Activity Report 2012

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: July 01, 2002 .

1. Members

Research Scientists
George Drettakis [Team Leader, Senior Researcher (DR) Inria, HdR]
Adrien Bousseau [Junior Researcher (CR) Inria]

External Collaborator
Ares Lagae [FWO Leuven, Belgium]

Engineers
Adrien David [Engineer ADT, 2 year contract Inria-funded, until July]
Rachid Guerchouche [Expert Engineer, VERVE European project, since February]

PhD Students
Emmanuelle Chapoulie [BDO scholarship]
Gaurav Chaurasia [MESR scholarship]
Sylvain Duchêne [VERVE European project]
Pierre-Yves Laffont [Inria CORDIS scholarship, until September]
Emmanuel Iarussi [ANR DRAO, since October]

Post-Doctoral Fellows
Jorge Lopez Moreno [Autodesk and VERVE European project, since February]
Stefan Popov [ANR ALTA, since October]
Christian Richardt [Inria AAR, since October]
Charles Verron [Inria AAR, until August]

Administrative Assistant
Sophie Honnorat [Administrative Assistant (AI) Inria, part-time]

Others
Felicitas Hetzelt [Intern, University of Erlangen, Germany, VERVE, March-August]
Emmanuel Iarussi [Intern, University of Buenos Aires, Argentina, March-August]
Laurent Lefebvre [Intern, Epitech, ANR ALTA, March-July]
Loic Sévèque [Intern, Polytech'Nice-Sophia, VERVE, May-September]

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

An image from the paper Optimizing Environment Maps for Material Depiction – by Emmanuelle Chapoulie and Adrien Bousseau – was selected to appear on the front cover of the 8 issues of the journal Computer Graphics Forum 2012.

Jorge Lopez Moreno was honored with the Outstanding Doctoral Thesis Award at Universidad de Zaragoza.

This year has been particularly productive for our group, with seven publications in the top journals of our field (ACM TOG, IEEE TVCG, CGF) [15], [20], [19], [14], [13], [17], [16]

3. Scientific Foundations

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" [37], 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 [27], but also simpler methods such as accessibility maps [34]. 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 [35] 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 [31], [30], [29] 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

Figure 1. Plausible rendering of an outdoors scene containing points, lines and polygons [26], 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.

of generators (vertices and edges); this approach has been published in the past [28]. Lazy evaluation, as well 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 [32] [35]), 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 [33] 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 [36].

4. Software

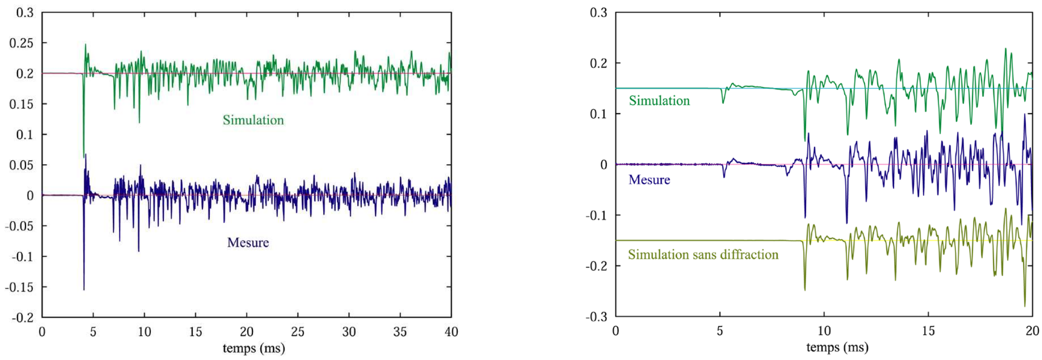

Figure 2. A comparison between a measurement (left) of the sound pressure in a given location of the "Bell Labs Box", a simple test environment built at Bell Laboratories, and a high-quality simulation based on a beam-tracing engine (right). Simulations include effects of reflections off the walls and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

4.1. RID: Rich Intrinsic Decomposer

Participants: Pierre-Yves Laffont, Adrien Bousseau, George Drettakis.

We developed a software platform to perform rich intrinsic decomposition methods from photographs of outdoor scenes, as described in [18]. 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.

4.2. Imerse: Inria Multi-Environment Realistic Simulation Engine

Participants: Adrien David, George Drettakis.

In the context of the ADT Interact3D and the ARC NIEVE, we developed Imerse, a middleware to be used as a VR engine, helping in the implementation of realistic simulations for immersive installations. Imerse provides a wrapper to OSG's (OpenSceneGraph) deep scene graph and its traversals abilities into an abstracted collection of high level objects which directly represent realistic entities (such as indoor elements, machines and realistic characters). It provides capacities such as skeletal animations or spatialized audio by interfacing with APF, while its clear composite pattern allows implementing more behaviors easily.

Finally, a generic design based on triggers and functors lets the final user implement complex scenarios of VR applications with the feeling of writing a script in C++. Applications developed on top of Imerse plug transparently into osgVR developed in the DREAM group (i.e., the research support development group of our Inria center). We are using osgVR to render OSG's scene graph in a distributed manner, since rendering clusters are available in an increasing number of installations. osgVR is a software layer developed by the DREAM research support group, ensuring synchronization and events/inputs distribution among a list of rendering slaves. These two libraries are available on GForge.

4.3. 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. 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, a previous 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 also been critical in establishing collaborations in the context of various grant proposals (EU and national).

4.4. 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.

5. New Results

5.1. Plausible Image Rendering

5.1.1. Rich Intrinsic Image Decomposition of Outdoor Scenes from Multiple Views

Participants: Pierre-Yves Laffont, Adrien Bousseau, George Drettakis.

Intrinsic image decomposition aims at separating photographs into independent reflectance and illumination layers. We show that this ill-posed problem can be solved by using multiple views of the scene from which we derive additional constraints on the decomposition.

Our first method uses pictures from multiple views at a *single time of day* to automatically reconstruct a 3D point cloud of an outdoor scene. Although this point cloud is sparse and incomplete, it is sufficient to compute plausible sky and indirect illumination at each oriented 3D point, given an environment map that represents incoming distant radiance. We introduce an optimization method to estimate sun visibility over the point cloud, which compensates for the lack of accurate geometry and allows the extraction of precise cast shadows. We finally use image-guided propagation algorithms to propagate the illumination computed over the sparse point cloud to every pixel, and to separate the illumination into distinct sun, sky, and indirect components. This *rich intrinsic image decomposition* enables advanced image manipulations, illustrated in Figure 3.

