Activity Report 2017

Project-Team GRAPHDECO

GRAPHics and DEsign with hEterogeneous COntent

RESEARCH CENTER
Sophia Antipolis - Méditerranée

THEME
Interaction and visualization

Table of contents

Project-Team GRAPHDECO

Creation of the Team: 2015 January 01, updated into Project-Team: 2015 July 01

Keywords:

Computer Science and Digital Science:

A3.1.4. - Uncertain data
A5. - Interaction, multimedia and robotics
A5.1. - Human-Computer Interaction
A5.1.1. - Engineering of interactive systems
A5.1.2. - Evaluation of interactive systems
A5.3.5. - Computational photography
A5.4.4. - 3D and spatio-temporal reconstruction
A5.5. - Computer graphics
A5.5.1. - Geometrical modeling
A5.5.2. - Rendering
A5.5.3. - Computational photography
A5.6. - Virtual reality, augmented reality
A5.9.1. - Sampling, acquisition
A6.3.5. - Uncertainty Quantification
A8.3. - Geometry, Topology
A9.2. - Machine learning
A9.3. - Signal analysis

Other Research Topics and Application Domains:

B5. - Industry of the future
B5.2. - Design and manufacturing
B5.7. - 3D printing
B8. - Smart Cities and Territories
B8.3. - Urbanism and urban planning
B9. - Society and Knowledge
B9.1.2. - Serious games
B9.2. - Art
B9.2.2. - Cinema, Television
B9.2.3. - Video games
B9.5. - Humanities
B9.5.6. - Archeology, History

1. Personnel

Research Scientists

George Drettakis [Team leader, Inria, Senior Researcher, HDR]
Adrien Bousseau [Inria, Researcher, HDR]
Frederic Durand [Inria, International Chair, Advanced Research Position]

Post-Doctoral Fellows

Gabriel Cirio [Inria]
Georgios Koulieris [Inria]
Abdelaziz Djelouah [Inria, until Feb 2017]
Miika Aittala [Université étrangère, until Jul 2017]

PhD Students

Johanna Delanoy [Inria]
Valentin Deschaintre [Optis CIFRE]
Jean-Dominique Favreau [Inria, from Oct 2017]
Julien Philip [Inria]
Simon Rodriguez [Université Nice Sophia-Antipolis]
Theo Thonat [Inria]

Technical staff

Laurent Boiron [Inria]
Sebastien Bonopera [Inria, from Sep 2017 until Nov 2017]
Yulia Gryaditskaya [Inria, from Feb 2017]
Sebastian Vizcay [Inria, from May 2017]
Bastien Wailly [Inria, from Nov 2017]

Interns

Sai Praveen Bangaru [Inria, from May 2017 until Jul 2017]
Mathieu Chambe [Ecole normale supérieure de Rennes, from May 2017 until Jul 2017]
Marina Menghetti [Inria, from Jul 2017 until Aug 2017]
Arijit Panigrahy [Inria, from May 2017 until Jul 2017]
Adele Saint-Denis [Inria, from Apr 2017 until Jul 2017]

Administrative Assistant

Sophie Honnorat [Inria]

2. Overall Objectives

2.1. General Presentation

In traditional Computer Graphics (CG) input is *accurately modeled* by hand by artists. The artists first create the 3D geometry – i.e., the polygons and surfaces used to represent the 3D scene. They then need to assign colors, textures and more generally material properties to each piece of geometry in the scene. Finally they also define the position, type and intensity of the lights. This modeling process is illustrated schematically in Fig. 1(left)). Creating all this 3D content involves a high level of training and skills, and is reserved to a small minority of expert modelers. This tedious process is a significant distraction for creative exploration, during which artists and designers are primarily interested in obtaining compelling imagery and prototypes rather than in accurately specifying all the ingredients listed above. Designers also often want to explore many variations of a concept, which requires them to repeat the above steps multiple times.

Once the 3D elements are in place, a *rendering* algorithm is employed to generate a shaded, realistic image (Fig. 1(right)). Costly rendering algorithms are then required to simulate light transport (or *global illumination*) from the light sources to the camera, accounting for the complex interactions between light and materials and the visibility between objects. Such rendering algorithms only provide meaningful results if the input has been *accurately* modeled and is *complete*, which is prohibitive as discussed above.

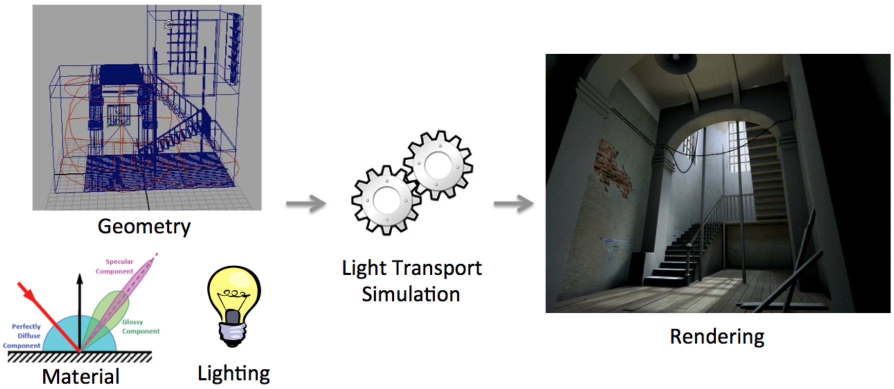

Figure 1. Traditional computer graphics pipeline. Rendering from www.thegnomonworkshop.com

A major recent development is that many alternative sources of 3D content are becoming available. Cheap depth sensors allow anyone to capture real objects but the resulting 3D models are often *uncertain*, since the reconstruction can be inaccurate and is most often incomplete. There have also been significant advances in casual content creation, e.g., sketch-based modeling tools. The resulting models are often approximate since people rarely draw accurate perspective and proportions. These models also often lack details, which can be seen as a form of uncertainty since a variety of refined models could correspond to the rough one. Finally, in recent years we have witnessed the emergence of new usage of 3D content for rapid prototyping, which aims at accelerating the transition from rough ideas to physical artifacts.

The inability to handle *uncertainty* in the data is a major shortcoming of CG today as it prevents the direct use of cheap and casual sources of 3D content for the design and rendering of high-quality images. The abundance and ease of access to *inaccurate*, *incomplete* and *heterogeneous* 3D content imposes the need to *rethink the foundations of 3D computer graphics* to allow *uncertainty* to be treated in inherent manner in Computer Graphics, from design all the way to rendering and prototyping.

The technological shifts we mention above, together with developments in computer vision, user-friendly sketch-based modeling, online tutorials, but also image, video and 3D model repositories and 3D printing represent a great opportunity for new imaging methods. There are several significant challenges to overcome before such visual content can become widely accessible.

In GraphDeco, we have identified two major scientific challenges of our field which we will address:

- First, the design pipeline needs to be revisited to **explicitly account for the variability and uncertainty of a concept and its representations**, from early sketches to 3D models and prototypes. Professional practice also needs to be adapted and facilitated to be accessible to all.
- Second, a new approach is required to **develop computer graphics models and algorithms capable of handling uncertain and heterogeneous data** as well as traditional synthetic content.

We next describe the context of our proposed research for these two challenges. Both directions address hetereogeneous and uncertain input and (in some cases) output, and build on a set of common methodological tools.

3. Research Program

3.1. Introduction

Our research program is oriented around two main axes: 1) Computer-Assisted Design with Heterogeneous Representations and 2) Graphics with Uncertainty and Heterogeneous Content. These two axes are governed by a set of common fundamental goals, share many common methodological tools and are deeply intertwined in the development of applications.

3.1.1. Computer-Assisted Design with Heterogeneous Representations

Designers use a variety of visual representations to explore and communicate about a concept. Figure 2 illustrates some typical representations, including sketches, hand-made prototypes, 3D models, 3D printed prototypes or instructions.

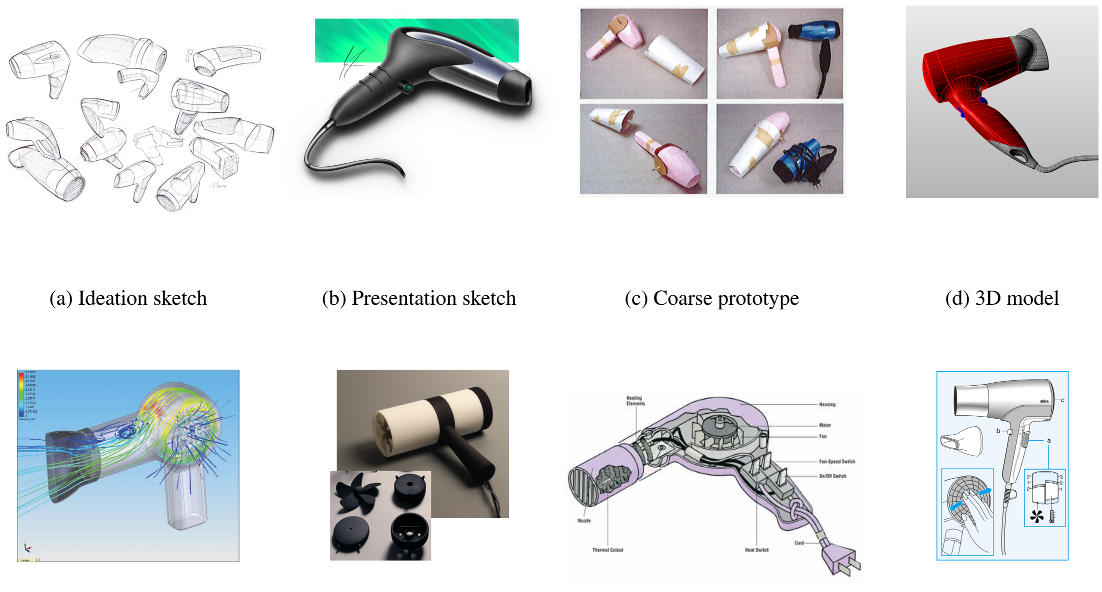

(a) Ideation sketch (b) Presentation sketch (c) Coarse prototype (d) 3D model

(f) Simulation (e) 3D Printing (g) Technical diagram (h) Instructions

Figure 2. Various representations of a hair dryer at different stages of the design process. Image source, in order: c-maeng on deviantart.com, shauntur on deviantart.com, "Prototyping and Modelmaking for Product Design" Hallgrimsson, B., Laurence King Publishers, 2012, samsher511 on turbosquid.com, my.solidworks.com, weilung tseng on cargocollective.com, howstuffworks.com, u-manual.com

