusion using socially realistic virtual environments

Instrument: Integrated Project (IP)

Duration: October 2011 - September 2014

Coordinator: Trinity College - Dublin (Ireland)

Others partners: DFKI (Germany), CNRS-ParisTech (France), CNRS-IRCAM (France), U. of Zaragoza (Spain), Testaluna (IT), KAINOS (UK)

See also: http://www.verveconsortium.eu/

Abstract

Social exclusion has many causes, but major factors are the fear and apathy that often accompany a disability. The European e-Inclusion policy stresses the importance of ICT in improving the quality of life in potentially disadvantaged groups, including older people and persons with disabilities. In this project, we will develop ICT tools to support the treatment of people who are at risk of social exclusion due to fear and/or apathy associated with a disability. These tools will be in the form of personalised VR scenarios and serious games specifically designed for therapeutic targets and made broadly available via a novel integration of interactive 3D environments directly into Web browsers. We will perform cutting edge research into rendering and simulating personalised and populated VR environments, 3D web graphics, and serious games. These technical efforts will be underpinned by our clinical/laboratory and industry partners, who will be fully involved throughout in the requirements, design and evaluation of VERVE, and liaison with the stakeholders (i.e., participants, carers/family, and health professionals). They will implement the VERVE interventions in three use-cases, each targeting a different group of participants: fear of falling, apathy related to cognitive decline and behavioural disturbances, and other emotional disturbances linked to anxiety. While developing clinical assessment methods and interventions for the first two patient groups is our primary focus, our results will be applicable to a much wider range of potentially disadvantaged individuals.

For the second period (October 2012 - September 2013), the consortium continued the work on implementing and improving the different solutions for the three use-cases: fear, apathy and anxieties. Different technologies were developed:

- Kitchen, a serious game for apathy.
- Freezing of Gait, a serious game for fear.
- Fear of Falling, a serious game for fear.
- Crowd-Phobia, a virtual reality application for anxieties.
- Memory Motivation Virtual Experience (MeMoVE), virtual reality application for apathy.

In particular REVES was mainly involved in the second use-case with the MeMoVE scenario. During this second period, the IBR technique was ported to the Immersive Space on a single screen of the CAVE. Experiments with healthy adults were performed in collaboration with the hospital of Nice (CHUN). The results of these experiments will be published in IEEE VR2014.

7.2.1.2. CR-PLAY – Capture Reconstruct Play

Type: COOPERATION (ICT)

Instrument: Specific Targeted Research Project

Objectif: Creativity

Duration: November 2013 - October 2016

Coordinator: Testaluna SA (IT)

Partner: TU Darmstadt (DE), UC London (UK), U. Patras (GR), Miniclip UK, Cursor Oy (FI)

Inria contact: George Drettakis

Abstract: The goal of this project is to use image- and video-based rendering and relighting techniques in the context of games and in particular mobile or casual games. The computer graphics and vision partners (UCL, TUD) are leaders in their fields, and have developed algorithms allowing easy capture of scenes using images and video, and reconstruction using vision algorithms. UCL and Inria have developed image- and video-based rendering algorithms which can be useful for games. These tools need to be perfected, reducing artifacts and difficulty of use so that they can be useful and productive for games companies. For evaluation, the HCI lab of the University of Patras will provide cutting-edge methodologies to make the resulting systems useable. The consortium is led by the games company Testaluna, based in Genova Italy, with whom we have a solid working relationship from our previous VERVE project (see above). Other industrial partners include Cursor Oy (a regional group of games companies in Finland, which is a leader in Europe in Casual games) and Miniclip, which is one of the major players in the online game market.

7.3. International Initiatives

7.3.1. Inria Associate Teams

- EA CRISP http://www-sop.inria.fr/reves/crisp/

  The goal of the CRISP associate team between REVES and University of California (UC) Berkeley is to investigate novel ways to create, render and interact with images based on the study of human perception. This novel and emerging area has been the focus of ongoing collaborations between researchers from the REVES research group at Inria (Adrien Bousseau, George Drettakis) and researchers in Computer Science and Vision Science at UC Berkeley (Maneesh Agrawala, Ravi Ramamoorthi, Martin S. Banks (Human Vision Science)). All of the researchers involved in CRISP share a common interest in creating and manipulating effective synthetic imagery. To achieve this goal we focus on understanding how people perceive complex material, lighting and shape, on developing new rendering algorithms based on this understanding, and on building interactive tools that enable users to efficiently specify the kind of image they wish to create. More specifically, we explore the following research directions :

  Perception: Images are generated from the interaction of lighting, material, and geometry. We evaluate how people perceive material, lighting, and geometry in realistic images such as photographs, and non realistic images such as drawings and paintings. This knowledge of human perception is essential for developing efficient rendering algorithms and interaction tools that focus on the most important perceptual features of an image.