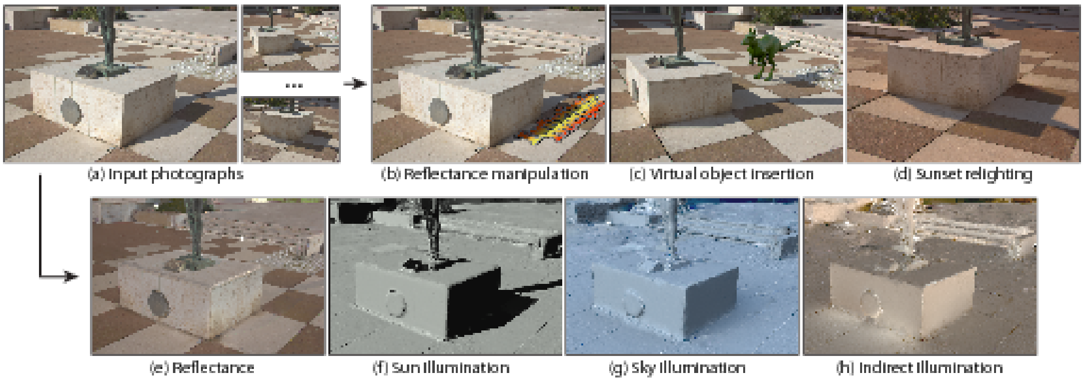

Figure 3. Starting from multiple views of the scene (a), our method decomposes photographs into four intrinsic layers — the reflectance (e), the illumination due to sun (f), the illumination due to sky (g) and the indirect illumination (h). Each layer can then be manipulated independently for advanced image editing applications (b-d).

This work has led to the RID software (Section 4.1) and to a technology transfer agreement with Autodesk (Section 6.1.1.1). A paper will be published in the IEEE Transactions on Visualization and Computer Graphics journal [18] (in press). It has also been presented at SIGGRAPH 2012 in the Poster and Talk sessions [22].

5.1.2. Coherent Intrinsic Images from Photo Collections

Participants: Pierre-Yves Laffont, Adrien Bousseau, George Drettakis.

We propose a second method to compute intrinsic images in the presence of varying lighting conditions. Our method exploits the rich information provided by *multiple viewpoints and illuminations* in an image collection to process complex scenes without user assistance, nor precise and complete geometry. Such collections can be gathered from photo-sharing websites, or captured indoors with a light source which is moved around the scene.

We use multi-view stereo to automatically reconstruct 3D points and normals, from which we derive relationships between reflectance values at different locations, across multiple views, and consequently across different lighting conditions. In addition, we propose an optimization approach which enforces coherent reflectance in all views of a scene.

The resulting *coherent intrinsic images* enable image-based illumination transfer between photographs of the collection, as illustrated in Figure 4.

This work is a collaboration with Frédo Durand (MIT) and Sylvain Paris (Adobe), and started with a visit of Pierre-Yves Laffont at MIT during Summer 2011. It has been published in the ACM Transactions on Graphics journal [19], and has been presented at SIGGRAPH Asia 2012.

5.1.3. Intrinsic Images by Clustering

Participant: Jorge Lopez Moreno.

Decomposing an input image into its intrinsic illumination and reflectance components is a long-standing ill-posed problem. We present a novel algorithm that requires no user strokes and works on a single image. Based on simple assumptions about its reflectance and luminance, we first find clusters of similar reflectance in the image, and build a linear system describing the connections and relations between them. Our assumptions are less restrictive than widely-adopted Retinex-based approaches, and can be further relaxed in conflicting situations. The resulting system is robust even in the presence of areas where our assumptions do not hold. We show a wide variety of results, including natural images, objects from the MIT dataset and texture images, along with several applications, proving the versatility of our method (see Figure 5).

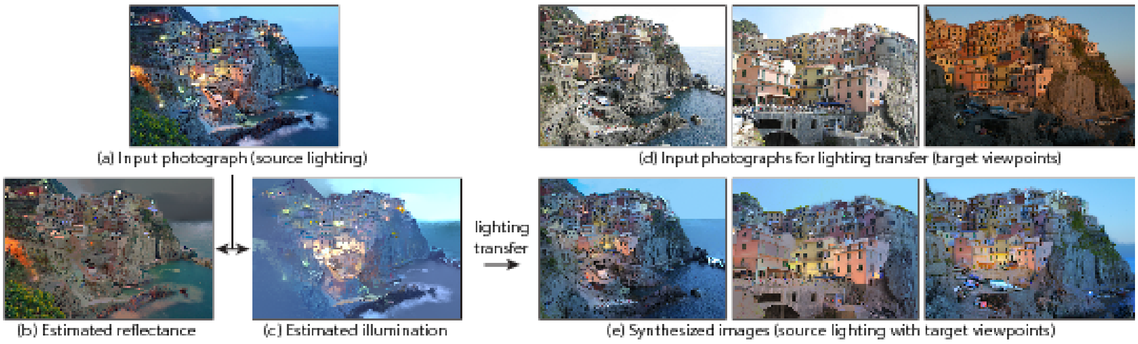

Figure 4. Our method automatically decomposes each image of a photo collection into reflectance and illumination (a-c). Transferring the illumination layer (c) to other viewpoints (d) yields synthetic images with novel viewpoint/lighting combinations (e).

Figure 5. Decomposition by our method of the input image (left) into illumination (center) and reflectance (right) components.

This work is a collaboration with Elena Garces, Adolfo Munoz and Diego Gutierrez from University of Zaragoza (Spain). The work was published in an special issue of the journal Computer Graphics Forum and presented at the Eurographics Symposium on Rendering 2012 [16].

5.1.4. *Relighting for Image Based Rendering*

Participants: Sylvain Duchêne, Jorge Lopez Moreno, Stefan Popov, George Drettakis.

Image-based rendering generates realistic virtual images from a small set of photographs. However, while current methods can simulate novel viewpoints from the input pictures, they cannot produce novel illumination conditions that differ from the lighting at the time of capture. The goal of this project is to provide such relighting capabilities. Our method first rely on multi-view stereo algorithms to estimate a coarse geometry of the scene. This geometry is often innacurate and incomplete. We complement it with image-based propagation algorithms that fill-in the missing data using the high-resolution input pictures. This combination of geometric and image-based cues allows us to generate plausible shadow motion and simulate novel sun directions.

5.1.5. *Depth Synthesis and Local Warps for Plausible Image-based Navigation*

Participants: Gaurav Chaurasia, Sylvain Duchêne, George Drettakis.

Modern multi-view stereo algorithms can estimate 3D geometry from a small set of unstrcutured photographs. However, the 3D reconstruction often fails on vegetation, vehicles and other complex geometry present in everyday urban scenes. We introduce a new Image-Based Rendering algorithm that is robust to unreliable geometry. Our algorithm segments the image into superpixels, *synthesizes* depth in superpixels with missing depth, warps them using a shape-preserving warp and blends them to create real-time plausible novel views for challenging target scenes, resulting in convincing immersive navigation experience.

This work is in collaboration with Dr. Olga Sorkine at ETH Zürich. and has been submitted to ACM Transactions on Graphics.

5.1.6. *Perception of Slant for Image-Based Rendering*

Participants: Christian Richardt, Peter Vangorp, George Drettakis.

Image-based rendering can create images with a high level of realism using simple geometry. However, as soon as the viewer moves away from the correct viewpoint, the image appears deformed. This work investigates the parameters which influence the perception of these image deformations. We propose a novel model of slant perception, which we validate using psychophysical experiments.