The early representations of a concept, such as rough sketches and hand-made prototypes, help designers formulate their ideas and test the form and function of multiple design alternatives. These low-fidelity representations are meant to be cheap and fast to produce, to allow quick exploration of the *design space* of the concept. These representations are also often approximate to leave room for subjective interpretation and to stimulate imagination; in this sense, these representations can be considered *uncertain*. As the concept gets more finalized, time and effort are invested in the production of more detailed and accurate representations, such as high-fidelity 3D models suitable for simulation and fabrication. These detailed models can also be used to create didactic instructions for assembly and usage.

Producing these different representations of a concept requires specific skills in sketching, modeling, manufacturing and visual communication. For these reasons, professional studios often employ different experts to produce the different representations of the same concept, at the cost of extensive discussions and numerous iterations between the actors of this process. The complexity of the multi-disciplinary skills involved in the design process also hinders their adoption by laymen.

Existing solutions to facilitate design have focused on a subset of the representations used by designers. However, no solution considers all representations at once, for instance to directly convert a series of sketches into a set of physical prototypes. In addition, all existing methods assume that the concept is unique rather than ambiguous. As a result, rich information about the variability of the concept is lost during each conversion step.

We plan to facilitate design for professionals and laymen by adressing the following objectives:

- We want to assist designers in the exploration of the *design space* that captures the possible variations of a concept. By considering a concept as a *distribution* of shapes and functionalities rather than a single object, our goal is to help designers consider multiple design alternatives more quickly and effectively. Such a representation should also allow designers to preserve multiple alternatives along all steps of the design process rather than committing to a single solution early on and pay the price of this decision for all subsequent steps. We expect that preserving alternatives will facilitate communication with engineers, managers and clients, accelerate design iterations and even allow mass personalization by the end consumers.
- We want to support the various representations used by designers during concept development. While drawings and 3D models have received significant attention in past Computer Graphics research, we will also account for the various forms of rough physical prototypes made to evaluate the shape and functionality of a concept. Depending on the task at hand, our algorithms will either analyse these prototypes to generate a virtual concept, or assist the creation of these prototypes from a virtual model. We also want to develop methods capable of adapting to the different drawing and manufacturing techniques used to create sketches and prototypes. We envision design tools that conform to the habits of users rather than impose specific techniques to them.
- We want to make professional design techniques available to novices. Affordable software, hardware and online instructions are democratizing technology and design, allowing small businesses and individuals to compete with large companies. New manufacturing processes and online interfaces also allow customers to participate in the design of an object via mass personalization. However, similarly to what happened for desktop publishing thirty years ago, desktop manufacturing tools need to be simplified to account for the needs and skills of novice designers. We hope to support this trend by adapting the techniques of professionals and by automating the tasks that require significant expertise.

3.1.2. Graphics with Uncertainty and Heterogeneous Content

Our research is motivated by the observation that traditional CG algorithms have not been designed to account for uncertain data. For example, global illumination rendering assumes accurate virtual models of geometry, light and materials to simulate light transport. While these algorithms produce images of high realism, capturing effects such as shadows, reflections and interreflections, they are not applicable to the growing mass of uncertain data available nowadays.

The need to handle uncertainty in CG is timely and pressing, given the large number of *heterogeneous sources of 3D content* that have become available in recent years. These include data from cheap depth+image sensors (e.g., Kinect or the Tango), 3D reconstructions from image/video data, but also data from large 3D geometry databases, or casual 3D models created using simplified sketch-based modeling tools. Such alternate content has varying levels of *uncertainty* about the scene or objects being modelled. This includes uncertainty in geometry, but also in materials and/or lights – which are often not even available with such content. Since CG algorithms cannot be applied directly, visual effects artists spend hundreds of hours correcting inaccuracies and completing the captured data to make them useable in film and advertising.

We identify a major scientific bottleneck which is the need to treat *heterogeneous* content, i.e., containing both (mostly captured) uncertain and perfect, traditional content. Our goal is to provide solutions to this bottleneck, by explicitly and formally modeling uncertainty in CG, and to develop new algorithms that are capable of mixed rendering for this content.

Figure 3. Image-Based Rendering (IBR) techniques use input photographs and approximate 3D to produce new synthetic views.

We strive to develop methods in which heterogeneous – and often uncertain – data can be handled automatically in CG with a principled methodology. Our main focus is on *rendering* in CG, including dynamic scenes (video/animations).

Given the above, we need to address the following challenges:

- Develop a theoretical model to handle uncertainty in computer graphics. We must define a new formalism that inherently incorporates uncertainty, and must be able to express traditional CG rendering, both physically accurate and approximate approaches. Most importantly, the new formulation must elegantly handle mixed rendering of perfect synthetic data and captured uncertain content. An important element of this goal is to incorporate *cost* in the choice of algorithm and the optimizations used to obtain results, e.g., preferring solutions which may be slightly less accurate, but cheaper in computation or memory.
- The development of rendering algorithms for heterogeneous content often requires preprocessing of image and video data, which sometimes also includes depth information. An example is the decomposition of images into intrinsic layers of reflectance and lighting, which is required to perform relighting. Such solutions are also useful as image-manipulation or computational photography techniques. The challenge will be to develop such "intermediate" algorithms for the uncertain and heterogeneous data we target.
- Develop efficient rendering algorithms for uncertain and heterogeneous content, reformulating rendering in a probabilistic setting where appropriate. Such methods should allow us to develop approximate rendering algorithms using our formulation in a well-grounded manner. The formalism should include probabilistic models of how the scene, the image and the data interact. These models should be data-driven, e.g., building on the abundance of online geometry and image databases, domain-driven, e.g., based on requirements of the rendering algorithms or perceptually guided, leading to plausible solutions based on limitations of perception.

4. Highlights of the Year

4.1. Highlights of the Year

This year marked the start of the ERC Starting grant D^3 coordinated by Adrien Bousseau, on interpreting drawing for 3D design. This activity already includes the principal investigator, one postdoc (Y. Gryaditskaya) and one engineer (B. Wailly) and will be growing over the next year. The scientific production this year included four ACM Transactions on graphics papapers (three at SIGGRAPH and one at SIGGRAPH Asia, of which 2 were work performed by our visiting International Chair F. Durand during his stay in our group), one paper at Eurographics and several other top-level publications.

5. New Software and Platforms

5.1. SGTDGP

Synthetic Ground Truth Data Generation Platform
KEYWORD: Graphics
FUNCTIONAL DESCRIPTION: The goal of this platform is to render large numbers of realistic synthetic images for use as ground truth to compare and validate image-based rendering algorithms and also to train deep neural networks developed in our team.

This pipeline consists of tree major elements that are:

- Scene exporter
- Assisted point of view generation
- Distributed rendering on Inria's high performance computing cluster

The scene exporter is able to export scenes created in the widely-used commercial modeler 3DSMAX to the Mitsuba opensource renderer format. It handles the conversion of complex materials and shade trees from 3DSMAX including materials made for VRay. The overall quality of the produced images with exported scenes have been improved thanks to a more accurate material conversion. The initial version of the exporter was extended and improved to provide better stability and to avoid any manual intervention.

From each scene we can generate a large number of images by placing multiple cameras. Most of the time those points of view has to be placed with a certain coherency. This task could be long and tedious. In the context of image-based rendering, cameras have to be placed in a row with a specific spacing. To simplify this process we have developed a set of tools to assist the placement of hundreds of cameras along a path.

The rendering is made with the open source renderer Mitsuba. The rendering pipeline is optimised to render a large number of point of view for single scene. We use a path tracing algorithm to simulate the light interaction in the scene and produce hight dynamic range images. It produces realistic images but it is computationally demanding. To speed up the process we setup an architecture that takes advantage of the Inria cluster to distribute the rendering on hundreds of CPUs cores.

The scene data (geometry, textures, materials) and the cameras are automatically transfered to remote workers and HDR images are returned to the user.

We already use this pipeline to export tens of scenes and to generate several thousands of images, which have been used for machine learning and for ground-truth image production.

- Contact: George Drettakis

5.2. Unity IBR

KEYWORD: Graphics
FUNCTIONAL DESCRIPTION: Unity IBR (for Image-Based Rendering in Unity) This is a software module that proceeds the development of IBR algorithms in Unity. In this case, algorithms are developed for the context of EMOTIVE EU project. The rendering technique was changed during the year to evaluate and compare which one produces better results suitable for Game Development with Unity (improvement of image quality and faster rendering). New features were also added such as rendering of bigger datasets and some debugging utilities. Software was also updated to keep compatibility with new released versions of Unity game engine.

