Activity Report 2014

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, Rendering

Creation of the Project-Team: 2002 July 01, end of the Project-Team: 2014 December 31.

1. Members

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

Engineers
Abdelaziz Djelouah [Inria, from Sept. 2014, granted by FP7 CRPLAY project]
Jérôme Esnault [Inria, from Jan. 2014, granted by FP7 CRPLAY project]
Rachid Guerchouche [Inria, until Sept. 2014, granted by FP7 VERVE project]
Clement Riant [Inria, until Sep 2014, granted by FP7 VERVE project]

PhD Students
Emmanuelle Chapoulie [Inria, until March 2014, Conseil Régional PACA]
Gaurav Chaurasia [Inria, until Mar 2014, granted by FP7 VERVE project]
Sylvain Duchêne [Inria, granted by FR7 VERVE project]
Emmanuel Iarussi [Inria, granted by ANR DRAO project]
Uditha Kasthuriarachchi [Inria, from Nov. 2014, granted by Inria Cordi-S]
Rodrigo Ortiz Cayon [Inria, granted by FP7 CR-PLAY project]

Post-Doctoral Fellows
Stefan Popov [Inria, until June 2014, granted by ANR ALTA project]
Christian Richardt [Inria, until Feb 2014]
Kenneth Vanhoey [Inria, from Oct. 2014]

Visiting Scientist
Fumio Okura [JSPS, from Apr. 2014]

Administrative Assistant
Sophie Honnorat [Inria]

Others
Ayush Tewari [IIIT Hyderabad, intern, June-Aug.2014]
Rahul Arora [IIT Kanpur, intern, May-July 2014]
George Koulieris [Technical University of Crete, student]
Ares Lagae [FWO Belgium, until Dec 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.

3. Research Program

3.1. Plausible 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. 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" [32], 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 [22], but also simpler methods such as accessibility maps [29]. 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.

3.1.2. 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.

Figure 1. Plausible rendering of an outdoors scene containing points, lines and polygons [21], 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.

3.2. High Quality Rendering Using Simulation

3.2.1. 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 [30] which have been recently developed for sound research.

3.2.2. 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 [26], [25], [24] 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 [23]. Lazy evaluation, as well as hierarchical solutions, are clearly interesting avenues of research, although are probably quite application dependent.

3.2.3. 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.

3.2.4. 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 [27] [30]), 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 [28] 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).

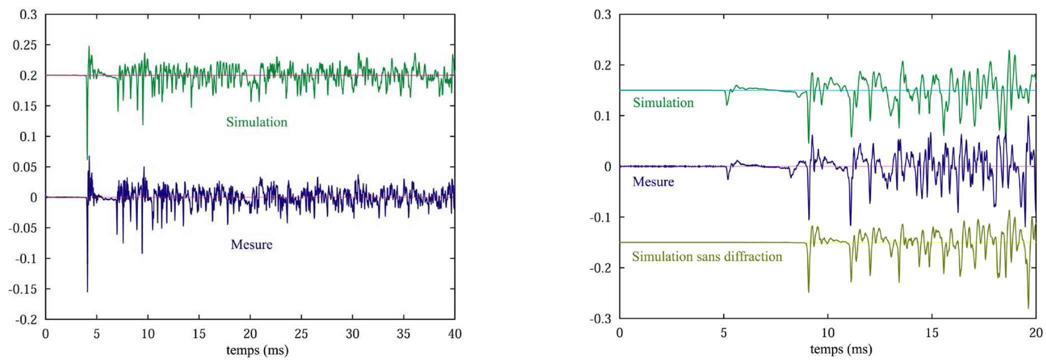

Figure 2. A comparison between a measurement (left) of the sound pressure in a given location of the "Bell Labs Box", a simple test environment built at Bell Laboratories, and a high-quality simulation based on a beam-tracing engine (right). Simulations include effects of reflections off the walls and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

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 [31].

4. Application Domains

4.1. Domain

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.

5. New Software and Platforms

5.1. Multi-View Image-Based Relighting Suite

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

We have continued our development of a set of libraries for handling multi-view image-based relighting algorithms. These constitute the basis for the relighting methods developed for the EU projects VERVE and CR-PLAY.

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 compute 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.

5.2. IBR-Common

Participants: Jerome Esnault, Gaurav Chaurasia, George Drettakis.

This framework provides common tools, utilities and pieces of code to facilitate prototyping of new ideas related to image-based rendering algorithms. Common features include loading shaders, loading images and 3D reconstructions, setting OpenGL context, basic user interface. The factored architecture of the framework allows users to quickly instantiate custom image-based renderers and test them on common datasets. In addition, a CMake structure automates the handling of cross-platform third-party libraries, file systems and compilation. The framework also allowed us to create a version of image-based rendering dedicated to the Immersive Space, in the context of the VERVE EU project.

5.3. IBR in Unity

Participants: Jerome Esnault, Gaurav Chaurasia, George Drettakis.

We have ported our image-based rendering algorithm to the Unity game engine, in collaboration with the Testaluna game company. This technology transfer is in the context of the CR-PLAY EU project.

Our implementation offers important features to game developers:

- Automatic generation of IBR datasets (calibrated cameras and 3D reconstruction) from multiple images of a scene.
- Ability to use different structure-from-motion (Bundler or VisualSFM) and multiview-stereo algorithms (PMVS or MVE from our partner TU Darmstadt).
- Integration of the rendering algorithm in Unity for game prototyping. This port required us to translate the algorithm from C++ to C♯ and to adapt shaders to be compatible with Unity requirements.

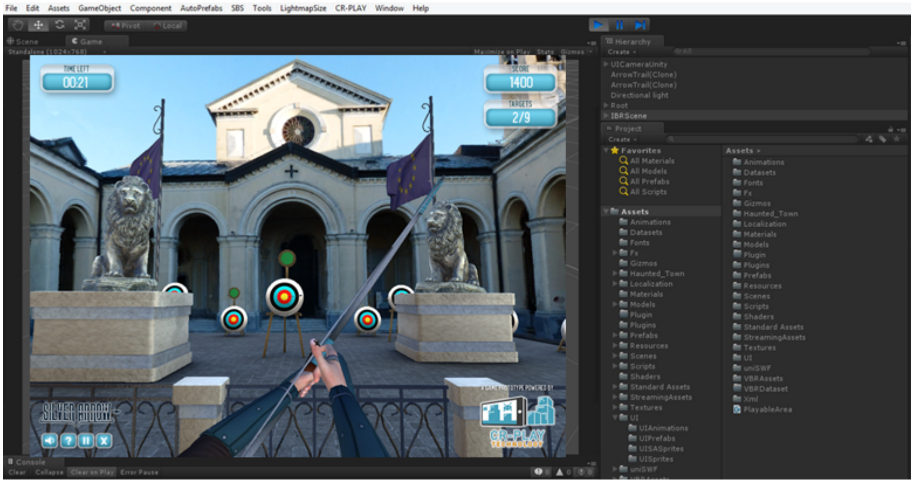

Figure 3. Screen capture of the Unity game development tool. The background buildings are rendered with our image-based rendering algorithm.