This work is a collaboration with Peter Vangorp at MPI Informatik, and Emily Cooper and Martin Banks from the University of California, Berkeley; in the context of the Associate Team CRISP (see also Section 7.3.1.1).

5.1.7. *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 uppon 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 Università di Roma.

5.1.8. *Example-Based Fractured Appearance*

Participants: Carles Bosch, George Drettakis.

A common weathering effect is the appearance of cracks due to material fractures. Previous exemplar-based aging and weathering methods have either reused images or sought to replicate observed patterns exactly. We propose an approach to exemplar-based modeling that creates weathered patterns by matching the statistics of fracture patterns in a photograph. We conducted a user study to determine which statistics are correlated to visual similarity and how they are perceived by the user. We describe a physically-based fracture model capable of producing similar crack patterns at interactive rates and an optimization method to determine its parameters based on key statistics of the exemplar. Our approach is able to produce a variety of fracture effects from simple crack photographs at interactive rates, as shown in Figure 6.

Figure 6. Application of our example-based fracturing method on different scenes. Photographs of input fracture patterns are shown in the insets.

This work is a collaboration with Loeiz Glondu, Maud Marchal and George Dumont from IRISA-INSA/Inria Rennes - Bretagne Atlantique, Lien Muguercia from the University of Girona, and Holly Rushmeier from Yale University. The work was published in the Computer Graphics Forum journal and presented at the 23rd Eurographics Symposium on Rendering [17].

5.1.9. Real-Time Rendering of Rough Refraction

Participant: Adrien Bousseau.

We propose an algorithm to render objects made of transparent materials with rough surfaces in real-time, under all-frequency distant illumination. Rough surfaces cause wide scattering as light enters and exits objects, which significantly complicates the rendering of such materials. We present two contributions to approximate the successive scattering events at interfaces, due to rough refraction: First, an approximation of the Bidirectional Transmittance Distribution Function (BTDF), using spherical Gaussians, suitable for real-time estimation of environment lighting using pre-convolution; second, a combination of cone tracing and macro-geometry filtering to efficiently integrate the scattered rays at the exiting interface of the object. We demonstrate the quality of our approximation by comparison against stochastic ray-tracing (see Figure 7). Furthermore we propose two extensions to our method for supporting spatially varying roughness on object surfaces and local lighting for thin objects.

This work is a collaboration with Charles De Rousiers , Kartic Subr, Nicolas Holzschuch from Inria Grenoble, and Ravi Ramamoorthi from UC Berkeley in the context of the Associate Team CRISP (see also Section 7.3.1.1). A paper describing the method was published in the IEEE Transactions on Visualization and Computer Graphics journal [14].

5.1.10. Gabor Noise by Example

Participants: Ares Lagae, George Drettakis.

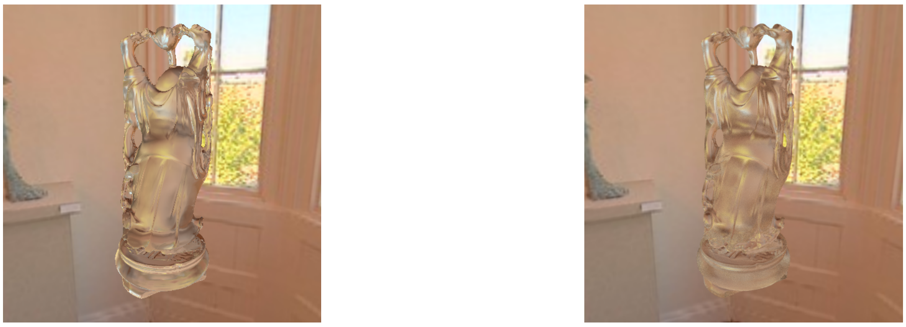

(a) Ground truth (b) Our method

Figure 7. Compared to an expensive ray-traced reference (a), our method produces plausible results in real-time (b).

Procedural noise is a fundamental tool in Computer Graphics. However, designing noise patterns is hard. In this project, we propose *Gabor noise by example*, a method to estimate the parameters of bandwidth-quantized Gabor noise, a procedural noise function that can generate noise with an arbitrary power spectrum, from exemplar Gaussian textures, a class of textures that is completely characterized by their power spectrum (see Figure 8).

More specifically, we introduce (i) bandwidth-quantized Gabor noise, a generalization of Gabor noise to arbitrary power spectra that enables robust parameter estimation and efficient procedural evaluation; (ii) a robust parameter estimation technique for quantized-bandwidth Gabor noise, that automatically decomposes the noisy power spectrum estimate of an exemplar into a sparse sum of Gaussians using non-negative basis pursuit denoising; and (iii) an efficient procedural evaluation scheme for bandwidth-quantized Gabor noise, that uses multi-grid evaluation and importance sampling of the kernel parameters. Gabor noise by example preserves the traditional advantages of procedural noise, including a compact representation and a fast on-the-fly evaluation, and is mathematically well-founded.

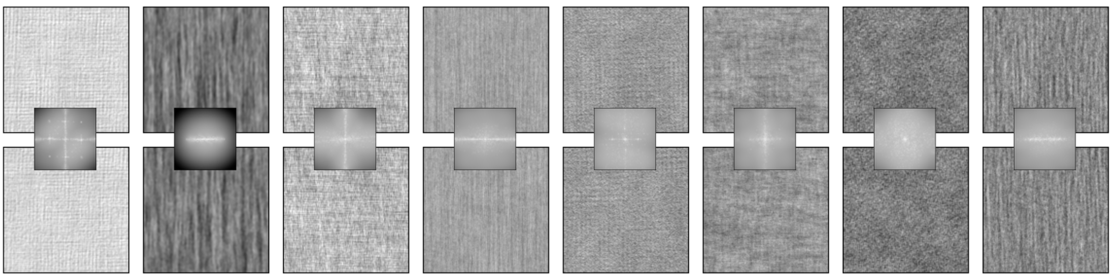

Figure 8. Gabor noise by example is a method to estimate the parameters of bandwidth-quantized Gabor noise, a procedural noise function that can generate noise with an arbitrary power spectrum, from exemplar Gaussian textures, a class of textures that is completely characterized by their power spectrum. (row 1) Gaussian texture. (row 2) Procedural noise. (insets) Estimated power spectrum.

This work is a collaboration with Bruno Galerne from MAP5, Université Paris Descartes and CNRS, Sorbonne Paris Cité; Ares Lagae from KU Leuven; and Sylvain Lefebvre from the ALICE project team, Inria Nancy - Grand Est. This work was presented at SIGGRAPH 2012 and published in ACM Transactions on Graphics [15].

5.1.11. Structured Gabor noise

Participants: Gaurav Chaurasia, Ares Lagae, George Drettakis.

Current procedural noise synthesis techniques [15] are limited to Gaussian random field textures. This project aims to generalize procedural noise to a broader class of structured textures.

