



Activity Report 2016

## **Project-Team GRAPHDECO**

GRAPHics and DEsign with hEterogeneous  
COntent

RESEARCH CENTER  
**Sophia Antipolis - Méditerranée**

THEME  
**Interaction and visualization**



## Table of contents

<b>1. Members</b>	<b>1</b>
<b>2. Overall Objectives</b>	<b>2</b>
<b>3. Research Program</b>	<b>3</b>
3.1.1. Computer-Assisted Design with Heterogeneous Representations	3
3.1.2. Graphics with Uncertainty and Heterogeneous Content	5
<b>4. Highlights of the Year</b>	<b>7</b>
<b>5. New Software and Platforms</b>	<b>7</b>
5.1. SWARPI-Unity	7
5.2. SIBR	7
5.3. MVIIR	7
5.4. SGTDBGP	8
<b>6. New Results</b>	<b>8</b>
6.1. Computer-Assisted Design with Heterogeneous Representations	8
6.1.1. How Novices Sketch and Prototype Hand-Fabricated Objects	8
6.1.2. Interactive Sketching of Urban Procedural Models	8
6.1.3. Fidelity vs. Simplicity: a Global Approach to Line Drawing Vectorization	9
6.1.4. SketchSoup: Exploratory Ideation using Design Sketches	10
6.1.5. Modeling Symmetric Developable Surfaces from a Single Photo	10
6.1.6. DeepSketch: Sketch-Based Modeling using Deep Volumetric Prediction	11
6.2. Graphics with Uncertainty and Heterogeneous Content	12
6.2.1. Cotemporal Multi-View Video Segmentation	12
6.2.2. Automatic 3D Car Model Alignment for Mixed Image-Based Rendering	13
6.2.3. Multi-View Inpainting for Image-Based Scene Editing and Rendering	13
6.2.4. Gaze Prediction using Machine Learning for Dynamic Stereo Manipulation in Games	14
6.2.5. A Feasibility Study with Image-Based Rendered Virtual Reality in Patients with Mild Cognitive Impairment	14
6.2.6. Scalable Inside-Out Image-Based Rendering	15
6.2.7. Measuring Accommodation and Comfort in Head-Mounted Displays	15
6.2.8. Beyond Gaussian Noise-Based Texture Synthesis	17
6.2.9. Fences in Image Based Rendering	17
6.2.10. Handling reflections in Image-Based Rendering	18
<b>7. Bilateral Contracts and Grants with Industry</b>	<b>18</b>
7.1. Bilateral Contracts with Industry	18
7.2. Bilateral Grants with Industry	18
<b>8. Partnerships and Cooperations</b>	<b>18</b>
8.1. National Initiatives	18
8.1.1.1. ANR DRAO	18
8.1.1.2. ANR SEMAPOLIS	19
8.2. European Initiatives	19
8.2.1.1. CR-PLAY – Capture Reconstruct Play	19
8.2.1.2. EMOTIVE	20
8.3. International Initiatives	20
8.3.1. Inria Associate Teams Not Involved in an Inria International Labs	20
8.3.2. Inria International Partners	20
8.4. International Research Visitors	21
8.4.1. Visits of International Scientists	21
8.4.2. Visits to International Teams	21
<b>9. Dissemination</b>	<b>21</b>
9.1. Promoting Scientific Activities	21

9.1.1. Scientific Events Organisation	21
9.1.2. Scientific Events Selection	21
9.1.2.1. Chair of Conference Program Committees	21
9.1.2.2. Member of Conference Program Committees	21
9.1.3. Journal	21
9.1.3.1. Member of the Editorial Boards	21
9.1.3.2. Reviewer - Reviewing Activities	21
9.1.4. Leadership within the Scientific Community	22
9.1.5. Scientific Expertise	22
9.2. Teaching - Supervision - Juries	22
9.2.1. Teaching	22
9.2.2. Supervision	22
9.2.3. Juries	22
9.3. Popularization	22
<b>10. Bibliography</b> .....	<b>23</b>

# Project-Team GRAPHDECO

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

## Keywords:

### Computer Science and Digital Science:

- 3.1.4. - Uncertain data
- 5. - Interaction, multimedia and robotics
- 5.1. - Human-Computer Interaction
- 5.1.1. - Engineering of interactive systems
- 5.1.2. - Evaluation of interactive systems
- 5.3.5. - Computational photography
- 5.4.4. - 3D and spatio-temporal reconstruction
- 5.5. - Computer graphics
- 5.5.1. - Geometrical modeling
- 5.5.2. - Rendering
- 5.5.3. - Computational photography
- 5.6. - Virtual reality, augmented reality
- 5.9.1. - Sampling, acquisition
- 6.3.5. - Uncertainty Quantification
- 7.5. - Geometry, Topology
- 8.2. - Machine learning
- 8.3. - Signal analysis

### Other Research Topics and Application Domains:

- 5. - Industry of the future
- 5.2. - Design and manufacturing
- 5.7. - 3D printing
- 8. - Smart Cities and Territories
- 8.3. - Urbanism and urban planning
- 9. - Society and Knowledge
- 9.1.2. - Serious games
- 9.2. - Art
- 9.2.2. - Cinema, Television
- 9.2.3. - Video games
- 9.5. - Humanities
- 9.5.6. - Archeology, History

## 1. Members

### Research Scientists

George Drettakis [Team leader, Inria, Research Scientist, HDR]

Adrien Bousseau [Inria, Research Scientist, HDR]

### Engineers

Georgios Kopanas [Inria, until Aug 2016]  
Sebastien Bonopera [Inria, from Sept 2015 to Sept 2016]

#### PhD Students

Johanna Delanoy [Inria]  
Valentin Deschaintre [Optis, from Nov 2016, CIFRE funding]  
Rodrigo Ortiz Cayon [Inria, until Nov 2016]  
Julien Philip [Inria, from Nov 2016, EU H2020 EMOTIVE funding]  
Simon Rodriguez [UCA (University), from Nov 2016]  
Theo Thonat [Inria]

#### Post-Doctoral Fellows

George Koulieris [Inria]  
Abdelaziz Djelouah [Inria]  
Kenneth Vanhoey [Inria, until Aug 2016]

#### Visiting Scientists

Frederic Durand [MIT, Visiting Scientist with an Inria Chair, from Sept 2016]  
Eugene Fiume [University of Toronto, from Jul 2016]  
Miika Aittala [MIT/Aalto University, from Nov 2016]

#### Administrative Assistant

Sophie Honnorat [Inria]

#### Others

Sai Praveen Bangaru [Inria, Intern, from May 2016 until Jul 2016]  
Srinivasa Rao Gadhamchetty [Inria, Intern, from May 2016 until Jul 2016]  
Amelie Fondevilla [Inria, Intern, from Feb 2016 to Jul 2016]

## 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.

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.