- Contact: George Drettakis

5.3. SIBR

Simple Image-Based Rendering
KEYWORD: Graphics

FUNCTIONAL DESCRIPTION: This is a framework containing libraries and tools used internally for research projects based on Image-Base Rendering. It includes both preprocessing tools (computing data used for rendering) and rendering utilities and serves as the basis for many research projects in the group.

It includes basic support for a large set of computer graphics and computer vision functionalities and includes implementations of several image-based rendering algorithms. The code base has become quite mature and is in the process of being used for tech transfer.

- Contact: George Drettakis

6. New Results

6.1. Computer-Assisted Design with Heterogeneous Representations

6.1.1. Patterns from Photograph: Reverse-Engineering Developable Products

Participants: Adrien Bousseau.

Developable materials are ubiquitous in design and manufacturing. Unfortunately, general-purpose modeling tools are not suited to modeling 3D objects composed of developable parts. We propose an interactive tool [13] to model such objects from a photograph (Fig. 4). Users of our system load a single picture of the object they wish to model, which they annotate to indicate silhouettes and part boundaries. Assuming that the object is symmetric, we also ask users to provide a few annotations of symmetric correspondences. The object is then automatically reconstructed in 3D. At the core of our method is an algorithm to infer the 2D projection of rulings of a developable surface from the traced silhouettes and boundaries. We impose that the surface normal is constant along each ruling, which is a necessary property for the surface to be developable. We complement these developability constraints with symmetry constraints to lift the curve network in 3D. In addition to a 3D model, we output 2D patterns enabling to fabricate real prototypes of the object on the photo. This makes our method well suited for reverse engineering products made of leather, bent cardboard or metal sheets.

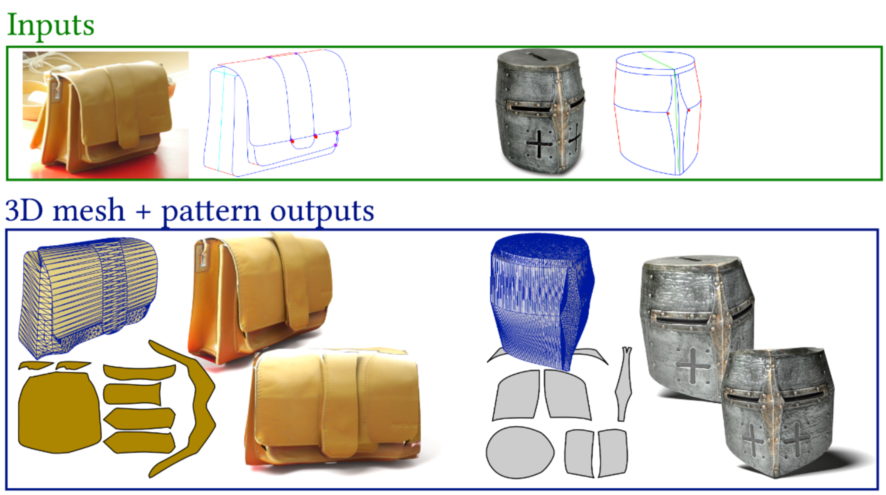

Figure 4. Given an annotated line drawing traced over the photograph of a piecewise-developable object (top), our method produces a 3D model of the visible parts of the object, and of some of the occluded parts inferred by symmetry. Our method guarantees developability of the model, as shown by the flattened patterns (bottom).

This work is a collaboration with Amélie Fondevilla, Damien Rohmer, Stefanie Hahmann and Marie-Paule Cani from the IMAGINE team at Inria Rhône Alpes. The work was published in the special issue of the journal Computers and Graphics (Elsevier), presented at the Shape Modeling International conference.

6.1.2. SketchSoup: Exploratory Ideation Using Design Sketches

Participants: Adrien Bousseau.

A hallmark of early stage design is a number of quick-and-dirty sketches capturing design inspirations, model variations, and alternate viewpoints of a visual concept. We developed SketchSoup [7], a workflow that allows designers to explore the design space induced by such sketches (Fig. 5). We take an unstructured collection of drawings as input, along with a small number of user-provided correspondences as input. We register them using a multi-image matching algorithm, and present them as a 2D interpolation space. By morphing sketches in this space, our approach produces plausible visualizations of shape and viewpoint variations despite the presence of sketch distortions that would prevent standard camera calibration and 3D reconstruction. In addition, our interpolated sketches can serve as inspiration for further drawings, which feed back into the design space as additional image inputs. SketchSoup thus fills a significant gap in the early ideation stage of conceptual design by allowing designers to make better informed choices before proceeding to more expensive 3D modeling and prototyping. From a technical standpoint, we describe an end-to-end system that judiciously combines and adapts various image processing techniques to the drawing domain – where the images are dominated not by color, shading and texture, but by sketchy stroke contours.

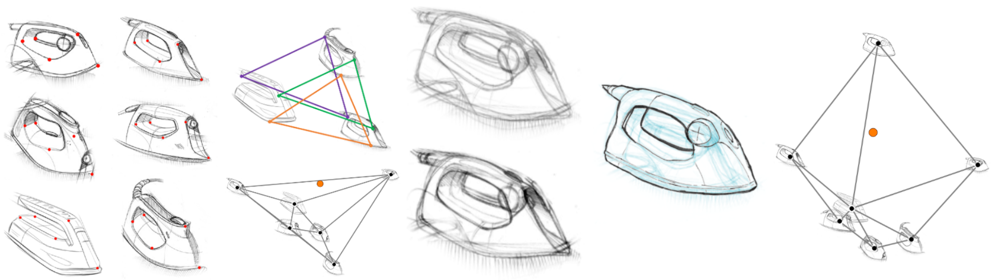

(a) *Input sketch images* **(b)** *Registration (top) and interpolation space (bottom)* **(c)** *Warping (top) and non-uniform blending (bottom)* **(d)** *Oversketching* **(e)** *Augmenting the interpolation space*

Figure 5. SketchSoup takes an unstructured set of sketches as input, along with a small number of correspondences (shown as red dots)(a), registers the sketches using an iterative match-warp algorithm harnessing matching consistency across images (b, top), and embeds the sketches into a 2D interpolation space based on their shape differences (b, bottom). Users can explore the interpolation space to generate novel sketches, which are generated by warping existing sketches into alignment (c, top), followed by spatially non-uniform blending (c, bottom). These interpolated sketches can serve as underlay to inspire new concepts (d), which can in turn be integrated into the interpolation space to iteratively generate more designs (e).

This work is a collaboration with Rahul Arora and Karan Singh from Toronto University (Canada) and Ishan Darolia and Vinay P. Namboodiri from IIT Kampur (India). The work was published in the journal Computer Graphics Forum and presented at the Eurographics conference.

6.1.3. Photo2ClipArt: Image Abstraction and Vectorization Using Layered Linear Gradients

Participants: Adrien Bousseau, Jean-Dominique Favreau.

We present a method to create vector cliparts from photographs [8]. Our approach aims at reproducing two key properties of cliparts: they should be easily editable, and they should represent image content in a clean, simplified way. We observe that vector artists satisfy both of these properties by modeling cliparts with linear color gradients, which have a small number of parameters and approximate well smooth color variations. In addition, skilled artists produce intricate yet editable artworks by stacking multiple gradients using opaque and semi-transparent layers. Motivated by these observations, our goal is to decompose a bitmap photograph into a stack of layers, each layer containing a vector path filled with a linear color gradient. We cast this problem as an optimization that jointly assigns each pixel to one or more layer and finds the gradient parameters of each layer that best reproduce the input. Since a trivial solution would consist in assigning each pixel to a different, opaque layer, we complement our objective with a simplicity term that favors decompositions made of few, semi-transparent layers. However, this formulation results in a complex combinatorial problem combining discrete unknowns (the pixel assignments) and continuous unknowns (the layer parameters). We propose a Monte Carlo Tree Search algorithm that efficiently explores this solution space by leveraging layering cues at image junctions. We demonstrate the effectiveness of our method by reverse-engineering existing cliparts and by creating original cliparts from studio photographs.

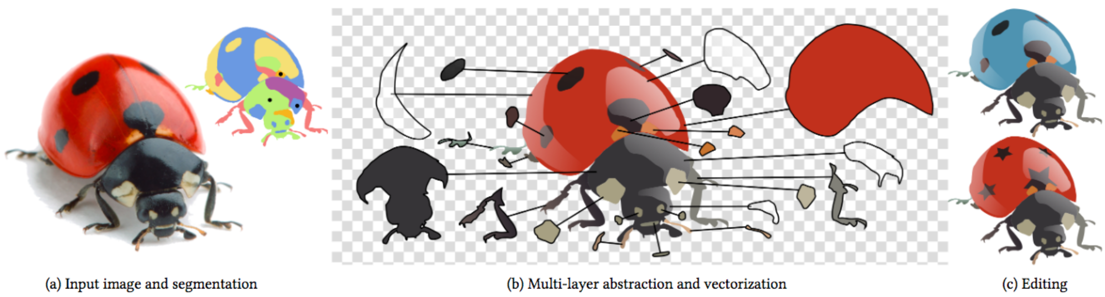

Figure 6. Given a segmented bitmap image as input (a), our method generates an abstract, layered vector clipart, where each layer is filled with an opaque or semi-transparent linear color gradient (b). By expressing the image as a stack of linear color gradients, our vector graphics reproduce the visual style of traditional cliparts and are easy to edit (c). In this example, we turned the lady bug blue (top) and replaced its dots by little stars (bottom). The black dots in the segmentation indicate opaque regions selected by the user to initialize or constrain our algorithm.