This work is in collaboration with Dr. Ares Lagae (Katholieke Universiteit Leuven, Belgium), Dr. Bruno Galerne (Université Paris Descartes) and Prof. Ravi Ramamoorthi (UC Berkeley), in the contect of the Associate Team CRISP (Section 7.3.1.1).

5.1.12. 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 9). 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.

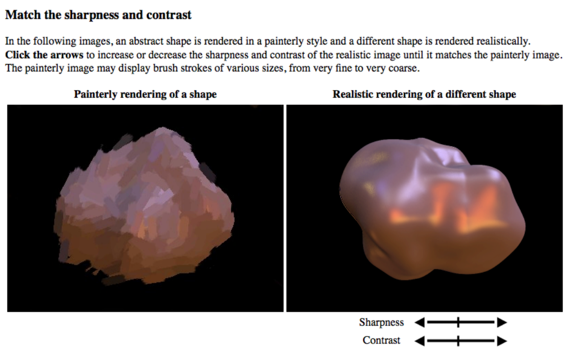

Figure 9. The experimental task used for studying gloss perception in stylized images.

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 like 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 7.3.1.1) and Frédo Durand from MIT. It will be published in ACM Transactions on Graphics 2013 [12] (in press).

5.2. Interaction and Design for Audiovisual Virtual Environments

5.2.1. Auditory-visual integration of emotional signals in a virtual environment for cynophobia

Participants: Emmanuelle Chapoulie, Adrien David, Rachid Guerchouche, George Drettakis.

Cynophobia (dog phobia) has both visual and auditory relevant components. In order to investigate the efficacy of virtual reality exposure-based treatment for cynophobia, we studied the efficiency of auditory-visual environments in generating presence and emotion. We conducted an evaluation test with healthy participants sensitive to cynophobia in order to assess the capacity of auditory-visual virtual environments to generate fear reactions. Our application involves both high fidelity visual stimulation displayed in an immersive space and 3D sound. This specificity enables us to present and spatially manipulate fearful stimuli in the auditory modality, the visual modality and both.

We conducted a study where participants were presented with virtual dogs in realistic environments. Dogs were presented in a progressive manner, from unimodal and static to audiovisual and dynamic. Participants were also submitted a Behavioral Assessment Test at the beginning and end of the experiment where they were presented a virtual dog walking towards them step by step until it was extremely close. Finally, they completed several questionnaires and were asked to comment on their experience. The participants reported higher anxiety levels in response to auditory-visual stimuli compared to unimodal stimuli. Our results strongly suggest that manipulating auditory-visual integration might be a good way to modulate affective reactions and that auditory-visual VR are a promising tool for the treatment of cynophobia.

This work is a collaboration with Marine TAFFOU and Isabelle VIAUD-DELMON from IRCAM, in the context of ARC NIEVE (see also Section 7.1.4). The work was published in the Annual Review of Cybertherapy and Telemedicine in 2012.

5.2.2. Procedural audio modeling for particle-based environmental effects

Participants: Charles Verron, George Drettakis.

In this project we proposed a sound synthesizer dedicated to particle-based environmental effects, for use in interactive virtual environments. The synthesis engine is based on five physically-inspired basic elements which we call sound atoms, that can be parameterized and stochastically distributed in time and space. Based on this set of atomic elements, models are presented for reproducing several environmental sound sources. Compared to pre-recorded sound samples, procedural synthesis provides extra flexibility to manipulate and control the sound source properties with physically-inspired parameters. The controls are used to simultaneously modify particle-based graphical models, resulting in synchronous audio/graphics environmental effects. The approach is illustrated with three models, that are commonly used in video games: fire, wind, and rain. The physically-inspired controls simultaneously drive graphical parameters (e.g., distribution of particles, average particles velocity) and sound parameters (e.g., distribution of sound atoms, spectral modifications) as illustrated on Figure 10 for fire. The joint audio/graphics control results in a tightly-coupled interaction between the two modalities that enhances the naturalness of the scene.

The work was presented at the 133rd AES convention in October 2012 [23].

5.2.3. Perception of crowd sounds

Participants: Charles Verron, George Drettakis.

Simulating realistic crowd scenes is an important challenge for virtual reality and games. Motion capture techniques allow to reproduce efficiently characters that look, move and sound realistic in virtual environments. However a huge amount of data is required to ensure that all agents behave differently in a big crowd. A common approach to solve this issue is to “clone" the same appearance, motion or sound several times, which can lead to perceived repetitions and break the realism of the scene. In this study we further investigate our perception of crowd scenes. Using a database of motions and sounds captured for 40 actors, along with a database of 40 different appearance templates, we propose an experimental framework to evaluate the perceptual degradations caused by clones. A particular attention is given to evaluate the influence of appearance,

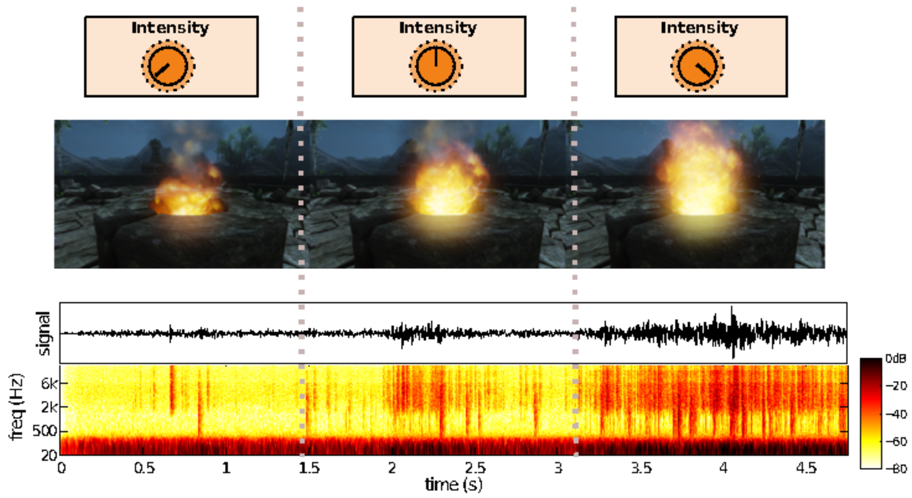

Figure 10. Audio/graphics high-level control of a fire. The control Intensity changes the rate and gain of noisy impacts, and the combustion noise of the fire sound model. Simultaneously, it controls the flame/smoke particle spawn rate for the graphics simulation.

motion and sound, either separately or in multimodal conditions. This study aims at providing useful insights on our perception of crowd scenes, and guidelines to designers in order to reduce the amount of resources to produce convincing crowd scenes.

This ongoing project is a collaboration between Inria, CNRS-LMA (Marseille, France) and Trinity College (Dublin, Ireland).

5.2.4. Walking in a Cube: Novel Metaphors for Safely Navigating Large Virtual Environments in Restricted Real Workspaces

Participants: Peter Vangorp, Emmanuelle Chapoulie, George Drettakis.