Figure 3 shows a screenshot of our Unity package in use for the creation of a simple game.

6. New Results

6.1. Highlights of the Year

Our work on sketch-based modeling for product designers (Sec. 6.4.4) has received significant attention. It appeared on the news page of University of British Columbia http://news.ubc.ca/2014/08/13/powerful-math-creates-3-d-shapes-from-simple-sketches/ and our video has been watched more than 7000 times on Youtube http://youtu.be/tbUljHJv4Rg. We filed a patent on this technology and we have contacts with several companies about a potential transfer.

Our poster on *C-LOD: Context-aware Material Level-of-Detail applied to Mobile Graphics* [16] received the 3rd place in the ACM's Graduate Student Research Competition at SIGGRAPH 2014. This work is a collaboration with George Alex Koulieris and Katerina Mania from the Technical University of Crete and Douglas Cunningham from the Technical University of Cottbus.
BEST PAPER AWARD :
[16] **C-LOD: Context-aware Material Level-of-Detail applied to Mobile Graphics in Computer Graphics Forum**. G. A. KOULIERIS, G. DRETTAKIS, D. W. CUNNINGHAM, K. MANIA.

6.2. Plausible and Realistic Image Rendering

6.2.1. Multi-View Intrinsic Images for Outdoors Scenes with an Application to Relighting

Participants: Sylvain Duchêne, Clement Riant, Gaurav Chaurasia, Stefan Popov, Adrien Bousseau, George Drettakis.

We introduce a method to compute intrinsic images for a multi-view set of outdoor photos with cast shadows, taken under the same lighting. We use an automatic 3D reconstruction from these photos and the sun direction as input and decompose each image into reflectance and shading layers, despite the inaccuracies and missing data of the 3D model. Our approach is based on two key ideas. First, we progressively improve the accuracy of the parameters of our image formation model by performing iterative estimation and combining 3D lighting simulation with 2D image optimization methods. Second we use the image formation model to express reflectance as a function of discrete visibility values for shadow and light, which allows us to introduce a robust visibility classifier for pairs of points in a scene. This classifier is used for shadow labeling, allowing us to compute high quality reflectance and shading layers. We then create shadow-caster geometry that preserves shadow silhouettes. Combined with the intrinsic layers, this approach allows multi-view relighting with moving cast shadows. We present results on several multi-view datasets, and show how it is now possible to perform image-based rendering with changing illumination conditions.

This work is part of an industrial partnership with Autodesk and is under revision for ACM Transactions On Graphics.

6.2.2. Compiler and Tiling Strategies for IIR Filters

Participants: Gaurav Chaurasia, George Drettakis.

We present a compiler for parallelizing IIR or recursive filters. IIR filters are frequently used for $O(1)$ convolutions, but they cannot exploit GPUs because they are very hard to parallelize and also exhibit poor memory locality which hinders performance on both CPUs and GPUs. We present algorithmic tiling strategies for IIR filters which overcome these limitations. Tiled IIR filters are notoriously hard to implement and hence largely ignored by programmers and hardware vendors. We present a compiler front-end that supports intuitive functional specification and tiling of IIR filters. We demonstrate that different tiling strategies may be optimal on different platforms and filter parameters; our compiler can express the exhaustive set of alternatives in just 10-20 lines of code. This enables programmers to easily explore a large variety of trade-offs at different levels of granularity, thereby making it easier and more likely to discover the optimal implementation, while also producing intuitive and maintainable code. Our initial results show that our compiler is as terse as vendor provided libraries, but it allows exploiting the algorithmic advantages of tiling which cannot be provided by any precompiled library.

For example, our compiler can compute a nearly 8 times faster summed area table (4096×4096 image) in 20 lines of code including a fully customized CUDA schedule, as compared to 10 lines in NVIDIA Thrust which does not allow tiling or customizing the CUDA schedule.

This ongoing work is a collaboration with Jonathan Ragan-Kelley (Stanford University), Sylvain Paris (Adobe) and Fredo Durand (MIT).

6.2.3. Video based rendering

Participants: Abdelaziz Djelouah, George Drettakis.

In this project our objective is to propose a new algorithm for novel view synthesis in the case of dynamic scene. The main difference compared to static image-based rendering is the limited number of viewpoints and the presence of the extra time dimension. In a configuration where the number of cameras is limited, segmentation becomes crucial to identify moving foreground regions. To facilitate the difficult task of multi-view segmentation, we currently target scenes captured with stereo cameras. Stereo pairs provide important information on the geometry of the scene while simplifying the segmentation problem.

This ongoing work is a collaboration with Gabriel Brostow from University College London in the context of the CR-PLAY EU project.

6.2.4. Temporally Coherent Video De-Anaglyph

Participants: Joan Sol Roo, Christian Richardt.