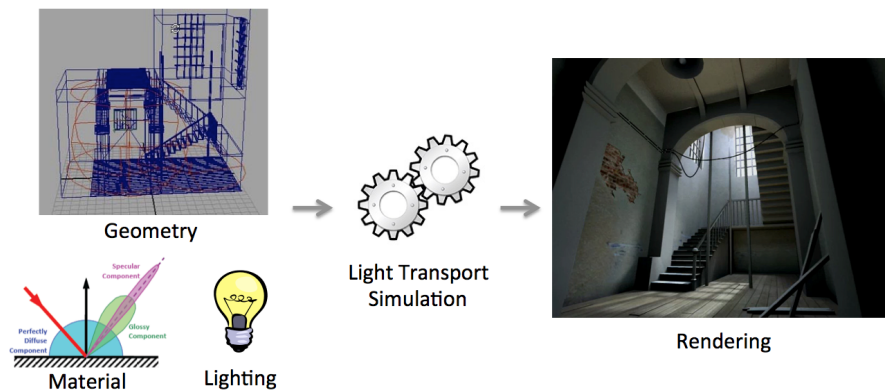


Figure 1. Traditional computer graphics pipeline. Rendering from [www.thegnomonworkshop.com](http://www.thegnomonworkshop.com)

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 heterogeneous 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.



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 addressing 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.



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

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.

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

In addition to publications in the leading conferences and journals in computer graphics (3 ACM Transactions on Graphics [5], [6], [8], 1 IEEE Virtual Reality), we made notable contributions to related fields such as human-computer interaction (1 ACM Conference on Human Factors in Computing Systems - CHI [9]) and computer vision (3 papers presented at the International Conference on 3D Vision [13], [12], [10]). Several of these results were developed in the context of the CR-Play project, which was completed in November with excellent reviews.

#### 4.1.1. Awards

Adrien Bousseau received a Young Researcher Award from the French National Research Agency (ANR) for the project ANR DRAO.

Adrien Bousseau obtained an ERC Starting Grant funding, the project will start in February 2017.

## 5. New Software and Platforms

### 5.1. SWARPI-Unity

SWARPI-Unity (for Superpixel Warp for Image-based rendering for Unity)

This is a software module developed in collaboration with Testaluna in the context of the CR-PLAY EU project. It involves an implementation of the Image-Based rendering algorithms of the group in the Unity3D framework. The software was improved this year to support mobile Android devices and was used in the evaluation step of the CR-PLAY project and for multiple demos.

- Participants: Sebastien Bonopera, Jerome Esnault, George Drettakis and Gaurav Chaurasia
- Contact: George Drettakis

### 5.2. SIBR

SIBR (for Simple Image-Based Rendering)

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. This new framework replaces the previously used IBR-COMMON tools.

- Participants: George Drettakis, Abdelaziz Djelouah, Rodrigo Ortiz Canyon, Theo Thonat, Sebastien Bonopera
- Contact: George Drettakis

### 5.3. MVIIR

MVIIR (for Multi-View Image Intrinsic Images and Relighting)

This package is the software implementation of the intrinsic image algorithm of Duchêne et al. It includes two libraries; one general-purpose that can be used to augment the functionalities of the previously mentioned SIBR framework, and another for specific logic concerning the relighting task. This package includes also programs to compute preprocess data required for the relighting of a dataset.

- Participants: George Drettakis, Sebastien Bonopera, Adrien Bousseau
- Contact: George Drettakis

## 5.4. SGTDP

SGTDP (for Synthetic Ground Truth Data Generation Platform)

We have started the development of a ground truth data generation platform based on complex and realistically rendered scenes built in 3D modelling packages such as 3DS Max. The platform includes an export module from 3DSMax with support for complex materials and shade trees such as those developed for the physically based rendering V-Ray platform. This module exports to the Mitsuba opensource renderer, and includes support for various operations using Mitsuba, as well as rendering on the Inria cluster. The platform is designed to generate ground truth data for learning as well as data for ground truth comparisons for image-based rendering projects in the group.

- Participants: George Drettakis, George Kopanas, Sai Bangaru
- Contact: George Drettakis

## 6. New Results

### 6.1. Computer-Assisted Design with Heterogeneous Representations

#### 6.1.1. How Novices Sketch and Prototype Hand-Fabricated Objects

**Participant:** Adrien Bousseau.

We are interested in how to create digital tools to support informal sketching and prototyping of physical objects by novices. Achieving this goal first requires a deeper understanding of how non-professional designers generate, explore, and communicate design ideas with traditional tools, i.e., sketches on paper and hands-on prototyping materials. We conducted a study framed around two all-day design charrettes where participants perform a complete design process: ideation sketching, concept development and presentation, fabrication planning documentation and collaborative fabrication of hand-crafted prototypes (Figure 4). This structure allows us to control key aspects of the design process while collecting rich data about creative tasks, including sketches on paper, physical models, and videos of collaboration discussions. Participants used a variety of drawing techniques to convey 3D concepts. They also extensively manipulated physical materials, such as paper, foam, and cardboard, both to support concept exploration and communication with design partners. Based on these observations, we propose design guidelines for CAD tools targeted at novice crafters.