Immersive spaces such as 4-sided displays with stereo viewing and high-quality tracking provide a very engaging and realistic virtual experience. However, walking is inherently limited by the restricted physical space, both due to the screens (limited translation) and the missing back screen (limited rotation). Locomotion techniques for such restricted workspaces should satisfy three concurrent goals: keep the user safe from reaching the translational and rotational boundaries; increase the amount of real walking; and finally, provide a more enjoyable and ecological interaction paradigm compared to traditional controller-based approaches.

We have proposed three novel locomotion techniques that attempt to satisfy these goals in innovative ways. We constrain traditional Wand locomotion by turning off the Wand controls for directions that can be reached by real walking instead, and we display warning signs when the user approaches the limits of the real workspace (Figure 11(a)). We also extend the Magic Barrier Tape paradigm with "blinders" to avoid rotation towards the missing back screen (Figure 11(b)). Finally, we introduce the "Virtual Companion", which uses a small bird to guide the user through virtual environments larger than the physical space (Figure 11(c,d)).

We evaluate the three new techniques through a user study with travel-to-target and path following tasks. The study provides insight into the relative strengths of each new technique for the three aforementioned goals. Specifically, if speed and accuracy are paramount, traditional controller interfaces augmented with our novel warning techniques may be more appropriate; if physical walking is more important, two of our paradigms, the extended Magic Barrier Tape and the Constrained Wand, should be preferred; and finally, fun and ecological criteria would favor the Virtual Companion.

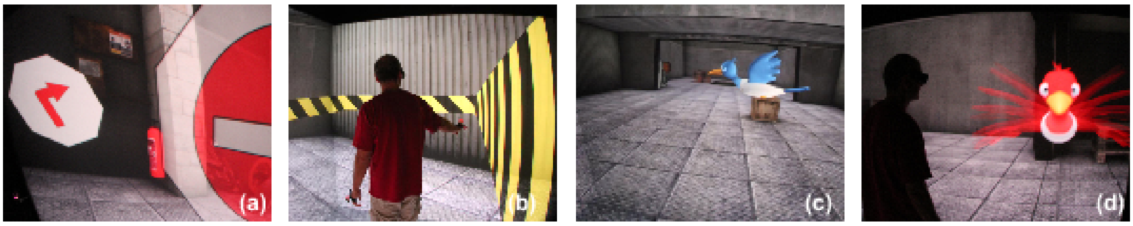

Figure 11. Screenshots illustrating the three novel locomotion techniques. From left to right: (a) Constrained Wand and signs: the "no-way" and "turn right" signs. (b) Extended Magic Barrier Tape: the tape and blinders. (c,d) Virtual Companion: the bird in "rest mode" (c) and "protection mode" (d).

This work is a collaboration with Gabriel Cirio, Maud Marchal and Anatole Lécuyer (VR4I project team, IRISA-INSA/Inria Rennes - Bretagne Atlantique) in the context of ARC NIEVE (see Section 7.1.4). The work was published in the special issue of the journal IEEE Transactions on Visualization and Computer Graphics (TVCG) [13], and presented at the IEEE Virtual Reality conference 2012.

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

Participants: Emmanuelle Chapoulie, Jean-Christophe Lombardo, 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. To do this, we develop a solution using 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 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 Evanthia Dimara and Maria Roussou from the University of Athens and with Maud Marchal from IRISA-INSA/Inria Rennes - Bretagne Atlantique. The work has been submitted to 3DUI 2013.

5.2.6. CrossShade: Shading Concept Sketches Using Cross-Section Curves

Participant: Adrien Bousseau.

We facilitate the creation of 3D-looking shaded production drawings from concept sketches. The key to our approach is a class of commonly used construction curves known as cross-sections, that function as an aid to both sketch creation and viewer understanding of the depicted 3D shape. In particular, intersections of these curves, or cross-hairs, convey valuable 3D information, that viewers compose into a mental model of the overall sketch. We use the artist-drawn cross-sections to automatically infer the 3D normals across the sketch, enabling 3D-like rendering (see Figure 12).

The technical contribution of our work is twofold. First, we distill artistic guidelines for drawing cross-sections and insights from perception literature to introduce an explicit mathematical formulation of the relationships between cross-section curves and the ge- ometry they aim to convey. We then use these relationships to develop an algorithm for estimating a normal field from cross-section curve networks and other curves present

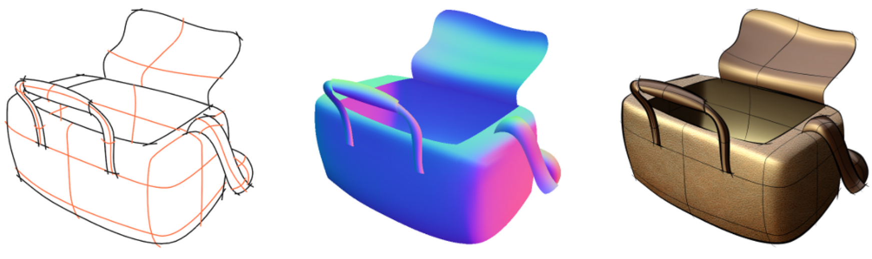

Figure 12. Concept sketches (a) frequently use cross-sections (drawn in orange) to convey 3D shape with just a handful of lines. We derive the mathematical properties of cross-section curves and leverage them to automatically estimate surface normals across the drawn objects (b). The resulting normal field allow users to shade the objects using a variety of shading styles and setups (c).

in concept sketches. We validate our formulation and algorithm through a user study and a ground truth normal comparison. These contributions enable us to shade a wide range of concept sketches with a variety of rendering styles.

This work is a collaboration with Cloud Shao and Karan Singh from the University of Toronto and Alla Sheffer from the University of British Columbia. It has been published at ACM Transactions on Graphics, proceedings of the SIGGRAPH 2012 conference.

5.2.7. CrossShape

Participant: Adrien Bousseau.

We facilitate the automatic creation of surfaced 3D models from design sketches that employ a commonly drawn network of cross-section curves. Our previous method generates 3D renderings of input sketches by creating a 3D surface normal field that interpolates the sketched cross-sections. This normal field however, incorporates the inevitable inaccuracy of sketched curves, making it inappropriate for 3D surface construction.

Successful construction of the 3D surface perceived from sketches requires cross-section properties and other perceived curve relationships such as symmetry and parallelism, to be met precisely. We present a novel formulation where these geometric constraints are satisfied while minimizing the difference between the sketch and the 3D cross-sections projected on it. We validate our approach by producing accurate surface reconstructions of existing 3D models represented using a network of cross-sections as well on a variety of sketch input. Finally we illustrate our surfacing solution within an interactive sketch based modeling framework.

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

5.2.8. Computer-assisted drawing

Participants: Emmanuel Iarussi, Adrien Bousseau.

A major challenge in drawing from observation is to trust what we *see* rather than what we *know*. Drawing books and tutorials provide simple techniques to gain consciousness of the shapes that we observe and their relationships. Common techniques include drawing simple geometrical shapes first – also known as *blocking in* – and checking for alignments and equal proportions. While very effective, these techniques are usually illustrated on few examples and it takes significant effort to generalize them to an arbitrary model. In addition, books and tutorials only contain static instructions and cannot provide feedback to people willing to practice drawing.