This work is a collaboration with Florent Lafarge from the Titane team at Inria Sophia Antipolis. The work was published in the journal ACM Transactions on Graphics and presented at the SIGGRAPH Asia conference.

6.1.4. Data Collection and Analysis of Industrial Design Drawings

Participants: Yulia Gryaditskaya (post-doctoral researcher), Adrien Bousseau (permanent researcher) and Fredo Durand (visiting researcher).

The goal of this project is to collect a dataset of industrial design drawings uniquely matched to a 3D geometry and accompanied by metadata such as each stroke position, time of creation and pressure. We are interested in creation of a classification of lines the designers use and analysis of their relation to the underlying shape. We are further interested in evaluation of correlation between the presence of construction techniques and perspective accuracy of the sketch. We are planning to provide a ground-truth labelling guided by experience of professional designers which will enable development of multiple algorithms, such as sketch beautification and vectorization, style transfer and 3D inference.

This work is a collaboration with Mark Sypesteyn, Jan Willem Hoftijzer and Sylvia Pont from TU Delft, Netherlands.

6.1.5. What You Sketch Is What You Get: 3D Sketching using Multi-View Deep Volumetric Prediction

Participants: Johanna Delanoy, Adrien Bousseau

Drawing is the most direct way for people to express their visual thoughts. However, while humans are extremely good at perceiving 3D objects from line drawings, this task remains very challenging for computers as many 3D shapes can yield the same drawing. Existing sketch-based 3D modeling systems rely on heuristics to reconstruct simple shapes, require extensive user interaction, or exploit specific drawing techniques and shape priors. Our goal is to lift these restrictions and offer a minimal interface to quickly model general 3D shapes with contour drawings. While our approach can produce approximate 3D shapes from a single drawing, it achieves its full potential once integrated into an interactive modeling system, which allows users to visualize the shape and refine it by drawing from several viewpoints. At the core of our approach is a deep convolutional neural network (CNN) that processes a line drawing to predict occupancy in a voxel grid. The use of deep learning results in a flexible and robust 3D reconstruction engine that allows us to treat sketchy bitmap drawings without requiring complex, hand-crafted optimizations. While similar architectures have been proposed in the computer vision community, our originality is to extend this architecture to a multiview context by training an updater network that iteratively refines the prediction as novel drawings are provided.

This work is a collaboration with Mathieu Aubry from Ecole des Ponts ParisTech and Alexei Efros and Philip Isola from UC Berkeley. It is supported by the CRISP Inria associate team.

6.1.6. Flatenning videos for stylization

Participants: Johanna Delanoy, Adrien Bousseau

Traditional 2D animation exhibits a very specific sense of motion where objects seems to move in a 2D world. Existing stylization methods for videos use the optical flow to ensure temporal consistency of the stylization process. Although the method produces very convincing results, the resulting video often looks like a 3D textured scene instead of a 2D animation. In this project, we propose to transform the input video into a new one with a simplified motion. To achieve this effect, we approximate the motion with 2D rigid patches (rigid motion and scaling): each frame is segmented into rigid motion patches that provide a good approximation of the initial motion. The final sequence exhibits a flattened motion and can be used in any stylization process. This produces a stylized video that has a feeling of 2D motion and is more similar to traditional animation.

This work is a collaboration with Aaron Hertzmann from Adobe Research.

6.2. Graphics with Uncertainty and Heterogeneous Content

6.2.1. Image based rendering of repetitive facades

Participants: Simon Rodriguez, Adrien Bousseau, Frederic Durand, George Drettakis.

Image Based Rendering techniques (IBR) rely on interpolation between multiple views of a scene to generate new viewpoints. One of the main requirements is that the density of input views be high enough to obtain satisfying coverage and resolution, in both preprocessing and rendering. In human-made scenes, repetitive elements can be used to alleviate this limitation when the baseline of input views is sparse. A small number of viewpoints of similar elements can be fused together to extract more information about the scene. We focus on buildings facades, which often exhibit repetitive architectural elements. We propose the following steps: such elements are extracted from input views, and used during pre-processing to generate an approximate geometry of the specific element, and to extract view-dependent effects. These data sources are then combined to perform improved rendering in an IBR context.

6.2.2. Plane-Based Multi-View Inpainting for Image-Based Rendering in Large Scenes

Participants: Julien Philip, George Drettakis.

Image-Based Rendering (IBR) allows high-fidelity free-viewpoint navigation using only a set of photographs and 3D reconstruction as input. It is often necessary or convenient to remove objects from the captured scenes to allow a form of scene editing for IBR. This requires multi-view inpainting of the input images. Previous methods suffer from several major limitations: they do not impose true multi-view coherence, resulting in artifacts such as blur, they do not preserve perspective during inpainting, provide inaccurate depth completion and can only handle scenes with a few tens of images. Our approach addresses these limitations by introducing a new multi-view method that performs inpainting in intermediate, locally common planes. Use of these planes results in correct perspective and multi-view coherence of inpainting results. For efficient treatment of large scenes, we present a fast planar region extraction method operating on small image clusters. We adapt the resolution of inpainting to that required in each input image of the multi-view dataset, and carefully handle image resampling between the input images and rectified planes. Our method can handle up to hundreds of input images, for indoors and outdoors environments.

6.2.3. Thin structures in Image Based Rendering

Participants: Theo Thonat, Abdelaziz Djelouah, Frederic Durand, George Drettakis.

One of the key problem in Image Based Rendering (IBR) methods is the rendering of regions with incorrect 3D reconstruction. Thin structures, with their lack of texture and distinctive features, are another important common source of 3D reconstruction errors. They are present in most urban pictures and represent a standard failure case for reconstruction algorithms, and state of the art rendering methods exhibit strong artifacts. In this project, we propose to detect and segment fences in urban setup for IBR applications. We use the assumption that thin structures lie on a 3D surface. We propose a multi-view approach to compute the thin structures images segmentation and its associated alpha matting. Finally, we propose also a new IBR algorithm to render these thin structures.

6.2.4. Handling reflections in Image-Based Rendering

Participants: Theo Thonat, Frederic Durand, George Drettakis.

In order to render new viewpoints, current Image Based Rendering (IBR) techniques use approximate geometry to warp and blend images from close viewpoints. They assume the scene materials are mostly diffuse, and they assume only a direct look at the geometry is enough. These assumptions fail in the case of specular surfaces such as windows. Dealing with reflections in an IBR context first requires identifying what are the diffuse and the specular color layers in the input images. The challenge is then to correctly warp the specular layers since the normals of the reflective surfaces might be not reliable.

6.2.5. Material capture

Participants Valentin Deschaintre, Miika Aittala, Frederic Durand, George Drettakis, Adrien Bousseau.
Convenient material acquisition is a complicated process, current methods are based on strong assumptions or limitations. Acquisition of spatially varying materials complete models is currently limited to specific materials or complex setups requiring multiple pictures with varying light and view positions.

This work aims at acquiring a material's Spatially-Varying BRDF using a single flash picture. We introduce procedural synthetic data generation and deep learning to mitigate the need of material and environment assumptions.

With fewer pictures used in the material acquisition process, comes more ambiguities in the explanation of the lighting behaviour. With one shot acquisition, important effects can be missed or misunderstood because of the lighting or view point. Current "lightweight" methods use various assumptions regarding materials. We use training to learn important ambiguities and how to solve them from the dataset, giving our network the capacity to handle the inherent uncertainty of one picture material acquisition.

6.2.6. Accommodation and Comfort in Head-Mounted Displays

Participants: George-Alex Koulieris, George Drettakis.

Head-mounted displays (HMDs) often cause discomfort and even nausea. Improving comfort is therefore one of the most significant challenges for the design of such systems. We evaluated the effect of different HMD display configurations on discomfort. We did this by designing a device to measure human visual behavior and evaluate viewer comfort. In particular, we focused on one known source of discomfort: the vergence-accommodation (VA) conflict. The VA conflict is the difference between accommodative and vergence response. In HMDs the eyes accommodate to a fixed screen distance while they converge to the simulated distance of the object of interest, requiring the viewer to undo the neural coupling between the two responses. Several methods have been proposed to alleviate the VA conflict, including Depth-of-Field (DoF) rendering, focus-adjustable lenses, and monovision. However, no previous work had investigated whether these solutions actually drive accommodation to the distance of the simulated object. If they did, the VA conflict would disappear, and we would expect comfort to improve. We designed the first device that allows us to measure accommodation in HMDs, and we used it to obtain accommodation measurements and to conduct a discomfort study, see Fig. 7. The results of the first experiment demonstrated that only the focus-adjustable-lens design drives accommodation effectively, while other solutions do not drive accommodation to the simulated distance and thus do not resolve the VA conflict. The second experiment measured discomfort. The results validated that the focus-adjustable-lens design improves comfort significantly more than the other solutions.

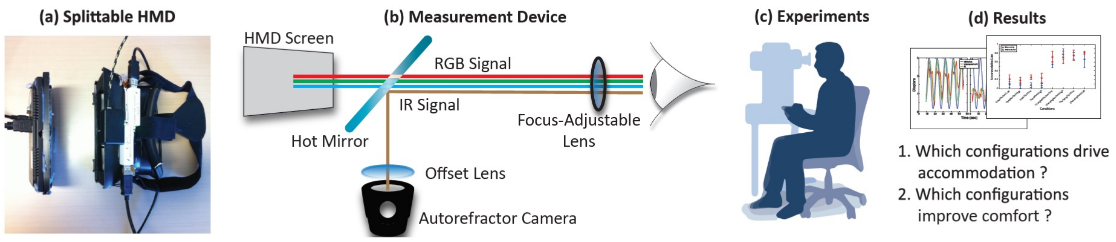

Figure 7. We evaluated the effect of different Head-Mounted Display (HMD) display configurations on discomfort by developing a splittable HMD with focus-adjustable lenses (a) that works with an autorefractor to objectively measure accommodation (b). (c) We run experiments to evaluate the effect of different display configurations. (d) Our results allowed us to answer fundamental questions about each configuration.