This work is a collaboration with Theophanis Tsandilas, Lora Oehlberg and Wendy Mackay from the ExSitu group, Inria Saclay. It has been published at ACM Conference on Human Factors in Computing Systems (CHI) 2016 [9].




Figure 4. We asked participants to design a costume, from an initial sketch to a physical prototype.

#### 6.1.2. Interactive Sketching of Urban Procedural Models

**Participant:** Adrien Bousseau.

3D modeling remains a notoriously difficult task for novices despite significant research effort to provide intuitive and automated systems. We tackle this problem by combining the strengths of two popular domains: sketch-based modeling and procedural modeling. On the one hand, sketch-based modeling exploits our ability to draw but requires detailed, unambiguous drawings to achieve complex models. On the other hand, procedural modeling automates the creation of precise and detailed geometry but requires the tedious definition and parameterization of procedural models. Our system uses a collection of simple procedural grammars, called snippets, as building blocks to turn sketches into realistic 3D models. We use a machine learning approach to solve the inverse problem of finding the procedural model that best explains a user sketch. We use non-photorealistic rendering to generate artificial data for training convolutional neural networks capable of quickly recognizing the procedural rule intended by a sketch and estimating its parameters. We integrate our algorithm in a coarse-to-fine urban modeling system that allows users to create rich buildings by successively sketching the building mass, roof, facades, windows, and ornaments (Figure 5). A user study shows that by using our approach non-expert users can generate complex buildings in just a few minutes.

This work is a collaboration with Gen Nishida, Ignacio Garcia-Dorado, Daniel G. Aliaga and Bedrich Benes from Purdue University. It has been published at ACM Transactions on Graphics (proc. SIGGRAPH) 2016 [8].




Figure 5. Our system allows novices to quickly create complex procedural 3D models of buildings by sketching.

### 6.1.3. Fidelity vs. Simplicity: a Global Approach to Line Drawing Vectorization

**Participant:** Adrien Bousseau.

Vector drawing is a popular representation in graphic design because of the precision, compactness and editability offered by parametric curves. However, prior work on line drawing vectorization focused solely on faithfully capturing input bitmaps, and largely overlooked the problem of producing a compact and editable curve network. As a result, existing algorithms tend to produce overly-complex drawings composed of many short curves and control points, especially in the presence of thick or sketchy lines that yield spurious curves at junctions. We propose the first vectorization algorithm that explicitly balances fidelity to the input bitmap with simplicity of the output, as measured by the number of curves and their degree. By casting this trade-off as a global optimization, our algorithm generates few yet accurate curves, and also disambiguates curve topology at junctions by favoring the simplest interpretations overall. We demonstrate the robustness of our algorithm on a variety of drawings, sketchy cartoons and rough design sketches (Figure 6).

The first author of this work, Jean-Dominique Favreau, is co-advised by Adrien Bousseau and Florent Lafarge (Titane team). The work was published at ACM Transactions on Graphics (proc. SIGGRAPH) 2016 [5].

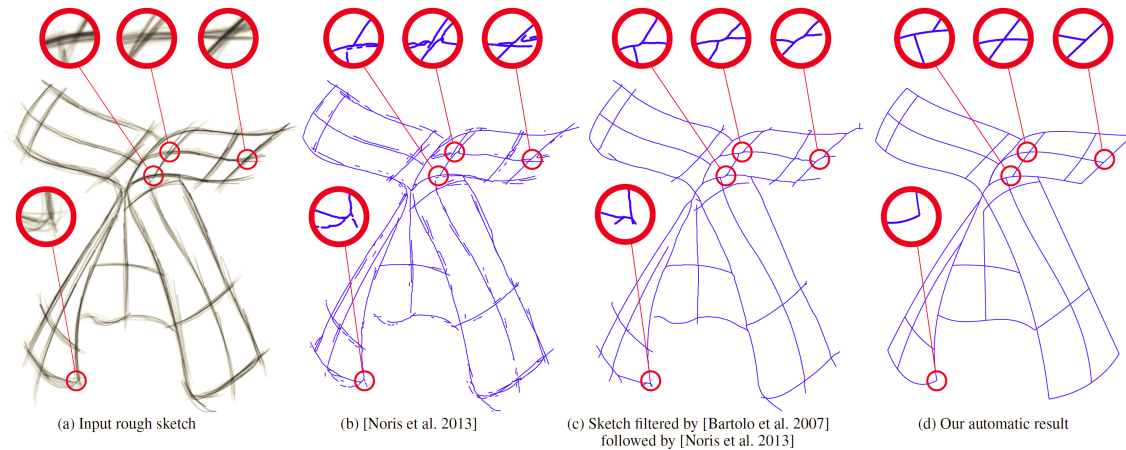


Figure 6. Rough sketches often contain overlapping strokes (a), which existing vectorization algorithms represent as multiple curves (b). Pre-filtering the drawing improves the vectorization, but produces spurious curve segments at junctions (c). Since existing algorithms analyze junctions locally, they cannot recover the proper topology of these seemingly similar line configurations. By adopting a global formulation that optimizes for both fidelity to the input sketch and simplicity of the output curve network, our algorithm recovers proper topology while significantly reducing the overall number of curves and control points. Design sketch after Sori Yanagi’s “Butterfly” stool.

#### 6.1.4. SketchSoup: Exploratory Ideation using Design Sketches

**Participant:** 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 present SketchSoup, a workflow that allows designers to explore the design space induced by such sketches. We take an unstructured collection of drawings as input, register them using a multi-image matching algorithm, and present them as a 2D interpolation space (Figure 7). 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.

This work is a collaboration with Rahul Arora and Karan Singh from University of Toronto and Vinay P. Namboodiri from IIT Kampur. The project was initiated while Rahul Arora was an intern in our group. It will be published in Computer Graphics Forum in 2017.