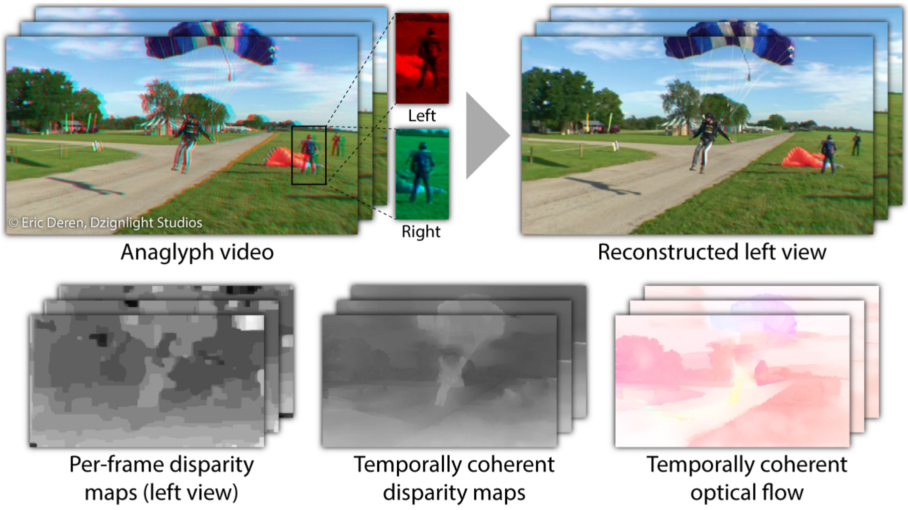

Figure 4. Top: We convert anaglyph videos (left) to temporally coherent full-color stereo videos (right). Bottom: Our approach starts with rough, per-frame disparity maps (left) and produces temporally coherent disparity maps and optical flow (center and right) that are used for reconstructing the stereo views.

This work investigates how to convert existing anaglyph videos to the full-color stereo format used by modern displays. Anaglyph videos only contain half the color information compared to the full-color videos, and the missing color channels need to be reconstructed from the existing ones in a plausible and temporally coherent fashion. In our approach, we put the temporal coherence of the stereo video results front and center (see Figure 4). As a result, our approach is both efficient and temporally coherent. In addition, it computes temporally coherent optical flow and disparity maps that can be used for various post-processing tasks. As a practical contribution, we also make the source code of our implementation available online under CeCILL-B license.

This work was carried out by Joan Sol Roo during his internship in the summer of 2013. The work was presented as a talk and poster at SIGGRAPH 2014 [20].

6.2.5. Probabilistic Connection Path Tracing

Participants: George Drettakis, Stefan Popov.

Bi-directional path tracing (BDPT) with Multiple Importance Sampling (MIS) is one of the most versatile unbiased rendering algorithms today. BDPT repeatedly generates sub-paths from the eye and the lights, which are connected for each pixel and then discarded. Unfortunately, many such bidirectional connections turn out to have low contribution to the solution. The key observation in this project, is that we can find better connections to an eye sub-path by considering multiple light sub-paths at once and creating connections probabilistically only with the most promising ones. We do this by storing light paths, and estimating probability density functions (PDF) of the discrete set of possible connections to all light paths. This has two key advantages: we efficiently create connections with high quality contributions by Monte Carlo sampling, and we reuse light paths across different eye paths. We also introduce a caching scheme for PDFs by deriving a low-dimensional approximation to sub-path contribution.

This ongoing work is a collaboration with Fredo Durand from MIT and Ravi Ramamoorthi from the University of California San Diego in the context of the CRISP associate team.

6.2.6. *Unified Color and Texture Transfer for By-Example Scene Editing*

Participants: Fumio Okura, Kenneth Vanhoey, Adrien Bousseau, George Drettakis.

Color and texture transfer methods are at the heart of by-example image editing techniques. Color transfer well represents the change of overall scene appearance; however it does not represent the change of texture and shape. On the other hand, by-example texture transfer expresses the texture change but it often destroys the target scene structure. We seek the best combination of by-example color and texture transfer to combine these transfer methods so as to selectively work where each method is suitable. Given the source and exemplar pair, the proposed algorithm learns local error metrics which describe if local change between the source and exemplar is best expressed by color or texture transfer. The metric provides us with a local prediction of where we need to synthesize textures using a texture transfer method. This work is a collaboration with Alexei Efros from UC Berkeley in the context of the associate team CRISP.

6.2.7. *Improved Image-Based Rendering*

Participants: Rodrigo Ortiz Cayon, Abdelaziz Djelouah, George Drettakis.

Image-based rendering algorithms based on warping present strong artifacts when rendering surfaces at grazing angle. We are working on a new IBR algorithm that overcomes this problem by rendering superpixel segments as piece-wise homography transformations. The input to our method is a set of images calibrated and a 3D point cloud generated from multi-view stereo reconstruction. In pre-processing we robustly fit planes to superpixel segments that contain reconstruction information and then propagate plausible depth and normal information for image-based rendering. Novel views are obtained by re-projecting superpixel segments as homography from different input views, then adaptively blending them according to distortion and confidence estimations.

6.2.8. *Structured Procedural Textures*

Participants: Kenneth Vanhoey, George Drettakis.

Textures form a popular tool to add visual detail to shapes, objects and scenes. Manual texture design is however a time-consuming process. An alternative is to generate textures from an input exemplar (*i.e.* an acquired photograph) automatically. The difficulty is to synthesize textures of arbitrary size from a single input, preferably with no repetition artifacts. State of the art synthesis techniques can be categorized in two: copy-based techniques and procedural noise-based ones. The first copy pixels using iterative algorithms. The latter deduce a continuous mathematical function from the exemplar, and evaluate it on the space to be textured. They have the advantage of continuity (no resolution-dependence, minimized memory storage, etc.) and fast local evaluation suitable for parallel GPU implementation. They are however tedious to define and manipulate. Current state of the art methods are limited to reproducing Gaussian patterns, that is, textures with no or few structure.

We investigate how to go beyond this limit. Noise-based methods constrain the Fourier power spectrum of a texture-generating noise function to resemble the spectrum of the exemplar. By also constraining the phase of the Fourier spectrum to resemble the exemplar, an exact reproduction is obtained, thus lacking variety and showing maximal repetition. By randomizing the phases, an unstructured "same-looking" image is obtained. This is suitable for noise-like patterns (*e.g.*, marble, wood veins, sand) but not for structured ones (*e.g.*, brick wall, mountain rocks, woven yarn).

In this project, we proceed by investigating the phase spectrum of an image. It contains the structure but identifying how and where is difficult. To characterize structure, we will exploit the splatting process of local random-phase noise and exhibit possible correlations between local phases and spatial placement.

This ongoing work is a collaboration with Ian Jermyn from Durham University.

6.3. Perception for Plausible Rendering

6.3.1. *An Automated High Level Saliency Predictor for Smart Game Balancing*

Participant: George Drettakis.

Successfully predicting visual attention can significantly improve many aspects of computer graphics: scene design, interactivity and rendering. Most previous attention models are mainly based on low-level image features, and fail to take into account high level factors such as scene context, topology, or task. Low-level saliency has previously been combined with task maps, but only for predetermined tasks. Thus, the application of these methods to graphics (e.g., for selective rendering) has not achieved its full potential.