This work [11] is a collaboration with Martin S. Banks and Bee Bui from University of California, Berkeley. The work was published in a regular issue of the journal Transactions on Graphics and presented at the ACM SIGGRAPH conference in Los Angeles, USA, 2017.

6.2.7. Focus-tunable and Fixed Lenses and Stereoscopic 3D Displays

Participants: George-Alex Koulieris, George Drettakis.

Stereoscopic 3D (S3D) displays provide an enhanced sense of depth by sending different images to the two eyes. But these displays do not reproduce focus cues (blur and accommodation) correctly. Specifically, the eyes must accommodate to the display screen to create sharp retinal images even when binocular disparity drives the eyes to converge to other distances. This mismatch causes discomfort, reduces performance, and distorts 3D percepts. We developed two techniques designed to reduce vergence-accommodation conflicts and thereby improve comfort, performance, and perception. One uses focus-tunable lenses between the display and viewer's eyes. Lens power is yoked to expected vergence distance creating a stimulus to accommodation that is consistent with the stimulus to vergence. This yoking should reduce the vergence-accommodation mismatch. The other technique uses a fixed lens before one eye and relies on binocularly fused percepts being determined by one eye and then the other, depending on simulated distance. This is meant to drive accommodation with one eye when simulated distance is far and with the other eye when simulated distance is near. We conducted performance tests and discomfort assessments with both techniques and with conventional S3D displays (see Fig. 8). We also measured accommodation. The focus-tunable technique, but not the fixed-lens technique, produced appropriate stimulus-driven accommodation thereby minimizing the vergence-accommodation conflict. Because of this, the tunable technique yielded clear improvements in comfort and performance while the fixed technique did not. The focus-tunable lens technique therefore offers a relatively easy means for reducing the vergence-accommodation conflict and thereby improving viewer experience.

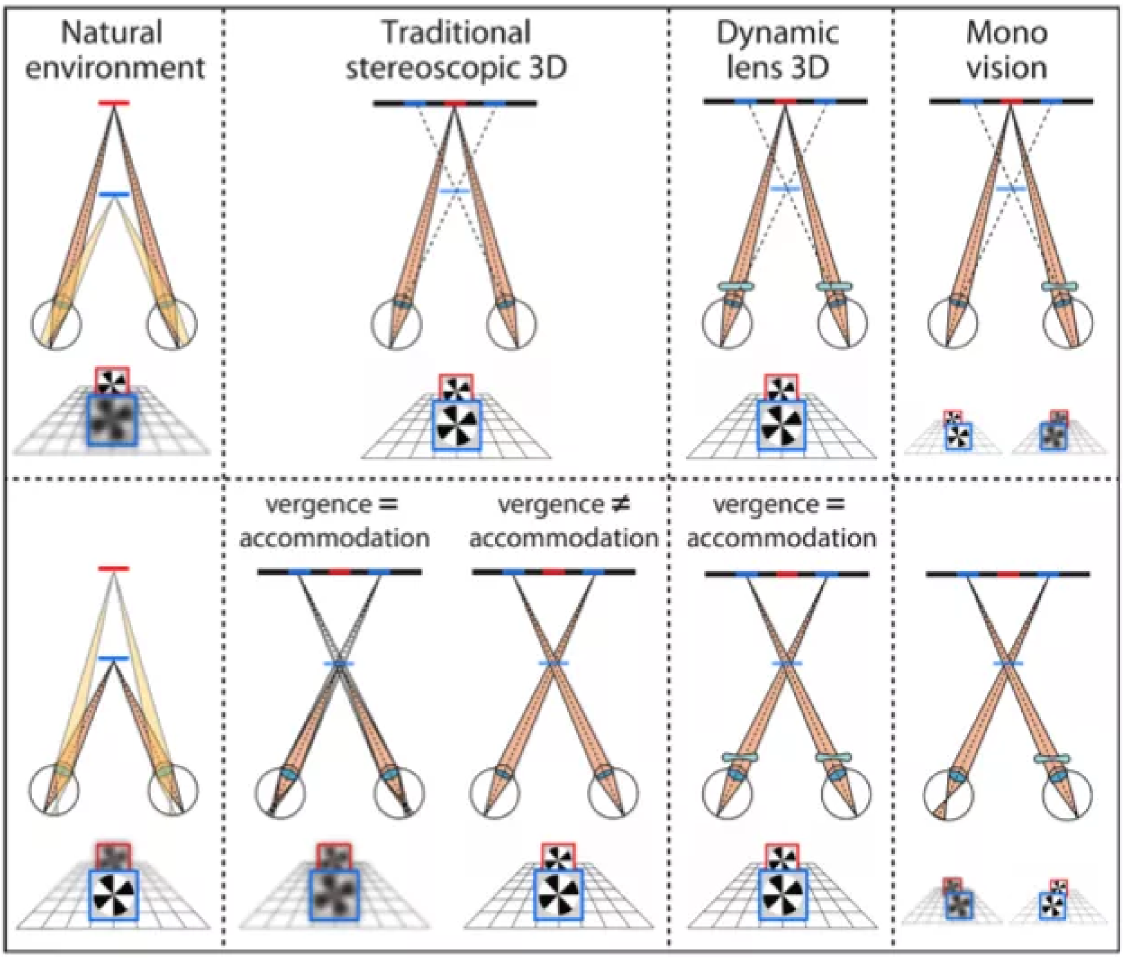

Figure 8. Focus cues in different viewing conditions for natural scenes, a conventional S3D display, and our proposed displays. In each panel, the upper part is an overhead view of the situation and the lower part is a schematic of what a viewer would see. In natural viewing (left column), vergence and accommodation distances are the same. In conventional stereoscopic 3D displays (second column), vergence distance varies with the disparity of the simulated object while accommodation distance is fixed at the display screen. In the proposed dynamic-lens display (third column), accommodation distance can be adjusted depending on the content being displayed. This is implemented by changing the power of a lens in front of each eye. In this case, accommodation and vergence distances are matched. In the proposed monovision display (right column), fixed lenses of different powers are placed in front of the two eyes. There are two accommodative distances that match the vergence distance, one for the left eye and one for the right.

This work is a collaboration with Martin S. Banks from University of California, Berkeley, Paul V. Johnson from Apple, Joohwan Kim, Nvidia, Jared AQ Parnell and Gordon D. Love, Durham University, UK. The work was published in the proceedings of SPIE and presented at SPIE Opto 2017 conference, San Francisco, USA.

6.2.8. Applications of Visual Perception to Virtual Reality Rendering

Participants: George-Alex Koulieris

Over the past few years, virtual reality (VR) has transitioned from the realm of expensive research prototypes and military installations into widely available consumer devices. But the high pixel counts and frame rates of current commodity devices more than double the rendering costs of 1920x1080 gaming, and next-generation HMDs could easily double or triple costs again. As a result, VR experiences are limited in visual quality, performance, and other capabilities. Human visual perception has repeatedly been shown to be important to creating immersion while keeping up with increasing performance requirements. Thus, an understanding of visual perception and its applications in real-time VR graphics is vital for HMD designers, application developers, and content creators. In this course we began with an overview of the importance of human perception in modern virtual reality. We accompanied this overview with a dive into the key characteristics of the human visual system and the psychophysical methods used to study its properties. After laying the perceptual groundwork, we presented three case studies outlining the applications of human perception to improving the performance, quality, and applicability of VR graphics. Finally, we concluded with a forward looking discussion, highlighting important future work and open challenges in perceptual VR, and a questions session for more in-depth audience interaction.

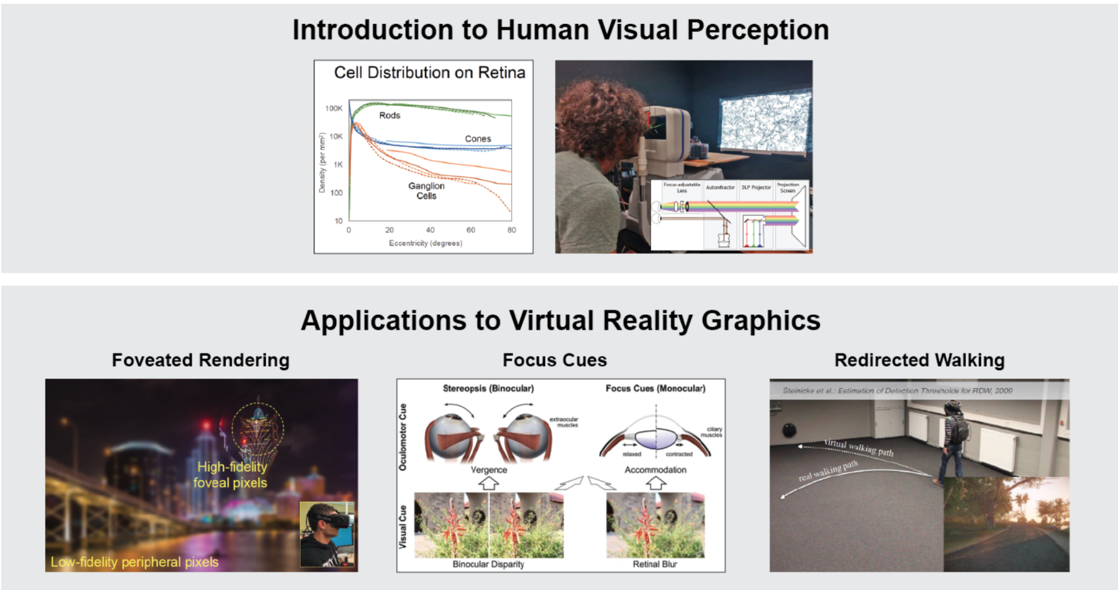

Figure 9. An overview of the topics covered in the course.