  Rendering: We develop rendering algorithms that generate images that are plausible with respect to the user's intent and allocate resources on the visual effects that best contribute to perception.

Interaction: We facilitate the creation of material, lighting, and geometric effects in synthetic images by developing novel user interfaces for novice and professional users.

Our contributions have the potential to benefit different applications of image creation such as illustration (archeology, architecture, education); entertainment (video games, movies) and design (sketching, photograph editing). This research naturally falls in Inria's strategic objective of interacting with real and virtual worlds.

The CRISP collaboration has resulted in three publications this year in ACM Transactions on Graphics, two being in the SIGGRAPH proceeding. These publications explore the perception of materials in stylized images [11], the perception of distortions in image-based rendering [18] and vector drawing tools for depicting stylized materials [14]. Ongoing projects include those described in Sec. 5.3.9 and Sec. 5.1.3.

7.3.2. Informal International Partners

7.3.2.1. France-USA

Participants: Gaurav Chaurasia, Adrien Bousseau, George Drettakis.

Beyond CRISP, we have an ongoing collaboration with Yale University (Holly Rushmeier and Julie Dorsey), on weathering, and we are continuing this collaboration on stone aging.

We also have an ongoing collaboration with Adobe Research (Sylvain Paris) and MIT (Fredo Durand) on parallel image-proccessing languages and global illumination (Fredo Durand).

7.3.2.2. France-Germany

Participant: George Drettakis.

We collaborate with the Max-Planck-Institut, Germany, where P. Vangorp (previously at REVES) is now a PostDoc. We collaborate on perception techniques for rendering see publication [18].

7.3.2.3. France-Canada

Participant: Adrien Bousseau.

We collaborate with K. Singh (University of Toronto) and Alla Scheffer (U. British Columbia, Vancouver), on sketching techniques for designers (see Sec. 5.3.10).

7.3.2.4. France-Greece

Participant: George Drettakis.

As mentioned in Sec. 5.2.3 we are collaborating with the Technical University of Crete on visual attention, in the context. of the Ph.D. of George Koulieris, supervised by Prof. Katerina Mania and the Un. of Cottburg (D. Cunningham).

7.4. International Research Visitors

7.4.1. Visits of International Scientists

7.4.1.1. Visitors

We hosted several researchers this year:

- George Koulieris (Tech. Univ. of Crete), in February.
- Eugene Fiume (Univ. of Toronto), in June.
- Peter Vangorp (MPI Informatik), in June.
- Wendy McKay, Theophanis Tsandilas and Lora Oehlberg (Insitu) in July.
- Floraine Berthouzoz (Berkeley), in April.
- Marty Banks (Berkeley), in September.
- Belen Masia (Zaragoza), in October.
- Pierre-Yves Laffont (Brown), in November.
- Holly Rushmeier (Yale), in November.
- Erik Reinhard (Technicolor), in November.

7.4.1.2. Internships

Participant: Joan Sol Roo.

Subject: Geometry Upsampling for Real-Time Rendering of Refractive Objects

Date: from May 2013 until Aug 2013

Institution: National University of the Center of the Buenos Aires Province (Argentina)

Participant: Arunim Samat.

Subject: Approximate Reflection Computation

Date: from Jul 2013 until Aug 2013

Institution: IIT Delhi (India)

Participant: Kritarth Anand.

Subject: Free-Viewpoint Image Based Rendering from Images With Dynamic Objects

Date: from May 2013 until Jul 2013

Institution: IIT Delhi (India)

8. Dissemination

8.1. Scientific Animation

8.1.1. Program Committees

- George Drettakis served on the SIGGRAPH 2013 program committee.
- Adrien Bousseau served on the Eurographics program committee, on the Eurographics Symposium on Rendering program committee and on the NPAR program committee.
- Christian Richardt participated in the Expressive 2013 program committee and created the website: http://www.cl.cam.ac.uk/conference/expressive-2013/.