In this project, we develop an interactive drawing tool that assists users in their practice of common drawing techniques. Our *drawing assistant* helps users to draw from any model photograph and provides corrective feedback interactively.

This work is a collaboration with Theophanis Tsandilas from the InSitu project team, Inria Saclay - Ile de France, in the context of the ANR DRAO project (see Section 7.1.2).

5.2.9. Depicting materials in vector graphics

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

Vector drawing tools like Illustrator and InkScape enjoy great popularity in illustration and design because of their flexibility, directness and distinctive look. Within such tools, skillful artists depict convincing material and lighting effects using 2D vector primitives like gradients and paths. However, it takes significant expertise to convey plausible material appearance in vector drawings. Instead, novice users often fill-in regions with a constant color, sacrifying plausibility for simplicity. In this project we present the first vector drawing tool that automates the depiction of material appearance. Users can use our tool to either fill-in regions automatically, or to generate an initial set of vector primitives that they can refine at will.

This work is a collaboration with Maneesh Agrawala from the University of Berkeley in the context of the Associate Team CRISP (see Section 7.3.1.1).

5.2.10. Gradient Art: Creation and Vectorization (survey)

Participant: Adrien Bousseau.

We survey the main two categories of methods for producing vector gradients. One is mainly interested in converting existing photographs into dense vector representations. By vector it is meant that one can zoom infinitely inside images, and that control values do not have to lie onto a grid but must represent subtle color gradients found in input images. The other category is tailored to the creation of images from scratch, using a sparse set of vector primitives. In this case, we still have the infinite zoom property, but also an advanced model of how space should be filled in-between primitives, since there is no input photograph to rely on. These two categories are actually extreme cases, and seem to exclude each other: a dense representation is difficult to manipulate, especially when one wants to modify topology; a sparse representation is hardly adapted to photo vectorization, especially in the presence of texture. Very few methods lie in the middle, and the ones that do require user assistance.

We published our survey in the book *Image and Video based Artistic Stylization* [25] editied by Springer. The survey was written in Collaboration with Pascal Barla from the MANAO project team, Inria Bordeaux - Sud Ouest, in the context of the ANR DRAO project (see Section 7.1.2).

6. Bilateral Contracts and Grants with Industry

6.1. Grants with Industry

6.1.1. Industrial Contracts and Donations

6.1.1.1. Autodesk

Participants: Pierre-Yves Laffont, Adrien Bousseau, George Drettakis.

We signed a technology transfer agreement with Autodesk RID technology on single-lighting condition intrinsic images. 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.

6.1.1.2. Adobe

Participants: Pierre-Yves Laffont, Adrien Bousseau, George Drettakis.

In the context of our collaboration with Adobe (project with S. Paris and F. Durand from MIT), we have received a cash donation in support of our research and software donations of Adobe CS6 Creative Suite.

7. Partnerships and Cooperations

7.1. National Initiatives

7.1.1. ANR ALTA

Participants: Emmanuelle Chapoulie, Adrien David, 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 solutions that can be used in Virtual Reality solutions, for example in the context of the immersive space.

7.1.2. ANR DRAO

Participants: Emmanuel Iarussi, Adrien Bousseau.

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 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.

7.1.3. ADT Interact3D

Participants: Adrien David, George Drettakis.

This ADT involves half time software development for ARC NIEVE, and the other half general support to the new Immersive Space Gouraud-Phong in Sophia-Antipolis (supervised by Jean-Christophe Lombardo of the DREAM service). The main contribution was the complete rewrite of our VR application environment with the development of the Imerse software. This platform will allow first experiments, and the development of a generic Virtual Reality framework addressing neuroscience/psychology applications. This generic platform is based on osgVR which aims at a high-quality context abstraction to be usable in several domains, as well as distributed rendering capacities. These improvements, deployable for a variety of applications to come, are tightly coupled with the current ARC NIEVE, thus contributing to its implementation. Future prospects for the ADT Interact 3D include developing novel multimodal interaction techniques for example for gesture-based interaction etc.

7.1.4. ARC NIEVE: Navigation and Interfaces in Emotional Virtual Environments

Participants: Peter Vangorp, Adrien David, George Drettakis, Gaurav Chaurasia, Emmanuelle Chapoulie.

The goal of this joint research project is to develop and evaluate improved interfaces for navigation in immersive virtual environments (VEs) such as the 4-wall stereoscopic ISpace system in the Immersive Space Gouraud-Phong.

There is evidence of significant overlap in brain structures related to spatial memory and orientation and those related to emotion. We examine the influence of high-quality 3D visual and auditory stimuli on the emotions evoked by the virtual environment. Our study focuses on the phobia of dogs as a way to modulate emotion in audiovisual VEs (see Figure 13).

Navigation in VEs involves the use of different views, i.e., egocentric ("first person") and allocentric ("bird's eye") views during navigation tasks. We study appropriate visual representations for each view (for example, the level of realism ranging from abstract map-like rendering for top-down views to photorealistic rendering for first-person views), and appropriate transitions between the different views.

We develop an appropriate methodology to evaluate such navigation interfaces in stressful environments, based on the insights gained by the emotion modulation study in phobic settings. This novel methodology can be seen as a "stress-test" for navigation interfaces: if the navigation interfaces developed are successful even in stressful setups, they will definitely be successful under "normal conditions".

ARC NIEVE has resulted in several publications this year: [21], [13].

Figure 13. A person immersed in a virtual environment where the behaviors of several dogs will evoke different levels of anxiety.

This is a joint research project with Isabelle Viaud-Delmon (IRCAM, CNRS), Anatole Lécuyer and Maud Marchal (VR4I project team, IRISA-INSA/Inria Rennes - Bretagne Atlantique), and Jean-Christophe Lombardo (DREAM / Inria Sophia Antipolis). Interact3D (Section 7.1.3) is associated with this ARC.

7.1.5. National French Bilateral Collaboration

We have ongoing collaborations with Maud Marchal and Anatole Lécuyer (VR4I project team, IRISA-INSA/Inria Rennes - Bretagne Atlantique) [13], [17], and Bruno Galerne (ENST/ENS Cachan) [15].

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/

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 first year period (October 2011 - September 2012), the consortium focused its effort on the following main actions:

- Designing scenarios for different situations, 3 main scenarios were designed:
  1. DogPhobia scenario (for phobias),
  2. Kitchen scenario (for Alzheimer patients),
  3. MeMoVE (for memory complaints).
- Ethical approvals submission for the different scenarios.
- Conducting different experiments in the context of these different scenarios, especially DogPhobia scenario.
- Development and adaptation of different technologies in order to implement the scenarios:
  - Image based rendering (IBR) for virtual realistic environment modeling,
  - Emotive avatars,
  - Crowds simulation,
  - Realistic human skin rendering.
- Development of different technical tool:
  - Virtual environment for mobile device serious game (Kitchen scenario),
  - Porting the IBR to immersive space for the MeMoVE scenario,
  - Adapting the partner's technologies to the different platforms within the consortium.