This work is a collaboration with Anjul Patney and Joohwan Kim, Nvidia, Marina Zanolli, Oculus VR, Gordon Wetzstein, Stanford University and Frank Steinicke, University of Hamburg. The course was presented at ACM SIGGRAPH 2017, Los Angeles, USA.

6.2.9. Real-time Binocular Tone Mapping

Participants: George-Alex Koulieris, George Drettakis

In real-world scenes, luminance values can simultaneously span several orders of magnitude. The human visual system has evolved such it can comfortably perceive this vast luminance gamut via adaptation mechanisms both in the eyes and the brain. Conventional displays can not reproduce such enormous dynamic ranges since their contrast ratio is severely limited. As a result, simulations of natural scenes significantly deviate from real-world percepts. Tone-mapping techniques approximate the appearance of high dynamic range scenes in limited dynamic range media. However, there exists an inherent trade-off between simultaneously inducing both a high dynamic range perception and preserving contrast in fine-scale details. We will alleviate this trade-off by capitalizing on the properties of the visual system's binocular fusion processes during dichoptic stimulation, i.e., when a different image is shown to each eye.

To be submitted: This work is an ongoing collaboration with Martin S. Banks, University of California, Berkeley, Rafal Mantiuk, University of Cambridge, and Fredo Durand, MIT.

6.2.10. A novel objective method for the assessment of binocular accommodative facility: A pilot study in a young adult population

Participants: George-Alex Koulieris

Accommodative facility (AF) is a clinical test used to evaluate the ability of the visual system to alter accommodation rapidly and accurately when the dioptric stimulus to accommodation is situated between two different distances. AF can be evaluated in monocular and binocular testing, providing a direct evaluation of the dynamics of the accommodative response. We investigated the validity of measuring the binocular accommodative facility using Hart charts in conjunction with the Grand Seiko Auto Ref/Keratometer in young participants. The main objective of the study was to propose a novel objective method to assess binocular accommodative facility in quantitative terms.

To be submitted: This work is an ongoing collaboration with Jesus Vera-Vilchez and Raimundo Jimenez Rodriguez, University of Granada, Spain.

6.2.11. Efficient Thin Shell Sounds

Participants: Gabriel Cirio, George Drettakis.

Thin shells, i.e. solids that are thin in one dimension compared to the other two, often produce rich and recognizable sounds when struck: from containers like a trash can or a plastic bottle, to musical instruments like a cymbal or a gong. Thin shells are notoriously difficult and expensive to simulate due to their nonlinear behavior under large excitations. To synthesize the sound generated by thin shells, we first reduce the problem to a small modal subspace, and compute all the required quantities directly in the subspace. We then further speed up the simulation by computing nonlinear dynamics using only a subset of low frequency vibrations, since these are responsible for most nonlinear phenomena, and use a frequency coupling to drive the remaining high frequency vibrations. Finally, we approximate the chaotic regime that can emerge in some thin shells (such as gongs and cymbals) by diffusing the power spectrum of the sound directly in the frequency domain, following a phenomenological approach to wave turbulence. We can produce rich and complex sounds for a wide range of behaviors with a computational cost orders of magnitude smaller than previous approaches.

This work is an ongoing collaboration with Changxi Zheng and Etian Grinspun from Columbia University in New York, and is supported by the EU H2020 Marie Sklodowska-Curie project PhySound.

6.2.12. Aether: An Embedded Domain Specific Sampling Language for Monte Carlo Rendering

Participant: Frederic Durand

Implementing Monte Carlo integration requires signifcant domain expertise. While simple samplers, such as unidirectional path tracing, are relatively forgiving, more complex algorithms, such as bidirectional path tracing or Metropolis methods, are notoriously diffcult to implement correctly.We propose Aether, an embedded domain specifc language for Monte Carlo integration, which offers primitives for writing concise and correct-by-construction sampling and probability code. The user is tasked with writing sampling code, while our compiler automatically generates the code necessary for evaluating PDFs as well as the book keeping and combination of multiple sampling strategies. Our language focuses on ease of implementation for rapid exploration, at the cost of run time performance. We demonstrate the effectiveness of the language by implementing several challenging rendering algorithms as well as a new algorithm, which would otherwise be prohibitively difficult.

The work [6] was published in a regular issue of the journal Transactions on Graphics and presented at the ACM SIGGRAPH conference in Los Angeles, USA, 2017.

6.2.13. Deep bilateral learning for real-time image enhancement

Participant: Frederic Durand

Performance is a critical challenge in mobile image processing. Given a reference imaging pipeline, or even human-adjusted pairs of images, we seek to reproduce the enhancements and enable real-time evaluation. For this, we introduce a new neural network architecture inspired by bilateral grid processing and local affine color transforms. Using pairs of input/output images, we train a convolutional neural network to predict the coefficients of a locally-affine model in bilateral space. Our architecture learns to make local, global, and content-dependent decisions to approximate the desired image transformation. At runtime, the neural network consumes a low-resolution version of the input image, produces a set of affine transformations in bilateral space, upsamples those transformations in an edge-preserving fashion using a new slicing node, and then applies those upsampled transformations to the full-resolution image. Our algorithm processes high-resolution images on a smartphone in milliseconds, provides a real-time viewfinder at 1080p resolution, and matches the quality of state-of-the-art approximation techniques on a large class of image operators. Unlike previous work, our model is trained off-line from data and therefore does not require access to the original operator at runtime. This allows our model to learn complex, scene-dependent transformations for which no reference implementation is available, such as the photographic edits of a human retoucher.

The work [10] was published in a regular issue of the journal Transactions on Graphics and presented at the ACM SIGGRAPH conference in Los Angeles, USA, 2017.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

7.1.1. Optis

Participants: Valentin Deschaintre, Adrien Bousseau, George Drettakis.

Valentin Deschaintre has a CIFRE PhD fellowship on Material Acquisition using Machine Learning, in collaboration with Optis, a company specialized in material acquisition and rendering.

7.1.2. Adobe

Participants: Adrien Bousseau, Johana Delanoy.

As part of a long standing collaboration with Adobe, J. Delanoy interned at Adobe Research with A. Hertzman, (San Fransisco). Adobe provides research and software donations as part of this collaboration.

7.1.3. Technicolor

Participants: George Drettakis, Adrien Bousseau.

We have initiated a collaboration with Technicolor on the use of deep learning for computational photography and video tasks. This involves the use of our synthetic ground truth data generation platform for graphics and vision tasks.

8. Partnerships and Cooperations

8.1. Regional Initiatives

Theo Thonat is funded in part by a Region PACA fellowship.

8.2. National Initiatives

8.2.1. ADT PicPlay

Participants: Sebastien Bonopera, George Drettakis.

The Technology Development Action (ADT) PicPlay a technology tranfer pre-maturation project, supported by Inria and by UCA Jedi. The objective is to create a startup company based on image based rendering technologies, taking benefit from the team's research and experience over the last 8 years. At this early stage, we evaluated the market and produced several Proof-of-Concept demonstrations for potential clients. One of the demonstrations is our new asset streaming capability that allows the use for huge datasets (see Fig. 10).

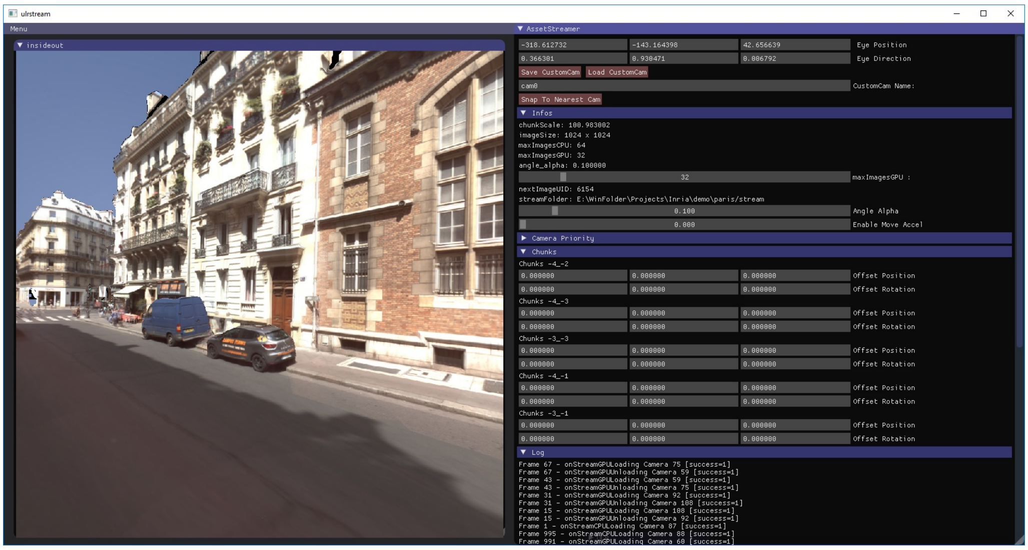

Figure 10. streamable dataset using 6500 pictures (usual not-streamable dataset has around 30 pictures)

We also developed a new solution to improve rendering quality. This solution uses a 3D mesh for each view and refines it according to this view only, before blending each view. Finally, PicPlay involved the development of several tools for converting and processing datasets.