8.1.2. Community Service

G. Drettakis chaired the Eurographics (EG) Awards Committee until May this year, and chairs the EG Working group on Rendering and is part of the EG conference steering committee and the Advisory Board for Eurographics 2013. He also served on several ad-hoc ACM and IEEE Committees, concerning policy and EIC searches. He is an associate editor-in-chief for IEEE Transactions on Computer Graphics and Visualization. He also heads the steering committee of the Gouraud-Phong Immersive Space at Sophia-Antipolis.

8.2. Teaching - Supervision - Juries

8.2.1. Teaching

- Masters: G. Drettakis organizes and teaches Computer Graphics at the ECP (Paris) (9h), A. Bousseau teaches 3h at the same course. G. Drettakis teaches 6h in the MAPI M1 Module (Jeux Video), A. Bousseau teaches 3h in this program. A. Bousseau organizes and teaches the Computer Graphics module in the Master 1 International in Computer Science of University Nice Sophia Antipolis (9h), G. Drettakis teaches 7.5h in this program.

8.2.2. Supervision

- PhD in progress: Rodirgo Ortiz-Cayon, Inpainting and Artifact Removal for Image Based Rendering, (started November 2013), Advisor George Drettakis.
- PhD in progress: Gaurav Chaurasia, Image-Based Rendering, since October 2010 (to defend February 2014), Advisor: George Drettakis
- PhD in progress: Emmanuelle Chapoulie, Gestures in VR, since October 2010 (to defend April 2014), Advisor: George Drettakis
- PhD in progress: Emmanuel Iarussi, Computer-Assisted Drawing, since October 2012, Advisor: George Drettakis, Co-advisor: Adrien Bousseau
- Internship: Julia Chatain (Ecole Polytechnique, from April 8 to Aug. 9, 2013). Depicting materials in vector graphics. Advisor: George Drettakis, Co-advisor: Adrien Bousseau, Jorge Lopez Moreno
- Internship: Uditha Kasthuriarachchi (Jul. to Sep. 2013). GPU-Based Refraction. Advisor: George Drettakis, Co-advisor: Stefan Popov

8.2.3. International Internship

- Joal Sol Roo (UNCPBA Argentina, from May 21 to Sept. 13, 2013). Advisor: Adrien Bousseau, Christian Richardt

8.2.4. Juries

George Drettakis served on Mauricio Delbracio's Ph.D. committee.

8.3. Popularization

8.3.1. Popularization, Invited Talks

- Adrien Bousseau presented the work on Coherent Intrinsic Images from Photo Collections (SIGGRAPH Asia 2012) in the Research Session of the FMX conference in Stuttgart.
- Adrien Bousseau presented the work on sketching tools for designers as a "Cafe'In" at the Inria Sophia Antipolis center and as an introduction to research for high school students from Nice (MathC2+) and Grasse. This work also appeared as an article on the Inria Sophia Antipolis website http://www.inria.fr/centre/sophia/actualites/des-outils-informatiques-pour-les-artistes-et-les-createurs
- Christian Richardt gave invited talks at UC Berkeley, Stanford, Nvidia Research (Santa Clara, California) and TU Berlin (Germany).

8.3.2. Demos

We gave demos (mainly of the immersive space) to high school students, representatives from the planned Museum of Cosquer in Marseille, Sup'Com Tunis, the companies Thales, Thales AleniaSpace and Optis.

8.3.3. Conferences and Invited Talks

- At SIGGRAPH 2013, Adrien Bousseau presented the work on Gloss Perception in Painterly and Cartoon Rendering [14]. Jorge Lopez Moreno also presented our work at the French chapter of SIGGRAPH in Paris. Our work also appeared on the Inria Sophia Antipolis website (http://www.inria.fr/centre/sophia/actualites/siggraph-2013-trois-equipes-du-centre-presentent-leurs-travaux).
- Adrien Bousseau and Emmanuel Iarussi attended the ACM CHI conference in Paris (May).
- Adrien Bousseau and Emmanuel Iarussi attended the ACM UIST conference in St Andrews (October) where Emmanuel Iarussi presented the paper on assisting drawing from photographs [19].
- Adrien Bousseau gave a talk on drawing tools for depicting stylized materials at the "Groupe de Travail sur le Rendu" in Paris http://gtrendu.blogspot.fr/.