#### 6.1.5. Modeling Symmetric Developable Surfaces from a Single Photo

**Participants:** Amelie Fondevilla, Adrien Bousseau.




Figure 7. SketchSoup takes an unstructured set of sketches as input, along with a small number of correspondences (shown as red dots), registers the sketches and embeds them into a 2D interpolation space based on their shape differences. Users can explore the interpolation space to generate novel sketches.

We propose to reconstruct 3D developable surfaces from a single 2D drawing traced and annotated over a side-view photo of a partially symmetrical object (Figure 8). Our reconstruction algorithm combines symmetry and orthogonality shapes cues within a unified optimization framework that solves for the 3D position of the Bézier control points of the drawn curves while being tolerant to drawing inaccuracy and perspective distortions. We then rely on existing surface optimization methods to produce a developable surface that interpolates our 3D curves. Our method is particularly well suited for the modeling and fabrication of fashion items as it converts the input drawing into flattened developable patterns ready for sewing.

This work is a collaboration with Damien Rohmer, Stefanie Hahmann and Marie-Paule Cani from the Imagine team (LJK/ Inria Grenoble Rhône Alpes). This work was presented at the AFIG French conference in November 2016, where it received the 3rd price for best student work.




Figure 8. Our method reconstructs a 3D mesh and 2D pattern pieces of a sewed object from a single annotated drawing.

### 6.1.6. DeepSketch: Sketch-Based Modeling using 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.

## 6.2. Graphics with Uncertainty and Heterogeneous Content

### 6.2.1. Cotemporal Multi-View Video Segmentation

**Participants:** Abdelaziz Djelouah, George Drettakis.

We address the problem of multi-view video segmentation of dynamic scenes in general and outdoor environments with possibly moving cameras. Multi-view methods for dynamic scenes usually rely on geometric calibration to impose spatial shape constraints between viewpoints. In this work, we show that the calibration constraint can be relaxed while still getting competitive segmentation results using multi-view constraints. We introduce new multi-view cotemporality constraints through motion correlation cues, in addition to common appearance features used by cosegmentation methods to identify co-instances of objects. We also take advantage of learning based segmentation strategies by casting the problem as the selection of monocular proposals that satisfy multi-view constraints. This yields a fully automated method that can segment subjects of interest without any particular pre-processing stage (see Fig. 9). Results on several challenging outdoor datasets demonstrate the feasibility and robustness of our approach.

This work is a collaboration with Jean-Sébastien Franco and Edmond Boyer from Morpheo team at Inria Grenoble, and Patrick Pérez from Technicolor. The work has been published in the International Conference on 3D Vision (3DV) - 2016 [10].

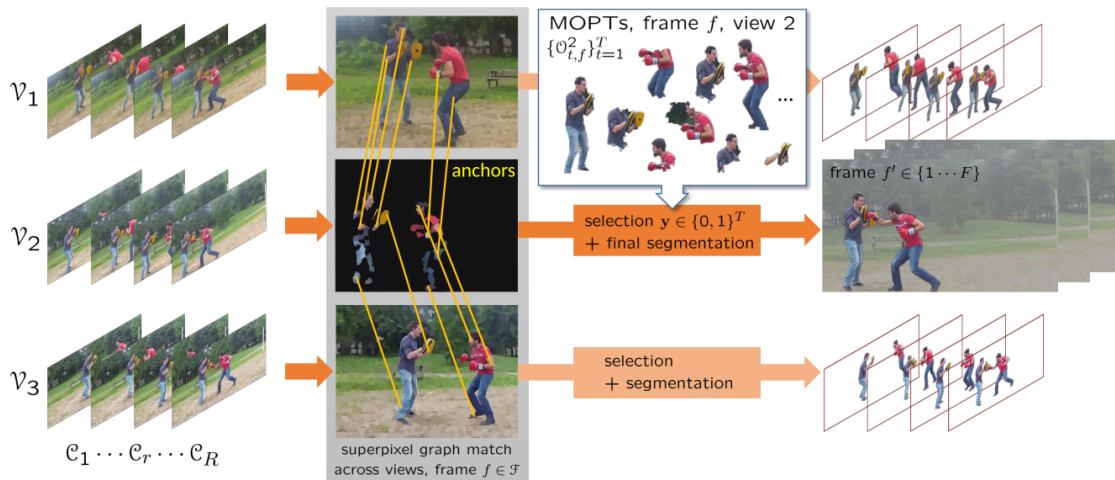


Figure 9. Synchronized videos of the same scene are partitioned into short clips. At a small number of instants where motion is sufficiently informative, cross-view correspondences between with similar appearance and motion are obtained by graph matching. In each view, matched superpixels, which are likely to lie on moving foreground objects, are used as sparse anchors to guide the selection process among a large pool of moving objects proposals extracted from all clips.



### 6.2.2. Automatic 3D Car Model Alignment for Mixed Image-Based Rendering

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

We present a method that improves quality of Image-Based Rendering of poorly reconstructed objects. We focus on the case of reflective objects which are hard to reconstruct, such as cars. The key insight is to replace these poorly reconstructed objects with models from existing rich 3D CAD databases, and subsequently align them to the input images. We use deep learning-based algorithms to obtain the 3D model present in the databases which is closest to the object seen in the images. We formulate two optimizations using all available information to finely position and orient the model and adapt it to image contours (see Fig. 10(1.)). Our method provides much higher quality rendering results of such objects compared to previous solutions as seen in Fig. 10(2.(b)).

This work is a collaboration with Francisco Massa and Mathieu Aubry from École des Ponts ParisTech. The work was published in the International Conference in 3D Vision (3DV) - 2016 [12].