In this work, we present the first automated high-level saliency predictor incorporating two hypotheses from perception and cognitive science that can be adapted to different tasks. The first states that a scene is comprised of objects expected to be found in a specific context as well objects out of context which are salient (scene schemata) while the other claims that viewer's attention is captured by isolated objects (singletons). We proposed a new model of attention by extending Eckstein's Differential Weighting Model. We conducted a formal eye-tracking experiment which confirmed that object saliency guides attention to specific objects in a game scene and determined appropriate parameters for a model. We presented a GPU-based system architecture that estimates the probabilities of objects to be attended in real-time (Figure 5). We embedded this tool in a game level editor to automatically adjust game level difficulty based on object saliency, offering a novel way to facilitate game design. We perform a study confirming that game level completion time depends on object topology as predicted by our system.

Figure 5. A low level saliency algorithm indicates that the most salient area of the image is the dark area behind the chair. Our tool highlights the vase at a consistent/singleton location as the most salient object in the image.

This work is a collaboration with George Alex Koulieris and Katerina Mania from the Technical University of Crete and Douglas Cunningham from the Technical University of Cottbus. The work was published in the ACM Transactions on Applied Perception (TAP) Journal [15] and presented as a Talk at SIGGRAPH 2014 in Vancouver.

6.3.2. C-LOD: Context-aware Material Level-of-Detail applied to Mobile Graphics

Participant: George Drettakis.

Attention-based Level-Of-Detail (LOD) managers downgrade the quality of areas that are expected to go unnoticed by an observer to economize on computational resources. The perceptibility of lowered visual fidelity is determined by the accuracy of the attention model that assigns quality levels. Most previous attention based LOD managers do not take into account saliency provoked by context, failing to provide consistently accurate attention predictions.

In this work, we extended a recent high level saliency model with four additional components yielding more accurate predictions: an object-intrinsic factor accounting for canonical form of objects, an object-context factor for contextual isolation of objects, a feature uniqueness term that accounts for the number of salient features in an image, and a temporal context that generates recurring fixations for objects inconsistent with the context. We conducted a perceptual experiment to acquire the weighting factors to initialize our model. We then designed C-LOD, a LOD manager that maintains a constant frame rate on mobile devices by dynamically re-adjusting material quality on secondary visual features of non-attended objects. In a proof of concept study we established that by incorporating C-LOD, complex effects such as parallax occlusion mapping usually omitted in mobile devices can now be employed, without overloading GPU capability and, at the same time, conserving battery power. We validated our work via eye-tracking (Figure 6)

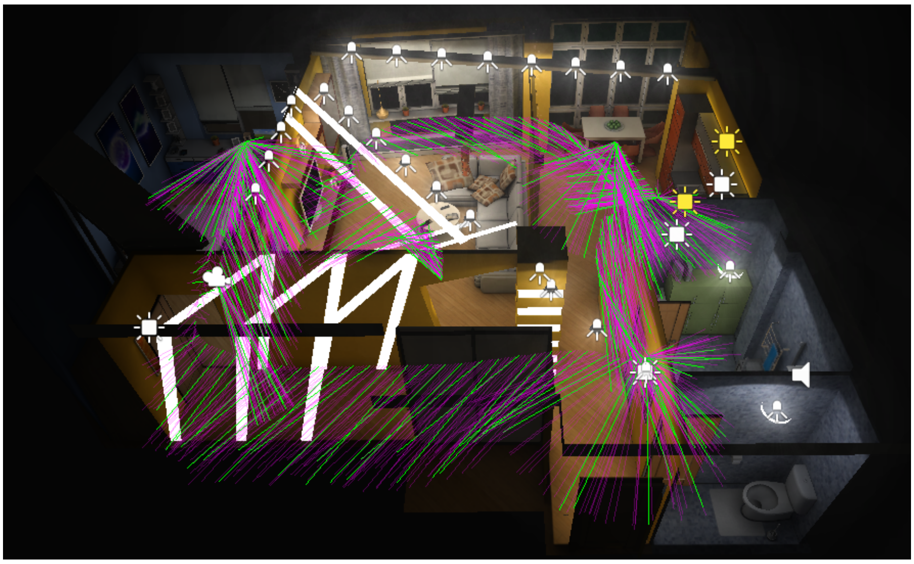

Figure 6. Our validation tool indicates the subject's gaze point with magenta colored beams. The green beams indicate predictions by our attention model.

This work is a collaboration with George Alex Koulieris and Katerina Mania from the Technical University of Crete and Douglas Cunningham from the Technical University of Cottbus. The work was published in a special issue of Computer Graphics Forum [16] and was presented at the Eurographics Symposium on Rendering 2014 in Lyon. It was also presented as a poster at SIGGRAPH 2014 in Vancouver winning the 3rd place in the ACM's Graduate Student Research Competition.

6.4. Interaction and Design for Virtual Environments

6.4.1. Evaluation of Direct Manipulation using Finger Tracking for Complex Tasks in an Immersive Cube

Participants: Emmanuelle Chapoulie, George Drettakis.

We present a solution for interaction using finger tracking in a cubic immersive virtual reality system (or immersive cube). Rather than using a traditional flystick device, users can manipulate objects with fingers of both hands in a close-to-natural manner for moderately complex, general purpose tasks. Our solution couples finger tracking with a real-time physics engine, combined with a heuristic approach for hand manipulation, which is robust to tracker noise and simulation instabilities. We performed a first study to evaluate our interface with tasks involving complex manipulations, such as balancing objects while walking in the cube. The users finger-tracked manipulation was compared to manipulation with a 6 degree-of-freedom flystick, as well as with carrying out the same task in the real world. Users were also asked to perform a free task, allowing us to observe their perceived level of presence in the scene. Our results showed that our approach provides a feasible interface for immersive cube environments and is perceived by users as being closer to the real experience compared to the flystick. However, the flystick outperforms direct manipulation in terms of speed and precision.

This work is a collaboration with Maria Roussou and Evanthia Dimara from the University of Athens, Maud Marchal from Inria Rennes, and Jean-Christophe Lombardo from Inria Sophia Antipolis. The work has been published in the journal Virtual Reality [13].

Figure 7. A user balancing a tray with both hands [13].