- George Drettakis attended Eurographics 2013 in Girona Spain (May) and EGSR 2013 in Zaragoza Spain (June).
- Gaurav Chaurasia attended Eurographics Symposium on Rendering at Zaragoza (Spain) June 19-21, 2013.
- Gaurav Chaurasia presented the paper “Depth Synthesis and Local Warps for Plausible Image-based Navigation” at SIGGRAPH at Annaheim (CA, USA) between July 21-25, 2013.
- Gaurav Chaurasia gave an invited talk titled “Algorithms and Perception for Interactive Free-Viewpoint Image-Based Navigation” at Microsoft Research Redmond, USA on July 18, 2013. He was hosted by Dr. Eyal Ofek in the Xtreme Computing Group.
- Gaurav Chaurasia visited MIT CSAIL between Aug 1, 2013 and Sept 30, 2013. He was hosted by Prof. Fredo Durand computer graphics group.
- Stefan Popov presented his paper at Eurographics 2013 in Girona, Spain
- Stefan Popov attended the Eurographics Symposium on Rendering 2013 in Zaragoza, Spain
- Christian Richardt attended the Berkeley-Inria-Stanford workshop in Stanford to represent the CRISP associate team.
- Christian Richardt attended Eurographics 2013 in Girona, Spain.
- Christian Richardt attended CVPR 2013 in Portland, Oregon, to present the "Megastereo" paper.
- Christian Richardt attended SIGGRAPH 2013 in Anaheim, California, and co-presented a paper with Peter Vangorp (Perception of Perspective Distortions in IBR).
- Emmanuel Iarussi participated in the ACM CHI Conference on Human Factors in Computing Systems. Paris, France - 27 April, 2 May 2013, and in the ACM Symposium on User Interface Software and Technology (UIST), St Andrews, UK, October 8-11, 2013.
- Sylvain Duchêne attended Siggraph Conference 2013 in Anaheim, USA, 18 hours Inria class on Computational Geometry Learning from David Cohen-Steiner, and university english course LAN02.

9. Bibliography

Major publications by the team in recent years

[1] N. BONNEEL, G. DRETTAKIS, N. TSINGOS, I. VIAUD-DELMON, D. JAMES. *Fast Modal Sounds with Scalable Frequency-Domain Synthesis*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", August 2008, vol. 27, n^{o} 3, http://www-sop.inria.fr/reves/Basilic/2008/BDTVJ08

[2] C. DACHSBACHER, M. STAMMINGER, G. DRETTAKIS, F. DURAND. *Implicit Visibility and Antiradiance for Interactive Global Illumination*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", August 2007, vol. 26, n^{o} 3, http://www-sop.inria.fr/reves/Basilic/2007/DSDD07

[3] O. DEUSSEN, C. COLDITZ, M. STAMMINGER, G. DRETTAKIS. *Interactive visualization of complex plant ecosystems*, in "Proceedings of the IEEE Visualization Conference", IEEE, October 2002, http://www-sop.inria.fr/reves/Basilic/2002/DCSD02/Vis2002Deussen.pdf

[4] F. DUGUET, G. DRETTAKIS. *Robust Epsilon Visibility*, in "Proceedings of ACM SIGGRAPH 2002", J. HUGHES (editor), ACM Press / ACM SIGGRAPH, July 2002, http://www-sop.inria.fr/reves/Basilic/2002/DD02/RobustEpsilonVisibility.pdf

[5] F. DUGUET, G. DRETTAKIS. *Flexible Point-Based Rendering on Mobile Devices*, in "IEEE Computer Graphics and Applications", July-August 2004, vol. 24, n° 4

[6] A. LAGAE, S. LEFEBVRE, G. DRETTAKIS, P. DUTRÉ. *Procedural Noise using Sparse Gabor Convolution*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", August 2009, vol. 28, n° 3, http://www-sop.inria.fr/reves/Basilic/2009/LLDD09