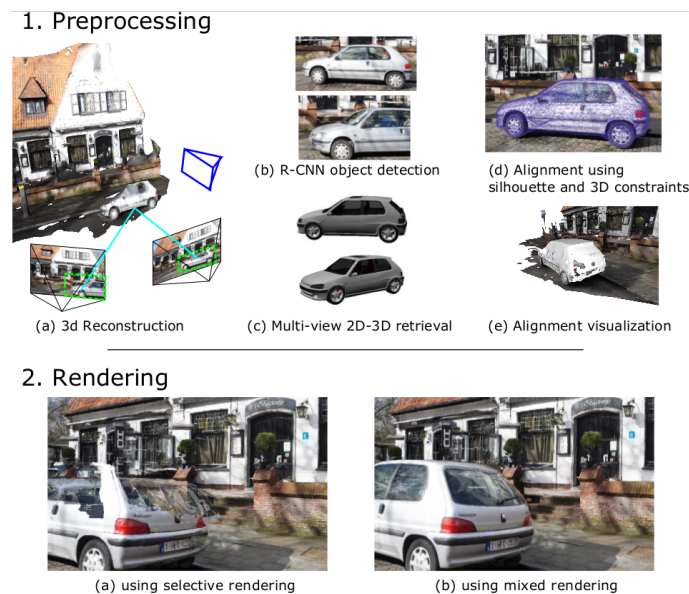


Figure 10. Overview of our pre-processing and rendering pipeline for Mixed Image-Based Rendering.

### 6.2.3. Multi-View Inpainting for Image-Based Scene Editing and Rendering

**Participants:** Theo Thonat, George Drettakis.

We propose a method to remove objects such as people and cars from multi-view urban image datasets (Figure 11), enabling free-viewpoint Image-Based Rendering (IBR) in the edited scenes. Our method combines information from multi-view 3D reconstruction with image inpainting techniques, by formulating the problem as an optimization of a global patch-based objective function. We use IBR techniques to reproject information from neighboring views, and 3D multi-view stereo reconstruction to perform multi-view coherent initialization for inpainting of pixels not filled by reprojection. Our algorithm performs multi-view consistent inpainting for color and 3D by blending reprojections with patch-based image inpainting. We run our algorithm on casually captured datasets, and Google Street View data, removing objects such as cars, people and pillars, showing that our approach produces results of sufficient quality for free-viewpoint IBR on “cleaned up” scenes, as well as IBR scene editing, such as limited displacement of real objects.

This work is a collaboration with Eli Shechtman and Sylvain Paris from Adobe Research. It has been published in the International Conference on 3D Vision (3DV) - 2016 [13].



Figure 11. Our system takes as input a set of images from the same scene (top row). The method then removes all the vehicules in a multi-view coherent way (bottom row).

#### 6.2.4. Gaze Prediction using Machine Learning for Dynamic Stereo Manipulation in Games

**Participants:** George Koulieris, George Drettakis.

Comfortable, high-quality 3D stereo viewing has become a requirement for interactive applications. Previous research shows that manipulating disparity can alleviate some of the discomfort caused by 3D stereo, but it is best to do this locally, around the object the user is gazing at. The main challenge is thus to develop a gaze predictor in the demanding context of real-time, heavily task-oriented applications such as games. Our key observation is that player actions are highly correlated with the present state of a game, encoded by game variables. Based on this, we trained a classifier to learn these correlations using an eye-tracker which provides the ground-truth object being looked at. The classifier is used at runtime to predict object category – and thus gaze – during game play, based on the current state of game variables. We used this prediction to propose a dynamic disparity manipulation method, which provided rich and comfortable depth. We evaluated the quality of our gaze predictor numerically and experimentally, showing that it predicts gaze more accurately than previous approaches. A subjective rating study demonstrates that our localized disparity manipulation is preferred over previous methods.

This work is a collaboration with Katerina Mania from the Technical University of Crete and Douglas Cunningham from the Technical University of Cottbus. The work was presented at the IEEE conference for Virtual Reality (IEEE VR) 2016 [14].

#### 6.2.5. A Feasibility Study with Image-Based Rendered Virtual Reality in Patients with Mild Cognitive Impairment

**Participant:** George Drettakis.




Figure 12. We propose a gaze predictor (a), used to perform localized stereo grading (b). We compare to no stereo grading (c) and prior work (d). We provide rich and comfortable depth. Please use red/cyan anaglyph glasses.

Virtual Reality (VR) has emerged as a promising tool in many domains of therapy and rehabilitation, and has recently attracted the attention of researchers and clinicians working with elderly people with MCI, Alzheimer’s disease and related disorders. In this study we tested the feasibility of using highly realistic image-based rendered VR with patients with MCI and dementia. We designed an attentional task to train selective and sustained attention, and we tested a VR and a paper version of this task (see Fig. 13) in a single-session within-subjects design. Results showed that participants with MCI and dementia reported to be highly satisfied and interested in the task, and they reported high feelings of security, low discomfort, anxiety and fatigue. In addition, participants reported a preference for the VR condition compared to the paper condition, even if the task was more difficult. Interestingly, apathetic participants showed a preference for the VR condition stronger than that of non-apathetic participants. These findings suggest that VR-based training can be considered as an interesting tool to improve adherence to cognitive training for elderly people with cognitive impairment.

This work was a collaboration with EA CoBTek/IA, CMRR (memory center) of the CHU (University Hospital) of Nice, Disney Research and Trinity College Dublin, as part of the (completed) VERVE EU project. The work was published in the PLoS ONE journal [7].

### 6.2.6. Scalable Inside-Out Image-Based Rendering

**Participant:** George Drettakis.

The goal of this project was to provide high-quality free-viewpoint rendering of indoors environments, captured with off-the-shelf equipment such as a high-quality color camera and a commodity depth sensor. Image-based Rendering (IBR) can provide the realistic imagery required at real-time speed. For indoor scenes however, two challenges are especially prominent. First, the reconstructed 3D geometry must be compact, but faithful enough to respect occlusion relationships when viewed up close. Second, man-made materials call for view-dependent texturing, but using too many input photographs reduces performance.

We customize a typical RGB-D 3D surface reconstruction pipeline to produce a coarse global 3D surface, and local, per-view geometry for each input image. Our tiled IBR preserves quality by economizing on the expected contributions that entire groups of input pixels make to a final image. The two components are designed to work together, giving real-time performance, while hardly sacrificing quality. Testing on a variety of challenging scenes shows that our inside-out IBR scales favorably with the number of input images.