The first year review of the VERVE project was hold on October 2nd, 2012, and the project were judged good and follows the defined plan.

7.3. International Initiatives

7.3.1. Inria Associate Teams

7.3.1.1. CRISP

Title: Creating and Rendering Images based on the Study of Perception

Inria principal investigator: Adrien Bousseau

International Partner (Institution - Laboratory - Researcher):

University of California Berkeley (United States) - Electrical Engineering and Computer Science - Maneesh Agrawala

Duration: 2011 - 2013

See also: 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 will 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 will explore the following research directions :

Perception: Images are generated from the interaction of lighting, material, and geometry. We will 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. We have started several projects on the perception of materials in realistic and non realistic images, with promising results.

Rendering: We will 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. Current projects on rendering include work on enhancing material variations in realistic and non realistic rendering.

Interaction: We will facilitate the creation of material, lighting, and geometric effects in synthetic images by developing novel user interfaces for novice and professional users. We are currently working on interfaces to draw object appearance and to relight photographs.

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.

7.4. Bilateral Collaborations

7.4.1. France-USA

Participants: Gaurav Chaurasia, Pierre-Yves Laffont, Adrien Bousseau, George Drettakis, Christian Richardt, Jorge Lopez-Moreno.

We have an ongoing collaboration with Yale University (Holly Rushmeier and Julie Dorsey), on weathering, resulting in the publication [17]. We continue this collaboration on stone aging.

We have an ongoing collaboration with Adobe Research (Sylvain Paris) and MIT (Fredo Durand) on intrinsic images for mutliple lighting conditions, resulting in the publication [19].

We also collaborate with M. Banks, R. Ramamoorthi and M. Agrawala from the University of California, Berkeley in the context of our CRISP associate team, resulting in the publications [14], [12]. Gaurav Chaurasia spent 6 weeks this summer at UC Berkeley in the context of this collaboration. Adrien Bousseau and George Drettakis also visited UC Berkeley for 3 days in August.

7.4.2. France-Switzerland

Participants: Gaurav Chaurasia, Sylvain Duchêne, George Drettakis.

We collaborate with O. Sorkine at ETH Zurich on image-based rendering, which resulted in a submission to ACM TOG.

7.4.3. France-Germany

Participant: George Drettakis.

We collaborate with the Max-Planck-Institut, Germany, where P. Vangorp is now a PostDoc. We collaborate on perception techniques for rendering and on interactions for virtual environments. This resulted in the following publication [13].

7.4.4. France-Spain

Participants: George Drettakis, Adrien Bousseau.

We collaborate with C. Bosch who is now at the University of Girona (Spain), on weathering.

7.4.5. France-Italy

Participant: Adrien Bousseau.

We collaborate with F. Pellacini from Sapienza Università di Roma on lightfield editing.

7.4.6. 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. This collaboration resulted in the publication [20] and in the 3 weeks visit of Xu Baoxuan (PhD student at U. British Columbia).

7.4.7. France-Belgium

Participant: George Drettakis.

We have continued the collaboration with A. Lagae from the Catholic University of Leuven, resulting in the publication [15].

7.5. International Research Visitors

7.5.1. Visits of International Scientists

We hosted several researchers this year:

- Maneesh Agrawala (Univ. of Berkeley), in May-June
- Brian Curless (Univ. of Washington), in October
- Eugene Fiume (Univ. of Toronto), in June
- Michael Gleicher(Univ. of Wisconsin), in June
- Diego Guttierez (Univ. of Zaragoza), in October
- Ares Lagae, (KU Leuven), in November
- Hendrik Lensch (Univ. of Ulm), in October
- Pierre Poulin (Univ. of Montreal), in May
- Alla Sheffer (Univ. of British Columbia), in May-June
- Karan Singh (Univ. of Toronto), in May-June
- Kartic Subr (Univ. College London), in March
- Peter Vangorp (Univ. Giessen), in September and November
- Romain Vergne (Univ. of Giessen), in March
- Brian Xu (Univ. of British Columbia), in September-October

7.5.1.1. Internships

Emmanuel IARUSSI (from Mar 2012 until Aug 2012), Inria Internship Program

Subject: Computer-assisted drawing lessons

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

Felicitas Hetzelt (from Mar 2012 until Aug 2012)

Subject: Computer-assisted drawing lessons

Institution: University of Erlangen (Germany)

8. Dissemination

8.1. Scientific Animation

8.1.1. Program Committees

G. Drettakis served on the ACM SIGGRAPH program committee. G. Drettakis and A. Bousseau served on the Eurographics Symposium on Rendering program committee. A. Bousseau served on the NPAR program committee.

8.1.2. Community service

G. Drettakis chairs the Eurographics (EG) Awards Committee and the EG Working group on Rendering and is part of the EG conference steering committee and the Advisory Board for Eurographics 2013. He is an associate editor-in-chief for IEEE Transactions on Computer Graphics and Visualization. G. Drettakis heads the steering committee of the Gouraud-Phong Immersive Space at Sophia-Antipolis.

8.1.3. Conference Presentations and Attendance

At SIGGRAPH 2012 in Los Angeles, G. Drettakis presented the paper [15] together with Bruno Galerne, A. Bousseau presented the paper [20] (August), and P-Y. Laffont presented a poster and a short talk on his work [22]. G. Chaurasia also attended the SIGGRAPH conference.

P-Y. Laffont presented his paper [19] at SIGGRAPH Asia 2012 in Singapore (December).

G. Drettakis attended EUROGRAPHICS 2012.

C. Verron presented the paper [23] at the AES Convention in San Francisco.

J. Lopez-Moreno, A. Bousseau and S. Duchêne attended the Eurographics Symposium on Rendering in Paris.

G. Chaurasia attended the *Human Activity and Vision Summer School* organized by Inria Sophia Antipolis from Oct 1-5, 2012.

8.2. Teaching - Supervision - Juries

8.2.1. Teaching

Licence: E. Chapoulie teaches in the L1 module on Analogical electronics (32h), at University Nice Sophia Antipolis EPU. P-Y. Laffont taught “Introduction to Object-Oriented Programming” (64h, L3), in the Computer Science and MUNDUS departments at Polytech’Nice.

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 (10h), G. Drettakis teaches 8h in this program.