[7] A. RECHE, I. MARTIN, G. DRETTAKIS. *Volumetric Reconstruction and Interactive Rendering of Trees from Photographs*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", July 2004, vol. 23, n° 3, http://www-sop.inria.fr/reves/Basilic/2004/RMD04/FinalPaper_0167.pdf

[8] M. STAMMINGER, G. DRETTAKIS. *Perspective Shadow Maps*, in "Proceedings of ACM SIGGRAPH 2002", J. HUGHES (editor), ACM Press/ ACM SIGGRAPH, July 2002, http://www-sop.inria.fr/reves/Basilic/2003/SDD03/

[9] N. TSINGOS, E. GALLO, G. DRETTAKIS. *Perceptual Audio Rendering of Complex Virtual Environments*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", July 2004, vol. 23, n° 3, http://www-sop.inria.fr/reves/Basilic/2004/TGD04/

[10] B. WALTER, G. DRETTAKIS, D. GREENBERG. *Enhancing and Optimizing the Render Cache*, in "Proceedings of the Eurographics Workshop on Rendering", P. DEBEVEC, S. GIBSON (editors), ACM Press, June 2002, http://www-sop.inria.fr/reves/Basilic/2002/WDG02

Publications of the year

Articles in International Peer-Reviewed Journals

[11] A. BOUSSEAU, J. P. O'SHEA, F. DURAND, R. RAMAMOORTHI, M. AGRAWALA. *Gloss Perception in Painterly and Cartoon Rendering*, in "ACM Transactions on Graphics", April 2013, vol. 2, http://hal.inria.fr/hal-00811637

[12] G. CHAURASIA, S. DUCHENE, O. SORKINE-HORNUNG, G. DRETTAKIS. *Depth Synthesis and Local Warps for Plausible Image-based Navigation*, in "ACM Transactions on Graphics", June 2013, vol. 32, n° 3 [*DOI :* 10.1145/2487228.2487238], http://hal.inria.fr/hal-00907793

[13] P.-Y. LAFFONT, A. BOUSSEAU, G. DRETTAKIS. *Rich Intrinsic Image Decomposition of Outdoor Scenes from Multiple Views*, in "IEEE Transactions on Visualization and Computer Graphics", 2013, vol. in press, http://hal.inria.fr/hal-00761121

[14] J. LOPEZ-MORENO, P. STEFAN, A. BOUSSEAU, M. AGRAWALA, G. DRETTAKIS. *Depicting Stylized Materials with Vector Shade Trees*, in "ACM Transactions on Graphics", July 2013, vol. 32, n° 4, http://hal.inria.fr/hal-00828067

[15] A. ORZAN, A. BOUSSEAU, P. BARLA, H. WINNEMÖLLER, J. THOLLOT, D. SALESIN. *Diffusion Curves: A Vector Representation for Smooth-Shaded Images*, in "Communications of the ACM", July 2013, vol. 56, n° 7, pp. 101-108, Voir aussi article publié dans "ACM Transactions on Graphics" (doi:10.1145/1360612.1360691) en 2008 [*DOI :* 10.1145/2483852.2483873], http://hal.inria.fr/hal-00840848

[16] C. SUIED, G. DRETTAKIS, O. WARUSFEL, I. VIAUD-DELMON. *Auditory-visual virtual reality as a diagnostic and therapeutic tool for cynophobia*, in "Journal of Cybertherapy and Rehabilitation", 2013, vol. 16, n^o 2, pp. 145-152, http://hal.inria.fr/hal-00920871

[17] M. TAFFOU, R. GUERCHOUCHE, G. DRETTAKIS, I. VIAUD-DELMON. *Auditory-Visual Aversive Stimuli Modulate the Conscious Experience of Fear*, in "Multisensory Research", 2013, vol. 26, n^o 4, pp. 347-370, http://hal.inria.fr/hal-00920876

[18] P. VANGORP, C. RICHARDT, E. A. COOPER, G. CHAURASIA, M. S. BANKS, G. DRETTAKIS. *Perception of Perspective Distortions in Image-Based Rendering*, in "ACM Transactions on Graphics", July 2013, vol. 32, n^o 4, pp. 58:1-11 [*DOI :* 10.1145/2461912.2461971], http://hal.inria.fr/hal-00907863