This work was a collaboration with P. Hedman, G. Brostow and T. Ritschel at UCL, as part of the CR-PLAY project. It was published in ACM Transactions on Graphics (Proc. SIGGRAPH Asia) [6].

### 6.2.7. Measuring Accommodation and Comfort in Head-Mounted Displays

**Participants:** George Koulieris, George Drettakis.



Figure 13. The VR and paper conditions of the study



Figure 14. Images from our method rendered in 1080p at 55 Hz on an Nvidia Titan X GPU. Input is an RGB-D video and 298 high-quality photos of 'Dr Johnson's house', London. With no wheelchair access to this floor, curators were keen to have their rooms digitized.

Head-mounted displays (HMDs) are rapidly becoming the preferred display for stereo viewing in virtual environments, but they often cause discomfort and even sickness. Previous studies have shown that a major cause of these adverse symptoms is the vergence-accommodation (VA) conflict. Specifically, the eyes use the distance to the screen to accommodate, while they use the distance to the fixated virtual object to converge. The VA conflict is the difference between those distances. The magnitude of the conflict is well correlated with subjective reports of discomfort. Many methods have been proposed for reducing the VA conflict and thereby reducing discomfort by making accommodation more consistent with vergence. But no one has actually measured accommodation in HMDs to see how well a given method is able to drive it to the desired distance. We built the first device for measuring accommodation in an HMD, using a modular design with off-the-shelf components, focus-adjustable lenses, and an autorefractor. We conducted experiments using the device to determine how well accommodation is driven with various combinations of HMD properties and viewing conditions: focus-adjustable lenses, depth-of-field rendering, binocular viewing, and “monovision” (setting the two eyes’ focal distances to quite different values). We found that focus-adjustable lenses drive accommodation appropriately across many conditions. The other techniques were much less effective in driving accommodation. We also investigated whether the ability to drive accommodation predicts viewer comfort. We did this by conducting a discomfort study with most of the conditions in the accommodation study. We found that the ability to drive accommodation did in fact predict the amount of discomfort. Specifically, the most comfortable conditions were the ones that generated accommodation consistent with vergence. Together, these results illustrate the potential benefit of focus-adjustable lenses: They enable stimulation of accommodation and thereby comfortable viewing. In contrast, monovision neither enable accurate accommodation nor comfortable viewing.

This work is an ongoing collaboration with Martin S. Banks from UC Berkeley, in the context of the CRISP Inria associate team.

#### **6.2.8. *Beyond Gaussian Noise-Based Texture Synthesis***

**Participants:** Kenneth Vanhoey, Georgios Kopanas, George Drettakis.

Texture synthesis methods based on noise functions have many nice properties: they are continuous (thus resolution-independent), infinite (can be evaluated at any point) and compact (only functional parameters need to be stored). A good method is also non-repetitive and aperiodic. Current techniques, like Gabor Noise, fail to produce structured content. They are limited to so-called “Gaussian textures”, characterized by second-order statistics like mean and variance only. 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). Other techniques, like Local Random-Phase noise, leverage some structure but as a trade-off with repetitiveness and periodicity.

In this project, we model higher-order statistics produced by noise functions in a parametric model. Then we define an algorithm for sampling of the noise functions’ parameters so as to produce a texture that meets prescribed statistics. This sampling ensures both the reproduction of higher-order visual features with high probability, like edges and ridges, and non-repetitiveness plus aperiodicity thanks to the stochastic sampling method. Moreover a (deep) learning algorithm has been established to infer the prescribed statistics from an input exemplar image.

This project is a collaboration with Ian H. Jermyn (Durham University, UK, former Inria) and Mathieu Aubry (ENPC, France).

#### **6.2.9. *Fences in Image Based Rendering***

**Participants:** 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. Some methods try to overcome the issue in the case of reflections and transparencies through the estimation of two planes or the usage of 3D stock models. Fences with their thin repetitive structures are another important common source of 3D reconstruction errors but have received very little attention in the context of image based rendering. 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 pictures for IBR applications. Similarly to related methods in image *de-fencing*, we use the assumptions that fences are thin repetitive structures lying on a 3D plane. To address this problem we consider the following steps: First we propose a multi-view approach to estimate the plane supporting the fences using repetition candidates. Second, we estimate image matting taking advantage of the multi-view constraints and the repetitive patterns. Finally, the estimated 3D plane and matting masks are used in a new rendering algorithm.

#### 6.2.10. Handling reflections in Image-Based Rendering

**Participants:** Theo Thonat, Frederic Durand, George Drettakis.

In order to render new viewpoints, current Image Based Rendering techniques (IBR) use an approximate geometry to warp and blend photographs from close viewpoints. They assume the scene materials are diffuse, so geometry colors are independent of the viewpoint, an assumption that fails 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 their associated geometry is not available and since the normals of the reflective surfaces might be not reliable.

## 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 is starting a CIFRE PhD in collaboration with Optis, a company specialized in material acquisition and rendering.

### 7.2. Bilateral Grants with Industry

#### 7.2.1. 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 will involve the use of our synthetic ground truth data generation platform (see Sec. 5.4) to tasks such as color grading and white balance. This is a collaboration with P. Pérez and E. Reinhard of Technicolor.

## 8. Partnerships and Cooperations

### 8.1. National Initiatives

#### 8.1.1. ANR

##### 8.1.1.1. ANR DRAO

**Participant:** 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.

While the ANR DRAO ended in 2015, it has resulted in a publication at ACM Conference on Human Factors in Computing Systems (CHI) 2016 [9]. Adrien Bousseau received an ANR Young Researcher Award for coordinating this project.

#### 8.1.1.2. ANR SEMAPOLIS

**Participants:** George Drettakis, Abdelaziz Djelouah, Theo Thonat.

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 of urban models, in particular using image-based rendering algorithms.