We have also worked on a followup study in which we examine a much more controlled context, studying only very limited movements, in 1D, 2D and 3D. To do this we designed specific devices that can be instantiated both in the virtual world and as physical objects. We compared finger manipulation to wand and to real configurations; the study demonstrated the feasibility of such a controlled comparison for the study of finger-based interaction. This work is in collaboration with InSitu, specifically F. Tsandilas, W. Mackay and L. Oehlberg, and has been accepted for publication in 2015 at IEEE 3DUI.

6.4.2. Reminiscence Therapy using Image-Based Rendering in VR

Participants: Emmanuelle Chapoulie, George Drettakis, Rachid Guerchouche, Gaurav Chaurasia.

We present a novel VR solution for Reminiscence Therapy (RT), developed jointly by a group of memory clinicians and computer scientists. RT involves the discussion of past activities, events or experiences with others, often with the aid of tangible props which are familiar items from the past; it is a popular intervention in dementia care. We introduced an immersive VR system designed for RT, which allows easy presentation of familiar environments. In particular, our system supports highly-realistic Image-Based Rendering in an immersive setting. To evaluate the effectiveness and utility of our system for RT, we performed a study with healthy elderly participants to test if our VR system could help with the generation of autobiographical memories. We adapted a verbal Autobiographical Fluency protocol to our VR context, in which elderly participants were asked to generate memories based on images they were shown. We compared the use of our image-based system for an unknown and a familiar environment. The results of our study showed that the number of memories generated for a familiar environment is higher than the number of memories obtained for an unknown environment using our system. This indicates that IBR can convey familiarity of a given scene, which is an essential requirement for the use of VR in RT. Our results also showed that our system is as effective as traditional RT protocols, while acceptability and motivation scores demonstrated that our system is well tolerated by elderly participants.

This work is a collaboration with Pierre-David Petit and Philippe Robert from the CMRR in Nice. The work has been published in the Proceedings of IEEE Virtual Reality [19].

Figure 8. Left: our hardware setup. Right: new point of view reconstructed from input cameras.

6.4.3. 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. We conducted a thorough study to evaluate different lightfield editing interfaces, tools and workflows from a user perspective. We additionally investigate the potential benefits of using depth information when editing, and the limitations imposed by imperfect depth reconstruction using current techniques. We perform two different experiments, collecting both objective and subjective data from a varied number of point-based editing tasks of increasing complexity: In the first experiment, we rely on perfect depth from synthetic lightfields, and focus on simple edits. This allows us to gain basic insight on lightfield editing, and to design a more advanced editing interface. This is then used in the second experiment, employing real lightfields with imperfect reconstructed depth, and covering more advanced editing tasks. Our study shows that users can edit lightfields with our tested interface and tools, even in the presence of imperfect depth. They follow different workflows depending on the task at hand, mostly relying on a combination of different depth cues. Last, we confirm our findings by asking a set of artists to freely edit both real and synthetic lightfields.

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. This work was published at ACM Transactions on Graphics 2014 (Proc. SIGGRAPH) [14].

6.4.4. True2Form: 3D Curve Networks from 2D Sketches via Selective Regularization

Participant: Adrien Bousseau.

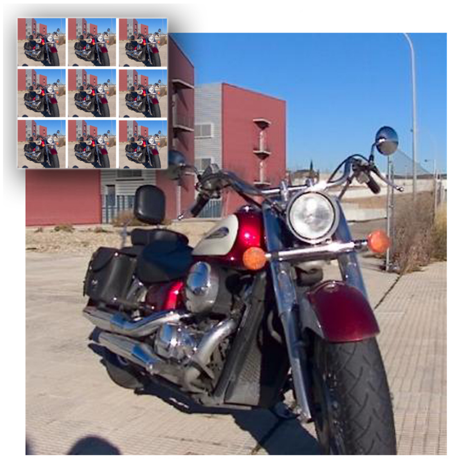

Figure 9. A lightfields represents multiple nearby views of a scene. We conducted a study to evaluate how people edit such data.

True2Form is a sketch-based modeling system that reconstructs 3D curves from typical design sketches. Our approach to infer 3D form from 2D drawings is a novel mathematical framework of insights derived from perception and design literature. We note that designers favor viewpoints that maximally reveal 3D shape information, and strategically sketch descriptive curves that convey intrinsic shape properties, such as curvature, symmetry, or parallelism. Studies indicate that viewers apply these properties selectively to envision a globally consistent 3D shape. We mimic this selective regularization algorithmically, by progressively detecting and enforcing applicable properties, accounting for their global impact on an evolving 3D curve network. Balancing regularity enforcement against sketch fidelity at each step allows us to correct for inaccuracy inherent in free-hand sketching. We perceptually validate our approach by showing agreement between our algorithm and viewers in selecting applicable regularities. We further evaluate our solution by: reconstructing a range of 3D models from diversely sourced sketches; comparisons to prior art; and visual comparison to both ground-truth and 3D reconstructions by designers.

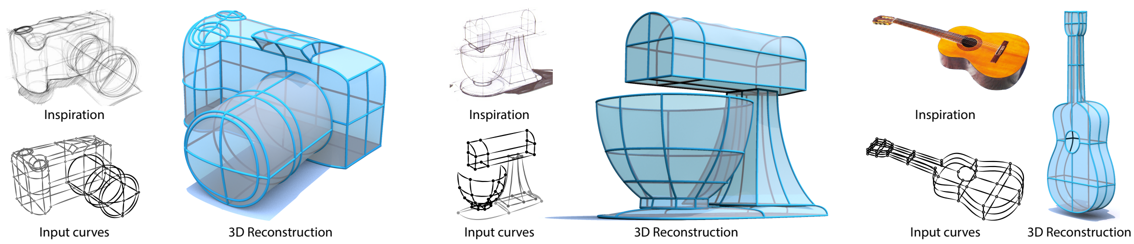

Figure 10. Our single-view modeling system allows us to reconstruct 3D models by tracing curves over existing sketches and photographs.

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. The paper was published at ACM Transactions on Graphics 2014 (Proc. SIGGRAPH) [18].

6.4.5. BendFields: Regularized Curvature Fields from Rough Concept Sketches

Participants: Adrien Bousseau, Emmanuel Iarussi.

