Activity Report 2018

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
A3.1.10. - Heterogeneous data
A3.4.1. - Supervised learning
A3.4.6. - Neural networks
A3.4.8. - Deep learning
A5.1. - Human-Computer Interaction
A5.1.1. - Engineering of interactive systems
A5.1.2. - Evaluation of interactive systems
A5.1.8. - 3D User Interfaces
A5.1.9. - User and perceptual studies
A5.3.5. - Computational photography
A5.4.4. - 3D and spatio-temporal reconstruction
A5.4.5. - Object tracking and motion analysis
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
A5.9.3. - Reconstruction, enhancement
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
B9.6. - Humanities
B9.6.6. - Archeology, History

1. Team, Visitors, External Collaborators

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

Post-Doctoral Fellows
Gabriel Cirio [Inria, until Feb 2018]
Rada Deeb [Inria, from Dec 2018]
Tibor Stanko [Inria, from Feb 2018]
Abdelaziz Djelouah [Disney Research, until Jan 2018]
Georgios Koulieris [Durham University, until Feb 2018]

PhD Students
Johanna Delanoy [Inria]
Valentin Deschaintre [Optis]
Jean-Dominique Favreau [Inria, until Mar 2018]
David Jourdan [Inria, from Oct 2018]
Julien Philip [Inria]
Simon Rodriguez [Ministère de l'Enseignement Supérieur et de la Recherche]
Theo Thonat [Inria]

Technical staff
Yulia Gryaditskaya [Inria]
Sebastian Vizcay [Inria]
Bastien Wailly [Inria]
Georgios Kopanas [Inria, Research Engineer, until Feb 2018]
Sebastien Morgenthaler [Inria, from Oct 2018]

Interns
Stavros Diolatzis [Inria, from Oct 2018]
Alex Auvolat Bernstein [Ecole Normale Supérieure Paris, from Feb 2018 until Jul 2018]
David Jourdan [Inria, from Mar 2018 until Aug 2018]
Adele Saint Denis [Inria, from May 2018 until Jul 2018]

Administrative Assistant
Sophie Honnorat [Inria]

Visiting Scientists
Miika Aittala [MIT, from Jun 2018 until Jul 2018]
Yotam Gingold [George Mason University, from Jun 2018 until Jul 2018]

External Collaborator
Frédéric Durand [MIT]

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.

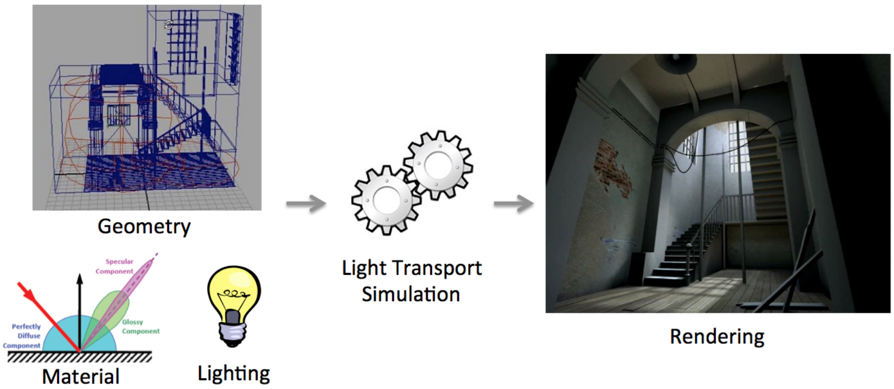

Figure 1. Traditional computer graphics pipeline. Rendering from www.thegnomonworkshop.com

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.