This year, the ANR SEMAPOLIS resulted in five publications on multi-view segmentation [10], multi-view inpainting [13], image-based rendering of cars [12] and interiors [6], and procedural modeling of buildings [8]. Two of these projects rely on a deep learning method from the ENPC group to identify semantic object categories in images (e.g., cars, people etc.) [13], [12]. In [12] we also collaborated with the ENPC group to use a deep learning method to allow the use of rendered images to identify objects in photographs. These collaborations have been extremely fruitful for our group, and have opened the way to several new collaborations with ENPC.

## 8.2. European Initiatives

### 8.2.1. FP7 & H2020 Projects

#### 8.2.1.1. CR-PLAY – Capture Reconstruct Play

<http://www.cr-play.eu>

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. 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.

This year we had four results related to CR-PLAY on multi-view segmentation [10], multi-view inpainting [13], image-based rendering of cars [12] and indoors [6]. The work on indoors rendering was led by the CR-PLAY partner UCL. The CR-PLAY project ended in October, and was successfully evaluated in November. The project has resulted in a number of technological developments related to the Ph.D. work of R. Ortiz-Cayon and T. Thonat, as well as the postdoc of A. Djelouah which will be the object of a market study in the goal of a technology transfer.

### 8.2.1.2. EMOTIVE

Type: COOPERATION (ICT)

Instrument: Research 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 will do this by providing the means to authors of cultural products to create high-quality, interactive, personalized digital stories.

GRAPHDECO will contribute by developing novel image-based rendering techniques to help museum curators and archeologists provide more engaging experiences.

## 8.3. International Initiatives

### 8.3.1. Inria Associate Teams Not Involved in an Inria International Labs

#### 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 (United States) - Electrical Engineering and Computer Science Department (EECS) - Maneesh Agrawala

Start year: 2014

See also: <http://www-sop.inria.fr/revs/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. In 2016, the CRISP collaboration supported the postdoc of George Koulouris, who spent 6 months at UC Berkeley and is now at Inria. Johanna Delanoy also spent 2 weeks at UC Berkeley to collaborate with Alexei Efros.

### 8.3.2. Inria International Partners

#### 8.3.2.1. Informal International Partners

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

**India.** A. Bousseau collaborates with Vinay Namboodiri from IIT Kanpur and hosted several of his students for internships (Rahul Arora and Srinivasa Rao Gadhamchetty).



**UK.** G. Drettakis collaborates with UCL in the context of the CR-PLAY projects, resulting in a publication on indoor image-based rendering [6].

**United States.** We regularly collaborate with Adobe Research. Most recently, G. Drettakis worked with Eli Shechtman and Sylvain Paris on multi-view inpainting [13]. We also collaborate with Daniel Aliaga from Purdue University on sketch-based procedural modeling [8].

## 8.4. International Research Visitors

### 8.4.1. Visits of International Scientists

Martin Banks (UC Berkeley) visited our group for two weeks in the context of the Associate Team CRISP (Sec. 8.3.1.1). We also hosted Philip Isola, Richard Zhang and Alexei Efros (UC Berkeley) for two days, also in the context of CRISP. Several international researchers made short visits and talks: Elena Garces (University of Zaragoza), Yulia Gryditskaya (MPI), Jan Jermyn (Durham University), Christian Lessig (Otto-von-Guericke Universitat Magdeburg), Marc Stamminger (Erlangen University). Finally, we hosted Frédo Durand from MIT (10 months) and Eugene Fiume from university of Toronto (6 months) for their sabbatical.

#### 8.4.1.1. Internships

Sai Praveen Bangaru and Srinivasa Rao Gadhamchetty were master interns from IIT Madras and IIT Kampur respectively.

### 8.4.2. Visits to International Teams

#### 8.4.2.1. Research Stays Abroad

George Koulieris spent 6 months at UC Berkeley (Feb. 1 - Jul. 31) to work with Martin S. Banks in the context of the CRISP Inria associate team. Johanna Delanoy also visited UC Berkeley for two weeks to work with Alexei Efros.

## 9. Dissemination

### 9.1. Promoting Scientific Activities

#### 9.1.1. Scientific Events Organisation

##### 9.1.1.1. General Chair, Scientific Chair

George Drettakis chairs the Eurographics working group on Rendering.

#### 9.1.2. Scientific Events Selection

##### 9.1.2.1. Chair of Conference Program Committees

Adrien Bousseau is co-chair of the Tutorial program of Eurographics 2017.

##### 9.1.2.2. Member of Conference Program Committees

Adrien Bousseau was on the program committee of Eurographics 2017 and of SIGGRAPH Asia 2016 technical briefs and posters.

#### 9.1.3. Journal

##### 9.1.3.1. Member of the Editorial Boards

Adrien Bousseau is Associate Editor of The Visual Computer journal (Springer)

George Drettakis is Associate Editor of ACM Transactions on Graphics (TOG) and of Computational Visual Media (CVM).

##### 9.1.3.2. Reviewer - Reviewing Activities

Adrien Bousseau was reviewer for ACM TOG, IEEE TVCG, SIGGRAPH, SIGGRAPH Asia, Computers & Graphics, ACM CHI, ACM UIST, AFIG.

George Drettakis was reviewer for SIGGRAPH, Computer Graphics Forum.

#### **9.1.4. Leadership within the Scientific Community**

George Drettakis chairs the local “Jacques Morgenstern” Colloquium organizing committee.

#### **9.1.5. Scientific Expertise**

G. Drettakis is a member of the Scientific Advisory Board (SAB) of Aalto University, Finland, and participated in the SAB meeting at Aalto in October.