Designers frequently draw curvature lines to convey bending of smooth surfaces in concept sketches. We present a method to extrapolate curvature lines in a rough concept sketch, recovering the intended 3D curvature field and surface normal at each pixel of the sketch. This 3D information allows us to enrich the sketch with 3D-looking shading and texturing. We first introduce the concept of *regularized curvature lines* that model the lines designers draw over curved surfaces, encompassing curvature lines and their extension as geodesics over flat or umbilical regions. We build on this concept to define the orthogonal cross field that assigns two regularized curvature lines to each point of a 3D surface. Our algorithm first estimates the projection of this cross field in the drawing, which is non-orthogonal due to foreshortening. We formulate this estimation as a scattered interpolation of the strokes drawn in the sketch, which makes our method robust to sketchy lines that are typical for design sketches. Our interpolation relies on a novel smoothness energy that we derive from our definition of regularized curvature lines. Optimizing this energy subject to the stroke constraints produces a dense non-orthogonal 2D cross field, which we then lift to 3D by imposing orthogonality. Thus, one central concept of our approach is the generalization of existing cross field algorithms to the non-orthogonal case. We demonstrate our algorithm on a variety of concept sketches with various levels of sketchiness. We also compare our approach with existing work that takes clean vector drawings as input.

This work is a collaboration with David Bommes from Titane project team, Inria Sophia-Antipolis. The manuscript has been accepted for publication with minor revisions at ACM Transactions on Graphics (TOG).

6.4.6. Line Drawing Interpretation in a Multi-View Context

Participant: Adrien Bousseau.

Many design tasks involve the creation of new objects in the context of an existing scene. Existing work in computer vision only provides partial support for such tasks. On the one hand, multi-view stereo algorithms allow the reconstruction of real-world scenes, while on the other hand algorithms for line-drawing interpretation do not take context into account. This work combines the strength of these two domains to interpret line drawings of imaginary objects drawn over photographs of an existing scene. The main challenge we face is to identify the existing 3D structure that correlates with the line drawing while also allowing the creation of new structure that is not present in the real world. We propose a labeling algorithm to tackle this problem, where some of the labels capture dominant orientations of the real scene while a free label allows the discovery of new orientations in the imaginary scene.

This work is a collaboration with Jean-Dominique Favreau and Florent Lafarge from Titane project team, Inria Sophia-Antipolis and is under submission for the CVPR conference.

6.4.7. Wrap It! Computer-Assisted Design and Fabrication of Wire Wrapped Jewelry

Participants: Adrien Bousseau, Emmanuel Iarussi.

We developed an interactive tool to assist the process of creating and crafting wire wrapped pieces of jewelry. In a first step, we guide the user in conceiving designs which are suitable to be fabricated with metal wire. In a second step, we assist fabrication by taking inspiration from jigs-based techniques, frequently used by craftsmen as a way to guide and support the wrapping process. Given a vector drawing composed of curves to be fabricated, it is crucial to first decompose it into segments that can be constructed with metal wire. Literature on jewelry-making provides a wide range of examples to perform this task, but they are hard to generalize to any input design. Based on the observation of these examples, we distill and generalize a set of design principles behind the finished pieces of jewelry. Relying on those principles, we propose an algorithm that generates a decomposition of the input where each piece is a single component of wire, that can be wrapped and gathered with the others. In addition, we also automate the design of custom physical jigs for fabrication of the jewelry piece. A jig consists of a board with holes on it, arranged in a regular grid structure. By placing a set of pins (of different radius) on the jig, the craftman builds a support structure that guides the wrapping process. The wire is bended and twisted around those pins to create the shape. Given the input design curves and the available jig parameters (size, number and radius of the pins), we propose an algorithm to automatically generate an arrangement of pins in order to better approximate the input curve with wire. Finally, users can follow automatically-generated step-by-step instructions to place the pins in the jig board and fabricate the end piece of jewelry.

This ongoing work is a collaboration with Wilmot Li from Adobe, San Francisco. The project was initiated by a 3-months visit of Emmanuel Iarussi at Adobe.

6.4.8. *Studying how novice designers communicate with sketches and prototypes*

Participant: Adrien Bousseau.

We performed a user study to better understand how novice designers communicate a concept during the different phases of its development. Our study was conducted as a one-day design contest where participants had to propose a concept, present it to a jury, describe it to an engineer and finally fabricate a prototype with the help of another participant. We collected sketches and videos for all steps of this exercise in order to evaluate how the concept evolves and how it is described to different audiences. We hope that our findings will inform the development of better computer-assisted design tools for novices.

This is an ongoing work in collaboration with Wendy McKay, Theophanis Tsandilas and Lora Oehlberg from the InSitu project team - Inria Saclay, in the context of the ANR DRAO project.

6.4.9. *Vectorising Bitmaps into Semi-Transparent Gradient Layers*

Participants: Christian Richardt, Adrien Bousseau, George Drettakis.

Figure 11. Our interactive vectorisation technique lets users vectorise an input bitmap (a) into a stack of opaque and semi-transparent vector layers composed of linear or radial colour gradients (b). Users can manipulate the resulting layers using standard tools to quickly produce new looks (c). Semi-transparent layers are outlined for visualisation; these edges are not part of our result.

Vector artists create complex artworks by stacking simple layers. We demonstrate the benefit of this strategy for image vectorisation, and present an interactive approach for decompositing bitmap drawings and studio photographs into opaque and semi-transparent vector layers. Semi-transparent layers are especially challenging to extract, since they require the inversion of the non-linear compositing equation. We make this problem tractable by exploiting the parametric nature of vector gradients, jointly separating and vectorising semi-transparent regions. Specifically, we constrain the foreground colours to vary according to linear or radial parametric gradients, restricting the number of unknowns and allowing our system to efficiently solve for an editable semi-transparent foreground.

We propose a progressive workflow, where the user successively selects a semi-transparent or opaque region in the bitmap, which our algorithm separates as a foreground vector gradient and a background bitmap layer. The user can choose to decompose the background further or vectorise it as an opaque layer. The resulting layered vector representation allows a variety of edits, as illustrated in Figure 11, such as modifying the shape of highlights, adding texture to an object or changing its diffuse colour. Our approach facilitates the creation of such layered vector graphics from bitmaps, and we thus see our method as a valuable tool for professional artists and novice users alike.

This work is a collaboration with Jorge Lopez-Moreno, now a postdoc at the University of Madrid, and Maneesh Agrawala from the University of California, Berkeley in the context of the CRISP Associated Team. The paper was presented at the Eurographics Symposium on Rendering (EGSR) 2014, and is published in a special issue of the journal Computer Graphics Forum [17].