8.2.2. Supervision

PhD : Pierre-Yves Laffont, Intrinsic Image Decomposition from Multiple Photographs [11], Université Nice Sophia Antipolis, defended in October 12th 2012, Advisor: George Drettakis, Co-advisor: Adrien Bousseau

PhD in progress : Emmanuelle Chapoulie, titre (provisoire) du mémoire, since October 2010, Advisor: George Drettakis

PhD in progress : Gaurav Chaurasia, titre (provisoire) du mémoire, since October 2010, Advisor: George Drettakis

PhD in progress : Sylvain Duchêne, titre (provisoire) du mémoire, since October 2011, Advisor: George Drettakis

PhD in progress : Emmanuel Iarussi, Computer-Assisted Drawing, since October 2012, Advisor: George Drettakis, Co-advisor: Adrien Bousseau

Internship : Laurent Lefebvre (EPITECH), Interactive Global Illumination, June-August 2012, Advisor: George Drettakis

Internship : Loic Sevecque (UNSA, MAPI), Modeling and Animation, June-August 2012, Advisor: George Drettakis

8.2.2.1. Thesis Committees

A. Bousseau was an examinateur for the PhD committee of Jiazhou Chen (MANAO project team, Inria Bordeaux - Sud Ouest). G. Drettakis was rapporteur for the HDR committee of I. Ihrke (MANAO project team, Inria Bordeaux - Sud Ouest), and on the PhD committee of Anthony Pajot (IRIT-CNRS/Université Paul Sabatier Toulouse).

8.3. Popularization

8.3.1. Invited Talks

A. Bousseau gave an invited talk at Adobe and at the “Cafés de l’Innovation du Centre d’Accueil d’Entreprises Innovantes (CAEI)” in Nice. He also gave a talk at the InSitu project team, Inria Saclay - Ile de France.

Pierre-Yves Laffont was invited to present his PhD work at University of Washington (USA), Paris ACM Siggraph (France) and KAIST (South Korea).

8.3.2. Demonstrations

Participants: Emmanuelle Chapoulie, Rachid Guerchouche, George Drettakis, Adrien Bousseau.

We performed many demonstrations this year, mostly in the Immersive Space but also of our various research results. Specifically, we performed demos to the companies Toyota, Novartis, to representatives of the Pole de Competitive Pegase, Chambre de Commerce et d’Industrie Nice Cote D’Azur, ERCIM and NII Tokyo. We also presented internal demonstrations to C. Kirchner and A. Petit of the Inria national management.

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^o 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^o 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^o 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^o 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

Doctoral Dissertations and Habilitation Theses

[11] P.-Y. LAFFONT. *Décomposition en images intrinsèques à partir de plusieurs photographies*, Université de Nice Sophia-Antipolis and Inria Sophia-Antipolis, October 2012, Accompanying materials described in the thesis can be downloaded from the author's website: http://thesis.py-laffont.info, http://hal.inria.fr/tel-00761119.

Articles in International Peer-Reviewed Journals

[12] A. BOUSSEAU, J. O'SHEA, R. RAMAMOORTHI, M. AGRAWALA. *Gloss Perception in Painterly and Cartoon Rendering*, in "ACM Transactions on Graphics (TOG)", 2013, to appear, http://www-sop.inria.fr/reves/publis/.

[13] G. CIRIO, P. VANGORP, E. CHAPOULIE, M. MARCHAL, A. LÉCUYER, G. DRETTAKIS. *Walking in a Cube: Novel Metaphors for Safely Navigating Large Virtual Environments in Restricted Real Workspaces*, in "IEEE Transactions on Visualization and Computer Graphics", April 2012, vol. 18, n^o 4, p. 546-554 [*DOI :* 10.1109/TVCG.2012.60], http://hal.inria.fr/hal-00759195.

[14] C. De Rousiers, A. Bousseau, K. Subr, N. Holzschuch, R. Ramamoorthi. *Real-Time Rendering of Rough Refraction*, in "IEEE Transactions on Visualization and Computer Graphics", October 2012, vol. 18, n^{o} 10, p. 1591-1602 [*DOI :* 10.1109/TVCG.2011.282], http://hal.inria.fr/hal-00652076.

[15] B. Galerne, A. Lagae, S. Lefebvre, G. Drettakis. *Gabor Noise by Example*, in "ACM Transactions on Graphics (TOG) - SIGGRAPH 2012 Conference Proceedings", July 2012, vol. 31, n^{o} 4, Article No. 73 [*DOI :* 10.1145/2185520.2185569], http://hal.inria.fr/hal-00695670.

[16] E. Garces, A. Munoz, J. Lopez-Moreno, D. Gutiérrez. *Intrinsic Images by Clustering*, in "Computer Graphics Forum", 2012, vol. 31, n^{o} 4, p. 1415-1425 [*DOI :* 10.1111/j.1467-8659.2012.03137.x], http://hal.inria.fr/hal-00761400.

[17] L. Glondu, L. Muguercia, M. Marchal, C. Bosch, H. Rushmeier, G. Dumont, G. Drettakis. *Example-Based Fractured Appearance*, in "Computer Graphics Forum", February 2012, vol. 31, n^{o} 4, p. 1547-1556 [*DOI :* 10.1111/j.1467-8659.2012.03151.x], http://hal.inria.fr/hal-00752381.

[18] 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, to appear, http://hal.inria.fr/hal-00761121.

[19] P.-Y. Laffont, A. Bousseau, S. Paris, F. Durand, G. Drettakis. *Coherent Intrinsic Images from Photo Collections*, in "ACM Transactions on Graphics", 2012, vol. 31, n^{o} 6 [*DOI :* 10.1145/2366145.2366221], http://hal.inria.fr/hal-00761123.

[20] C. Shao, A. Bousseau, A. Sheffer, K. Singh. *CrossShade: Shading Concept Sketches Using Cross-Section Curves*, in "ACM Transactions on Graphics", 2012, vol. 31, n^{o} 4 [*DOI :* 10.1145/2185520.2185541], http://hal.inria.fr/hal-00703202.

[21] M. Taffou, E. Chapoulie, A. David, R. Guerchouche, G. Drettakis, I. Viaud-Delmon. *Auditory-Visual Integration of Emotional Signals in a Virtual Environment for Cynophobia*, in "Annual Review of Cybertherapy and Telemedicine", 2012 [*DOI :* 10.3233/978-1-61499-121-2-238], http://hal.inria.fr/hal-00760726.

International Conferences with Proceedings

[22] P.-Y. Laffont, A. Bousseau, G. Drettakis. *Rich Intrinsic Image Decomposition of Outdoor Scenes from Multiple Views*, in "SIGGRAPH", ACM SIGGRAPH, 2012, Talk & Poster, http://www-sop.inria.fr/reves/Basilic/2012/LBD12.

[23] C. Verron, G. Drettakis. *Procedural audio modeling for particle-based environmental effects*, in "133rd AES Convention", San Francisco, United States, 2012, http://hal.inria.fr/hal-00759818.

[24] C. Zanni, P. Bares, A. Lagae, M. Quiblier, M.-P. Cani. *Geometric Details on Skeleton-based Implicit Surfaces*, in "Eurographics 2012", Cagliari, Italy, May 2012, Short Paper, http://hal.inria.fr/hal-00694504.

Scientific Books (or Scientific Book chapters)

[25] P. Barla, A. Bousseau. *Gradient Art: Creation and Vectorization*, in "Image and Video based Artistic Stylization", J. Colomosse, P. Rosin (editors), Springer, November 2012, http://hal.inria.fr/hal-00760068.

References in notes

[26] 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.

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

[28] 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.

[29] 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, p. 245–257.

[30] 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.

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

[32] 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.

[33] 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.

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

[35] 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.

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

[37] 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.