## **9.2. Teaching - Supervision - Juries**

### **9.2.1. Teaching**

Licence : Abdelaziz Djelouah, TP système d’information, niveau L1, Université Nice Sophia Antipolis, 24h

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

Masters II: G. Drettakis, A. Bousseau: Foundations of Image Synthesis, CentralSupélec, Paris, 28h

Masters I: A. Bousseau, G. Drettakis: Image Synthesis International Masters I, University of Nice Sophia-Antipolis, 12h

Masters I: G. Drettakis, A. Bousseau: Image Synthesis, Masters MAPI Cannes, University of Nice Sophia-Antipolis, 12h

### **9.2.2. Supervision**

PhD in progress : Rodrigo Ortiz-Cayon, Mixed image-based rendering, Nice Sophia Antipolis University, started Dec. 2013, to be 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, advisor: A. Bousseau and F. Lafarge (Titane)

PhD in progress : Valentin Deschaintre, lightweight material capture, Nice Sophia Antipolis University, started Nov. 2016, advisor: 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. Juries**

Adrien Bousseau was in the PhD committees of Boris Raymond (Inria Bordeaux) and Elena Garces (University of Zaragoza)

George Drettakis was in the PhD committees Oriel Frigo (Paris Descartes) and was an evaluator (rapporteur) for Kaan Yucer (ETH Zurich) and Yitzchak Lockerman (Yale).

## **9.3. Popularization**

George Koulieris described his research in an interview for the Inria@SiliconValley Newsletter <sup>1</sup>.

<sup>1</sup><https://project.inria.fr/siliconvalley/2016/06/21/post-doc-george-koulieris/>

Adrien Bousseau described his research in two interviews for the Inria website <sup>2</sup> <sup>3</sup>.

Adrien Bousseau gave demos of his tool to help people practice drawing-by-observation techniques during the “Rencontres du numérique” organized by ANR (Nov. 16th) and during the “Rencontres Inria-Industrie” on e-learning (Dec. 1st).

## 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/revs/Basilic/2013/CDS13>
- [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/revs/Basilic/2015/DRCLLP15>
- [3] S. POPOV, R. RAMAMOORTHY, 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<sup>o</sup> 4, <http://www-sop.inria.fr/revs/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<sup>o</sup> 4, <http://www-sop.inria.fr/revs/Basilic/2014/XCSBMS14>

### Publications of the year

#### Articles in International Peer-Reviewed Journals

- [5] J.-D. FAVREAU, F. LAFARGE, A. BOUSSEAU. *Fidelity vs. Simplicity: a Global Approach to Line Drawing Vectorization*, in "ACM Transactions on Graphics", 2016 [DOI : 10.1145/2897824.2925946], <https://hal.inria.fr/hal-01309271>
- [6] P. HEDMAN, T. RITSCHER, G. DRETTAKIS, G. BROSTOW. *Scalable Inside-Out Image-Based Rendering*, in "ACM Transactions on Graphics", December 2016, vol. 35, n<sup>o</sup> 6, <https://hal.inria.fr/hal-01369382>
- [7] V. MANERA, E. CHAPOULIE, J. BOURGEOIS, R. GUERCHOUCHE, R. DAVID, J. ONDREJ, G. DRETTAKIS, P. ROBERT. *A Feasibility Study with Image-Based Rendered Virtual Reality in Patients with Mild Cognitive Impairment and Dementia*, in "PLoS ONE", March 2016, vol. 11, n<sup>o</sup> 3, 14 p. [DOI : 10.1371/JOURNAL.PONE.0151487], <https://hal.inria.fr/hal-01292254>
- [8] G. G. NISHIDA, I. G. GARCIA-DORADO, D. G. ALIAGA, B. BENES, A. BOUSSEAU. *Interactive sketching of urban procedural models*, in "ACM Transactions on Graphics", 2016, vol. 35, n<sup>o</sup> 4, pp. 1 - 11 [DOI : 10.1145/2897824.2925951], <https://hal.inria.fr/hal-01378388>

<sup>2</sup><https://www.inria.fr/en/centre/sophia/news/adrien-bousseau-receives-an-erc-starting-grant>

<sup>3</sup><https://www.inria.fr/centre/sophia/actualites/adrien-bousseau-laureat-du-prix-jeune-chercheur>

### International Conferences with Proceedings

- [9] A. BOUSSEAU, T. TSANDILAS, L. OEHLBERG, W. E. MACKAY. *How Novices Sketch and Prototype Hand-Fabricated Objects*, in "Conference on Human Factors in Computing Systems (CHI)", San Jose, United States, ACM, May 2016 [DOI : 10.1145/2858036.2858159], <https://hal.inria.fr/hal-01272187>
- [10] A. DJELOUAH, J.-S. FRANCO, E. BOYER, P. PÉREZ, G. DRETTAKIS. *Cotemporal Multi-View Video Segmentation*, in "International Conference on 3D Vision", Stanford, United States, October 2016, <https://hal.inria.fr/hal-01367430>
- [11] F. NEGIN, J. BOURGEOIS, E. CHAPOULIE, P. ROBERT, F. BREMOND. *Praxis and Gesture Recognition*, in "The 10th World Conference of Gerontechnology (ISG 2016)", Nice, France, September 2016, <https://hal.inria.fr/hal-01416372>
- [12] R. ORTIZ-CAYON, A. DJELOUAH, F. MASSA, M. AUBRY, G. DRETTAKIS. *Automatic 3D Car Model Alignment for Mixed Image-Based Rendering*, in "2016 International Conference on 3D Vision (3DV)", Stanford, United States, October 2016, <https://hal.inria.fr/hal-01368355>
- [13] T. THONAT, E. SHECHTMAN, S. PARIS, G. DRETTAKIS. *Multi-View Inpainting for Image-Based Scene Editing and Rendering*, in "2016 International Conference on 3D Vision (3DV)", Stanford, United States, October 2016, <https://hal.inria.fr/hal-01369554>

### Conferences without Proceedings

- [14] G. A. KOULIERIS, G. DRETTAKIS, D. CUNNINGHAM, K. MANIA. *Gaze Prediction using Machine Learning for Dynamic Stereo Manipulation in Games*, in "IEEE Virtual Reality", Greenville, United States, March 2016, <https://hal.inria.fr/hal-01290273>