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

7.1.1. Autodesk

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

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

7.2. Bilateral Grants with Industry

7.2.1. Adobe

Participants: George Drettakis, Gaurav Chaurasia.

Adobe has offered a small donation in the context of our collaboration on compilers for image processing (Sec. 6.2.2).

8. Partnerships and Cooperations

8.1. National Initiatives

8.1.1. 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. 6.2.5.

8.1.2. 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 3D modeling from sketches [18] and on vectorization of photographs [17].

8.1.3. 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.

8.2. European Initiatives

8.2.1. 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.

8.2.2. 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.

We have started specific scientific collaborations with TUD on capture guidance and UCL on video-based rendering, which will continue in 2015.

8.3. International Initiatives

8.3.1. Inria Associate Teams

8.3.1.1. CRISP2

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

International Partner (Institution - Laboratory - Researcher):

University of California Berkeley

Duration: 2011 - Present

See also: http://www-sop.inria.fr/reves/crisp/

The CRISP collaboration aims at developing novel techniques to create and manipulate effective numerical imagery. We adopt a multidisciplinary approach, focusing on understanding how people create and perceive images, on developing new rendering algorithms based on this understanding, and on building interactive tools that enable users to efficiently produce the images they have in mind. The participants of CRISP share complementary expertise in computer graphics, human computer interaction and human visual perception.

After a very productive year in 2013, we continued our work on drawing and manipulating materials in vector graphics in 2014. This work was published in the Computer Graphics Forum journal and presented at the Eurographics Symposium on Rendering (EGSR) [17]. We are currently working on two collaborative projects in the context of CRISP. One project is on light transport simulation (with Ravi Ramamoorthi, now at UC San Diego), the other project is on appearance transfer between photographs (with Alyosha Efros, who recently joined UC Berkeley). We also have several project ideas to start with Martin S. Banks (Human Vision Science).

8.3.2. Informal International Partners

8.3.2.1. France-USA

Participants: Gaurav Chaurasia, Emmanuel Iarussi, Adrien Bousseau, George Drettakis.

Beyond the CRISP associate team, we have an ongoing collaboration with Adobe Research (Sylvain Paris) and MIT (Fredo Durand) on parallel image-proccessing languages and global illumination (Fredo Durand). We also have another collaboration with Adobe Research (Wilmot Li) on jewelry design. Emmanuel Iarussi did a 3-months visit at Adobe in the context of this collaboration.

8.3.2.2. 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. 6.4.4).

8.3.2.3. France-Greece

Participant: George Drettakis.

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) (see Sec. 6.3.2 and 6.3.1).

8.4. International Research Visitors

8.4.1. Visits of International Scientists

8.4.1.1. Visitors

We hosted several researchers this year:

- Gordon Wetzstein (MIT), in January
- Wendy McKay and Theophanis Tsandilas in February
- Kenneth Vanhoey (Univ. de Strasbourg), in February
- Fredo Durand (MIT), in February
- Jean Ponce (ENS), in February
- Marcus Magnor (TU Braunschweig), in February
- Christian Theobalt (Max Planck Institut), in February
- Markus Gross (ETH Zurich), in April
- Abdelaziz Djelouah (Inria Grenoble), in May
- Indira Thouvenin (UT Compiegne), in June
- Josef Sivic (Inria and ENS), in July
- Wenzel Jakob (ETH Zurick), in September
- Marty Banks (Berkeley part of EA CRISP), in June and November
- Gaurav Chaurasia (MIT), in November

8.4.1.2. Internships

Arora Rahul

Date: May 2014 - July 2014

Institution: IITK (India)

Ayush Tewari

Date: June 2014 - Aug 2014

Institution: IIIT

Uditha Kasthuriarachchi

Date: April 2014 - Sept 2014

Institution: UNSA

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Scientific events organisation

9.1.1.1. general chair, scientific chair

G. Drettakis chairs the EG Working group on Rendering. He also heads the steering committee of the Gouraud-Phong Immersive Space at Sophia-Antipolis.

9.1.1.2. member of the organizing committee

G. Drettakis served on several ad-hoc ACM and IEEE Committees, concerning policy and EIC searches.

9.1.2. Scientific events selection

9.1.2.1. member of the conference program committee

George Drettakis and Adrien Bousseau served on the Eurographics 2014 international program committee. Adrien Bousseau served on the SIGGRAPH 2014 program committee and the SIBGRAPI 2014 program committee. George Drettakis served on the EGSR 2014 international program committee.

9.1.3. Journal

9.1.3.1. member of the editorial board

G. Drettakis is an associate editor-in-chief for IEEE Transactions on Computer Graphics and Visualization.

9.1.3.2. reviewer

G. Drettakis and A. Bousseau are frequent reviewers for ACM Transactions on Graphics, IEEE Transactions on Computer Graphics and Visualization, Computer Graphics Forum.

9.1.4. Conferences

9.1.4.1. Attendence

- Sylvain Duchêne, Rodrigo Ortiz Cayon, Adrien Bousseau and George Drettakis attended the Eurographics Symposium on Rendering in Lyon.
- Emmanuel Iarussi, Rodrigo Ortiz Cayon and Adrien Bousseau attended the Eurographics conference in Strasbourg.
- Adrien Bousseau and George Drettakis attended the SIGGRAPH conference in Vancouver.
- Kenneth Vanhoey attended the ACM Siggraph Asia 2014 conference in Shenzhen, China

9.1.5. Summer Schools

9.1.5.1. Attendence

Rodrigo Ortiz Cayon attended a summer school on virtual reality at ESTIA - Bidart, France.

9.2. Teaching - Supervision - Juries

9.2.1. 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 4.5h in this program.

Licence: Emmanuel Iarussi is a teaching assistant for the course on Human Computer Interaction at IUT - University Nice Sophia-Antipolis.