International Conferences with Proceedings

[19] E. IARUSSI, A. BOUSSEAU, T. TSANDILAS. *The Drawing Assistant: Automated Drawing Guidance and Feedback from Photographs*, in "ACM Symposium on User Interface Software and Technology (UIST)", St Andrews, United Kingdom, ACM, October 2013, http://hal.inria.fr/hal-00840853

[20] C. RICHARDT, Y. PRITCH, H. ZIMMER, A. SORKINE-HORNUNG. *Megastereo: Constructing High-Resolution Stereo Panoramas*, in "CVPR - International Conference on Computer Vision and Pattern Recognition - 2013", Portland, Oregon, United States, June 2013, pp. 1256-1263, Oral presentation (3.3% acceptance rate) [*DOI :* 10.1109/CVPR.2013.166], http://hal.inria.fr/hal-00907858

Research Reports

[21] E. CHAPOULIE, M. MARCHAL, E. DIMARA, M. ROUSSOU, J.-C. LOMBARDO, G. DRETTAKIS. , *Evaluation of Direct Manipulation using Finger Tracking for Complex Tasks in an Immersive Cube*, Inria, September 2013, n^o RT-0440, 23 p. , http://hal.inria.fr/hal-00857534

References in notes

[22] O. DEUSSEN, C. COLDITZ, M. STAMMINGER, G. DRETTAKIS. *Interactive visualization of complex plant ecosystems*, in "Proceedings of the IEEE Visualization Conference", IEEE, October 2002, http://www-sop.inria.fr/reves/Basilic/2002/DCSD02/Vis2002Deussen.pdf

[23] J. DORSEY, H. K. PEDERSEN, P. HANRAHAN. *Flow and Changes in Appearance*, in "ACM Computer Graphics (SIGGRAPH'96 Proceedings)", Aout 1996, pp. 411–420

[24] F. DUGUET, G. DRETTAKIS. *Robust Epsilon Visibility*, in "ACM Transactions on Computer Graphics (Proceedings of ACM SIGGRAPH 2002)", July 2002, http://www-sop.inria.fr/reves/Basilic/2002/DD02/RobustEpsilonVisibility.pdf

[25] F. DURAND, G. DRETTAKIS, C. PUECH. *The 3D Visibility Complex, a new approach to the problems of accurate visibility*, in "Rendering Techniques'96 (7th Eurographics Workshop on Rendering)", Springer Verlag, June 1996, pp. 245–257

[26] F. DURAND, G. DRETTAKIS, C. PUECH. *The Visibility Skeleton: A Powerful and Efficient Multi-Purpose Global Visibility Tool*, in "ACM Computer Graphics (SIGGRAPH'97 Conference Proceedings)", Aout 1997

[27] F. DURAND, G. DRETTAKIS, C. PUECH. *Fast and Accurate Hierarchical Radiosity Using Global Visibility*, in "ACM Transactions on Graphics", April 1999, vol. 18, pp. 128–170

[28] T. FUNKHOUSER, I. CARLBOM, G. ELKO, G. PINGALI, M. SONDHI, J. WEST. *A Beam Tracing Approach to Acoustic Modeling for Interactive Virtual Environments*, in "ACM Computer Graphics (SIGGRAPH'98 Proceedings)", July 1998

[29] T. FUNKHOUSER, N. TSINGOS, I. CARLBOM, G. ELKO, M. SONDHI, J. WEST. *Modeling Sound Reflection and Diffraction in Architectural Environments with Beam Tracing*, in "Forum Acusticum", September 2002

[30] G. MILLER. *Efficient Algorithms for Local and Global Accessibility Shading*, in "ACM Computer Graphics (SIGGRAPH'94 Proceedings)", July 1994, pp. 319–326

[31] N. TSINGOS, T. FUNKHOUSER, I. CARLBOM. *Modeling Acoustics in Virtual Environments Using the Uniform Theory of Diffraction*, in "ACM Computer Graphics (SIGGRAPH 2001 Proceedings)", July 2001

[32] N. TSINGOS. *Artifact-free asynchronous geometry-based audio rendering*, in "ICASSP'2001", May 2001

[33] B. WALTER, G. DRETTAKIS, S. PARKER. *Interactive Rendering using the Render Cache*, in "Rendering Techniques'99 (10th Eurographics Workshop on Rendering)", Springer Verlag, June 1999