8.2.2. ANR

8.2.2.1. ANR SEMAPOLIS

Participants: George Drettakis, Abdelaziz Djelouah, Theo Thonat.

This ANR project ended in September 2017. The goal was to use semantic information to improve urban reconstruction and rendering. The consortium was led by ENPC (R. Marlet) and includes the Inria Willow team and the GREY-C laboratory on image processing. Our contribution was in the rendering of urban models, in particular using image-based rendering algorithms. Our contribution resulted in several publications or planned publications (e.g., those described in Sec. 6.2.2, 6.2.3)

8.3. European Initiatives

8.3.1. FP7 & H2020 Projects

8.3.1.1. D^3: Interpreting Drawings for 3D Design

Type: ERC

Instrument: Starting Grant

Duration: February 2017 - February 2023

Participants: Adrien Bousseau, Yulia Gryaditskaya, Bastien Wailly.

Abstract. Designers draw extensively to externalize their ideas and communicate with others. However, drawings are currently not directly interpretable by computers. To test their ideas against physical reality, designers have to create 3D models suitable for simulation and 3D printing. However, the visceral and approximate nature of drawing clashes with the tediousness and rigidity of 3D modeling. As a result, designers only model finalized concepts, and have no feedback on feasibility during creative exploration. Our ambition is to bring the power of 3D engineering tools to the creative phase of design by automatically estimating 3D models from drawings. However, this problem is ill-posed: a point in the drawing can lie anywhere in depth. Existing solutions are limited to simple shapes, or require user input to ”explain” to the computer how to interpret the drawing. Our originality is to exploit professional drawing techniques that designers developed to communicate shape most efficiently. Each technique provides geometric constraints that help viewers understand drawings, and that we shall leverage for 3D reconstruction.

Our first challenge is to formalize common drawing techniques and derive how they constrain 3D shape. Our second challenge is to identify which techniques are used in a drawing. We cast this problem as the joint optimization of discrete variables indicating which constraints apply, and continuous variables representing the 3D model that best satisfies these constraints. But evaluating all constraint configurations is impractical. To solve this inverse problem, we will first develop forward algorithms that synthesize drawings from 3D models. Our idea is to use this synthetic data to train machine learning algorithms that predict the likelihood that constraints apply in a given drawing. In addition to tackling the long-standing problem of single-image 3D reconstruction, our research will significantly tighten design and engineering for rapid prototyping.

8.3.1.2. PhySound

- Type: Training (ICT)
- Instrument: Marie-Curie Postdoctoral fellowship
- Partner: Columbia
- Abstract: Sound is as important as visuals in modern media (films, video-games). Yet, little effort has been devoted to the rendering of sound from digital environments, compared to the phenomenal advances of visual rendering. Sound is added to virtual scenes through the ad-hoc edition of real sounds, requiring recording phases and manual synchronization between recorded clips and visuals, while yielding limited and repetitive sounds. This project addresses this problem by generating sounds from virtual environments through physically based simulation, and focuses on a challenging family of objects: thin shells. Characteristic thin shell sounds include tearing cloth and paper, crushing cans and plastic bottles, and crumpling a piece of paper and a plastic bag. The high quality, offline simulation and rendering of thin shell sound will be addressed through a set of modeling approaches and computational tools (model reduction, high frequency bandwidth extension and pre-computed sound databases), while the real-time but computationally constrained sound rendering will rely on data-driven approaches. This research will considerably widen the number of real life object sounds that can be digitally generated, and will contribute to the young research field of physically based sound rendering, which has the potential of becoming the next key technology of the media industry.

8.3.1.3. EMOTIVE

Type: COOPERATION (ICT)

Instrument: Reseach Innovation Action

Objectif: Virtual Heritage

Duration: November 2016 - October 2019

Coordinator: EXUS SA (UK)

Partner: Diginext (FR), ATHENA (GR), Noho (IRL), U Glasgow (UK), U York (UK)

Inria contact: George Drettakis

Abstract: Storytelling applies to nearly everything we do. Everybody uses stories, from educators to marketers and from politicians to journalists to inform, persuade, entertain, motivate or inspire. In the cultural heritage sector, however, narrative tends to be used narrowly, as a method to communicate to the public the findings and research conducted by the domain experts of a cultural site or collection. The principal objective of the EMOTIVE project is to research, design, develop and evaluate methods and tools that can support the cultural and creative industries in creating Virtual Museums which draw on the power of 'emotive storytelling'. This means storytelling that can engage visitors, trigger their emotions, connect them to other people around the world, and enhance their understanding, imagination and, ultimately, their experience of cultural sites and content. EMOTIVE does this by providing the means to authors of cultural products to create high-quality, interactive, personalized digital stories. GRAPHDECO contributes by developing novel image-based rendering techniques to help museum curators and archeologists provide more engaging experiences, and in particular for the offsite experience for one of the sites (see Fig. 11).

8.4. International Initiatives

8.4.1. Inria International Partners

8.4.1.1. Informal International Partners

Canada. A. Bousseau collaborates regularly with the University of Toronto (K. Singh) and the University of British Columbia (A. Sheffer).

UK. G. Drettakis collaborates with UCL (G. Brostow, P. Hedman) and with R. Mantiuk (Cambridge).

United States. We regularly collaborate with Adobe Research (A. Hertzman, S. Paris). We also collaborate with Daniel Aliaga from Purdue University. We collaborate with M. Banks and A. Efros from University of California, Berkeley.

8.5. International Research Visitors

8.5.1. Visits of International Scientists

Several students and postdocs of F. Durand visited from MIT during 2017:

8.5.2. Visits to International Teams

8.5.2.1. Sabbatical programme

Fredo Durand was the recipient of the Inria International Chair and spent the academic year 2016-2017 in the group.

8.5.2.2. Research Stays Abroad

Johanna Delanoy spent 6 months at Adobe Research as an intern to collaborate with Aaron Hertzmann. S. Rodriguez and T. Thonat visited the MIT CSAIL Computer Graphic Lab, in Boston, USA. V. Deschaintre Visited Frederic Durand and Miika Aittala at MIT October/November. Y. Gryaditskaya visited the research group of Daniel Sykora, CTU Prague, Czech Republic, and the Industrial Design Faculty of TU Delft, Netherlands in June.

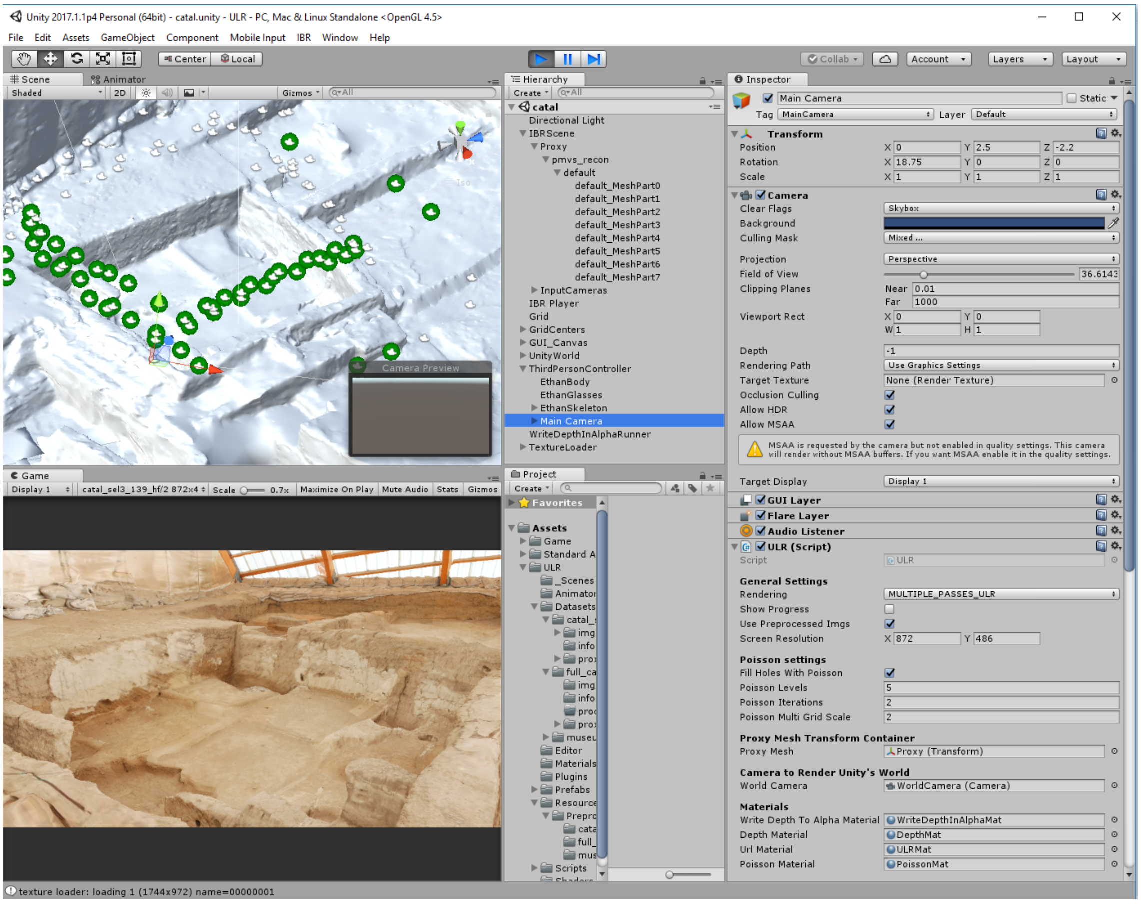

Figure 11. Screenshot of Unity IBR system developed for EMOTIVE.