9.2.2. Supervision

- PhD in progress: Rodirgo Ortiz-Cayon, Inpainting and Artifact Removal for Image Based Rendering, (started November 2013), Advisor: George Drettakis.
- PhD in progress: Emmanuel Iarussi, Computer-Assisted Drawing, since October 2012, Advisor: George Drettakis, Co-advisor: Adrien Bousseau
- PhD in progress: Sylvain Duchene, Image-Based Relighting, since October 2011 (to defend in March 2015), Advisor: George Drettakis.
- PhD in progress: Uditha Kasthuriarachchi, Mixed rendering of captured and virtual content, since November 2014, Advisor: George Drettakis.
- PhD defended: Gaurav Chaurasia, Algorithms and Perceptual Analysis for Interactive Free Viewpoint Image-Based Navigation [12], since October 2010 (defended February 2014), Advisor: George Drettakis
- PhD defended: Emmanuelle Chapoulie, Gestures and direct manipulation for immersive virtual reality [11], since October 2010 (defended June 2014), Advisor: George Drettakis
- Internship: Ayush Tewari. Advisor: George Drettakis, Co-advisor: Adrien Bousseau
- Internship: Rahul Arora. Advisor: Adrien Bousseau
- Internship: Uditha Kasthuriarachchi. Advisor: George Drettakis

9.2.3. Juries

- G. Drettakis was a member of the Ph.D. committee of J. Ardouin in Rennes (December).

9.3. Popularization

9.3.1. Popularization, Invited Talks

- We participated in the VERVE workshop on virtual reality technologies for ageing, which took place at the Institut Claude Pompidou on Nov. 6th. Over 200 participants (patients, families and caregivers) visited during this day. George Drettakis gave a short interview on our research for local television.
- Adrien Bousseau presented his work on sketch-based modeling at Lycée de Vence during the Fête de la science 2014. He also participated to a "cafe science" during the Semaine du Cerveau event in Marseille.
- Adrien Bousseau gave an invited talk at the Amadeus Global Technical Forum, December 2014.
- Adrien Bousseau and George Drettalis gave invited talks at UC Berkeley in the context of the associate team CRISP.

9.3.2. Demos

We gave demos to companies visiting the site, often including presentations in the immersive space: Visteon, Optis. We also gave demos and presentations at the Mediterranean Days event, to university and high school students ("stages 3eme"), as well as representatives from the Sustainable Design School in Nice.

10. 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] G. CHAURASIA, S. DUCHÊNE, O. SORKINE-HORNUNG, G. DRETTAKIS. *Depth Synthesis and Local Warps for Plausible Image-based Navigation*, in "ACM Transactions on Graphics", 2013, vol. 32, to be presented at SIGGRAPH 2013, http://www-sop.inria.fr/reves/Basilic/2013/CDSD13

[3] 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

[4] 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)", ACM, 2013, http://www-sop.inria.fr/reves/Basilic/2013/IBT13

[5] P.-Y. LAFFONT, A. BOUSSEAU, S. PARIS, F. DURAND, G. DRETTAKIS. *Coherent Intrinsic Images from Photo Collections*, in "ACM Transactions on Graphics (SIGGRAPH Asia Conference Proceedings)", 2012, vol. 31, http://www-sop.inria.fr/reves/Basilic/2012/LBPDD12

[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] C. SHAO, A. BOUSSEAU, A. SHEFFER, K. SINGH. *CrossShade: Shading Concept Sketches Using Cross-Section Curves*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", 2012, vol. 31, n^o 4, http://www-sop.inria.fr/reves/Basilic/2012/SBSS12

[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] E. CHAPOULIE. *Gestures and direct manipulation for immersive virtual reality*, Université Nice Sophia Antipolis, June 2014, https://tel.archives-ouvertes.fr/tel-01070736

[12] G. CHAURASIA. *Algorithms and Perceptual Analysis for Interactive Free Viewpoint Image-Based Navigation*, Université Nice Sophia Antipolis, February 2014, https://tel.archives-ouvertes.fr/tel-00979913

Articles in International Peer-Reviewed Journals

[13] E. CHAPOULIE, M. MARCHAL, E. DIMARA, M. ROUSSOU, J.-C. LOMARDO, G. DRETTAKIS. *Evaluation of Direct Manipulation using Finger Tracking for Complex Tasks in an Immersive Cube*, in "Virtual Reality", August 2014, https://hal.inria.fr/hal-01060303

[14] A. JARABO, B. MASIA, A. BOUSSEAU, F. PELLACINI, D. GUTIERREZ. *How Do People Edit Light Fields?*, in "ACM Transactions on Graphics", 2014, vol. 33, n^o^ 4 [*DOI :* 10.1145/2601097.2601125], https://hal.inria.fr/hal-01060860

[15] G. A. KOULIERIS, G. DRETTAKIS, D. W. CUNNINGHAM, K. MANIA. *An Automated High Level Saliency Predictor for Smart Game Balancing*, in "ACM Transactions on Applied Perception", October 2014, https://hal.inria.fr/hal-01062718

[16] *Best Paper*
G. A. KOULIERIS, G. DRETTAKIS, D. W. CUNNINGHAM, K. MANIA. *C-LOD: Context-aware Material Level-of-Detail applied to Mobile Graphics*, in "Computer Graphics Forum", July 2014, https://hal.inria.fr/hal-01062710.

[17] C. RICHARDT, J. LOPEZ-MORENO, A. BOUSSEAU, M. AGRAWALA, G. DRETTAKIS. *Vectorising Bitmaps into Semi-Transparent Gradient Layers*, in "Computer Graphics Forum", 2014, vol. 33, n^o^ 4, https://hal.inria.fr/hal-01060820

[18] B. XU, W. CHANG, A. SHEFFER, A. BOUSSEAU, J. MCCRAE, K. SINGH. *True2Form: 3D Curve Networks from 2D Sketches via Selective Regularization*, in "ACM Transactions on Graphics", 2014, vol. 33, n^o^ 4 [*DOI :* 10.1145/2601097.2601128], https://hal.inria.fr/hal-01060850

International Conferences with Proceedings

[19] E. CHAPOULIE, R. GUERCHOUCHE, P.-D. PEIT, G. CHAURASIA, P. ROBERT, G. DRETTAKIS. *Reminiscence Therapy using Image-Based Rendering in VR*, in "IEEE Virtual Reality", Minneapolis, United States, September 2014, https://hal.inria.fr/hal-01060300

Scientific Popularization

[20] J. SOL ROO, C. RICHARDT. *Temporally Coherent Video De-Anaglyph*, September 2014 [*DOI :* 10.1145/2614106.2614125], https://hal.inria.fr/hal-01060864

References in notes

[21] 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

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

[23] 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

[24] 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

[25] 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

[26] 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

[27] 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

[28] 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

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

[30] 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

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

[32] 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