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Scientific Events Organisation

9.1.1.1. Member of the Organizing Committees

G. Drettakis was a member of the ACM SIGGRAPH 2017 papers advisory board.

9.1.2. Scientific Events Selection

9.1.2.1. Member of the Conference Program Committees

G. Cirio was on the program committee of IEEE World Haptics 2017. A. Bousseau was a Papers Committee member of Eurographics 2018 and SIGGRAPH Asia 2017 technical briefs and posters, and a member of Eurographics STAR committee. G. Drettakis was a member of the ACM SIGGRAPH 2017 papers committee and of the Eurographics STAR committee.

9.1.2.2. Reviewer

A. Bousseau was a reviewer for ACM Transactions on Graphics, SIGGRAPH, SIGGRAPH Asia, Eurographics, ACM CHI Gabriel Cirio was reviewer for Eurographics, ACM SIGGRAPH, IEEE World Haptics, IEEE Transactions on Haptics, ACM UIST, IEEE 3DUI, IEEE VR. Y. Gryaditskaya was a reviewer for SIGGRAPH Asia 2017, Eurographics 2018, ACM TOG and JEI.

9.1.3. Journal

9.1.3.1. Member of the Editorial Boards

A. Bousseau is an Associate Editor of The Visual Computer journal (Springer) G. Drettakis is Associate Editor of ACM Transactions on Graphics (TOG) and of Computational Visual Media (CVM).

9.1.4. Invited Talks

Y. Gryaditskaya visted the Charles University of Prague in May 2017. Adrien Bousseau gave a talk about 3D reconstruction of design drawings at Delft University.

9.1.5. Leadership within the Scientific Community

G. Drettakis chairs the Eurographics working group on Rendering. G. Drettakis chairs the local "Jacques Morgenstern" Colloquium organizing committee.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

Licence : J. Delanoy, TP Bases de donnees and personal project, niveau L1, IUT - Université Nice Sophia Antipolis, 64h

Licence : S. Rodriguez, niveau L1, 64h as a DCCE in Polytech'Nice (UCA).

Licence : J. Philip, Introduction au Web & Application du Web, niveau L2, Polytech Nice-Sophia - Université Côte d'Azur, 64h

Licence: T. Thonat, Informatique Theorique, niveau L3, Polytech Nice, Sophia-Antipolis, 64h

Master : G. Drettakis, Computer Graphics, 12h, CentraleSupelec; 6h Université Nice Sophia Antipolis, 64h

Master : A. Bousseau, Computer Graphics, 9h, CentraleSupelec; 9h Université Nice Sophia Antipolis, 64h

9.2.2. Others

Y. Gryaditskaya and V. Deschaintre attended MAM and EGSR 2017 V. Deschaintre, S. Rodriguez, J. Philip and T. Thonat attended UCA Deep Learning School, 12-15 June 2017, Sophia-Antipolis, France. S. Rodriguez, J. Delanoy, T. Thonat and J. Philip attended Eurographics 2017, 24-28 April 2017, Lyon, France. Johanna Delanoy participated in the Doctoral Consortium at Eurographics 2017. Johanna Delanoy was a Student Volunteer at Siggraph 2017.

9.2.3. Presentations

Theo Thonat presented his 3DV-2016 multi view inpainting work at the AFIG French conference.

9.2.3.1. Supervision

PhD defended: Rodrigo Ortiz-Cayon, Mixed image-based rendering, Nice Sophia Antipolis University, started Dec. 2013, defended in Feb. 2017, advisor: G. Drettakis

PhD in progress: Théo Thonat, Multi-view image processing for image-based rendering, Nice Sophia Antipolis University, started Oct. 2015, advisor: G. Drettakis

PhD in progress: Johanna Delanoy, Data-driven sketch-based modeling, Nice Sophia Antipolis University, started Oct. 2015, advisor: A. Bousseau

PhD in progress: Jean-Dominique Favreau, geometric analysis of line drawings, Nice Sophia Antipolis University, started Oct. 2014, advisors: A. Bousseau and F. Lafarge (Titane)

PhD in progress: Valentin Deschaintre, lightweight material capture, Nice Sophia Antipolis University, started Nov. 2016, advisors: A. Bousseau and G. Drettakis

PhD in progress: Julien Philip, Mixed rendering for cultural heritage, Nice Sophia Antipolis University, started Nov. 2016, advisor: G. Drettakis

PhD in progress: Simon Rodriguez, Combining image-based and procedural modeling, Nice Sophia Antipolis University, started Nov. 2016, advisor: G. Drettakis

9.2.3.2. Juries

A. Bousseau was a reviewer for the PhD of Yulia Gryaditskaya (MPI Saarbruecken), Benoit Arbelot (Inria Rhône Alpes) and T. Stanko (Inria Rhône Alpes).

9.3. Popularization

Theo Thonat gave demos of Image Based Renderings during the 2017 "Fete de la science" at the Inria stand (Oct. 7.) for school children and adults audience.

10. Bibliography

Major publications by the team in recent years

[1] 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

[2] S. DUCHÊNE, C. RIANT, G. CHAURASIA, J. LOPEZ-MORENO, P.-Y. LAFFONT, S. POPOV, A. BOUSSEAU, G. DRETTAKIS. *Multi-View Intrinsic Images of Outdoors Scenes with an Application to Relighting*, in "ACM Transactions on Graphics", 2015, http://www-sop.inria.fr/reves/Basilic/2015/DRCLLPD15

[3] S. POPOV, R. RAMAMOORTHI, F. DURAND, G. DRETTAKIS. *Probabilistic Connections for Bidirectional Path Tracing*, in "Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)", 2015, vol. 34, n^{o} 4, http://www-sop.inria.fr/reves/Basilic/2015/PRDD15b

[4] 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 (SIGGRAPH Conference Proceedings)", 2014, vol. 33, n^{o} 4, http://www-sop.inria.fr/reves/Basilic/2014/XCSBMS14

Publications of the year

Doctoral Dissertations and Habilitation Theses

[5] R. ORTIZ CAYÓN. *Improving speed and image quality of image-based rendering*, Université Côte d'Azur, February 2017, https://tel.archives-ouvertes.fr/tel-01536072

Articles in International Peer-Reviewed Journals

[6] L. ANDERSON, T.-M. LI, J. LEHTINEN, F. DURAND. *Aether: An Embedded Domain Specific Sampling Language for Monte Carlo Rendering*, in "ACM Transactions on Graphics", July 2017, vol. 36, n^{o} 4, pp. 1 - 16 [*DOI :* 10.1145/3072959.3073704], https://hal.inria.fr/hal-01676191

[7] R. ARORA, I. DAROLIA, V. P. NAMBOODIRI, K. SINGH, A. BOUSSEAU. *SketchSoup: Exploratory Ideation Using Design Sketches*, in "Computer Graphics Forum", 2017, vol. 0, pp. 1 - 11 [*DOI :* 10.1111/CGF.13081], https://hal.inria.fr/hal-01470706

[8] J.-D. FAVREAU, F. LAFARGE, A. B. BOUSSEAU. *Photo2ClipArt: Image Abstraction and Vectorization Using Layered Linear Gradients*, in "ACM Transactions on Graphics", 2017, vol. 36, n^{o} 6, https://hal.inria.fr/hal-01581981

[9] A. FONDEVILLA, A. BOUSSEAU, D. ROHMER, S. HAHMANN, M.-P. CANI. *Patterns from Photograph: Reverse-Engineering Developable Products*, in "Computers and Graphics", August 2017, vol. 66, pp. 4-13 [*DOI :* 10.1016/J.CAG.2017.05.017], https://hal.inria.fr/hal-01525747

[10] M. GHARBI, J. CHEN, J. T. BARRON, S. W. HASINOFF, F. DURAND. *Deep Bilateral Learning for Real-Time Image Enhancement*, in "ACM Transactions on Graphics", July 2017, vol. 36, n^{o} 4, pp. 1 - 12 [*DOI :* 10.1145/3072959.3073592], https://hal.inria.fr/hal-01676188

[11] G.-A. KOULIERIS, B. BUI, M. S. BANKS, G. DRETTAKIS. *Accommodation and Comfort in Head-Mounted Displays*, in "ACM Transactions on Graphics", July 2017, vol. 36, n^{o} 4, 11 p. [*DOI :* 10.1145/3072959.3073622], https://hal.inria.fr/hal-01528458

International Conferences with Proceedings

[12] M. S. BANKS, P. JOHNSON, J. KIM, G. A. KOULIERIS, G. DRETTAKIS, J. A. Q. PARNELL, L. GORDON. *Focus-tunable and Fixed Lenses and Stereoscopic 3D Displays*, in "Proc. SPIE 10125, Emerging Liquid Crystal Technologies XII", San Fransisco, United States, April 2017, https://hal.inria.fr/hal-01585892

Other Publications

[13] A. FONDEVILLA, A. BOUSSEAU, D. ROHMER, S. HAHMANN, M.-P. CANI. *Towards developable products from a sketch*, April 2017, EUROGRAPHICS Poster Proceedings, Poster, https://hal.inria.fr/hal-01647123

[14] A. PATNEY, M. ZANNOLI, G.-A. KOULIERIS, J. KIM, G. WETZSTEIN, F. STEINICKE. *Applications of Visual Perception to Virtual Reality Rendering*, July 2017, pp. 1 - 38, SIGGRAPH 2017 [*DOI :* 10.1145/3084873.3086551], https://hal.inria.fr/hal-01585878