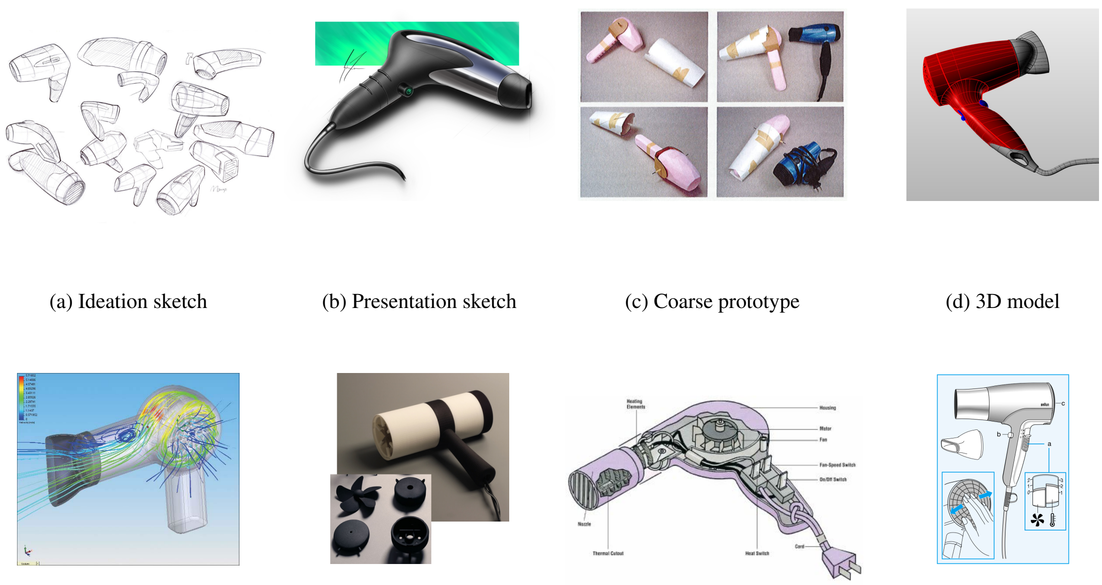

(a) Ideation sketch (b) Presentation sketch (c) Coarse prototype (d) 3D model

(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.

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

This year marked the start of the ERC Starting grant FunGraph coordinated by George Drettakis, on managing uncertainty in rendering of captured content. This activity already includes the principal investigator, one engineer (S. Morgenthaler), one postdoc (R. Deeb), and an intern (S. Diolatzis). The scientific production this year included three papers in ACM Transactions on Graphics (two at SIGGRAPH and one at SIGGRAPH Asia), three papers in Computer Graphics Forum (two at EGSR and one at Eurographics), and two papers at the ACM Symposium on Interactive 3D Graphics and Games.

4.1.1. Awards

George Drettakis received a medal from University Côte d'Azur for his ERC grant.

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.

We have recently integrated the platform with the SIBR software library, allowing us to read mitsuba scenes. We have written a tool to allow camera placement to be used for rendering and for reconstruction of synthetic scenes, including alignment of the exact and reconstructed version of the scenes. This dual-representation scenes can be used for learning and as ground truth. We can also perform various operations on the ground truth data within SIBR, e.g., compute shadow maps of both exact and reconstructed representations etc.

- 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. In addition, in order to develop a demo showcasing the technology, a multiplayer VR scene was created proving the integration of IBR with the rest of the 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

5.4. SynthDraw

KEYWORDS: Non-photorealistic rendering - Vector-based drawing
FUNCTIONAL DESCRIPTION: The SynthDraw library extracts occluding contours and sharp features over a 3D shape, computes all their intersections using a binary space partitionning algorithm, and finally makes a raycast to determine each sub-contour visibility. The resulting lines can then be exported as an SVG file for subsequent processing, for instance to stylize the drawing with different brush strokes. The library can also export various attributes for each line, such as its visibility and type. SynthDraw is based on the geometry processing library libIGL.

RELEASE FUNCTIONAL DESCRIPTION: This first version extracts occluding contours and creases, and computes their visibility with brute-force ray casting.

- Contact: Bastien Wailly

5.5. DeepSketch

KEYWORDS: 3D modeling - Sketching - Deep learning
FUNCTIONAL DESCRIPTION: DeepSketch is a sketch-based modeling system that runs in a web browser. It relies on deep learning to recognize geometric shapes in line drawings. The system follows a client/server architecture, based on the Node.js and WebGL technology. The application's main targets are iPads or Android tablets equipped with a digital pen, but it can also be used on desktop computers.

RELEASE FUNCTIONAL DESCRIPTION: This first version is built around a client/server Node.js application whose job is to transmit a drawing from the client's interface to the server where the deep networks are deployed, then transmit the results back to the client where the final shape is created and rendered in a WebGL 3D scene thanks to the THREE.js JavaScript framework. Moreover, the client is able to perform various camera transformations before drawing an object (change position, rotate in place, scale on place) by interacting with the touch screen. The user also has the ability to draw the shape's shadow to disambiguate depth/height. The deep networks are created, trained and deployed with the Caffe framework.

- Contact: Adrien Bousseau

6. New Results

6.1. Computer-Assisted Design with Heterogeneous Representations

6.1.1. 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 are 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 (Figure 4). 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. The work was published in Proceedings of the ACM on Computer Graphics and Interactive Techniques and presented at the ACM SIGGRAPH I3D Symposium on Interactive Computer Graphics and Games [12].

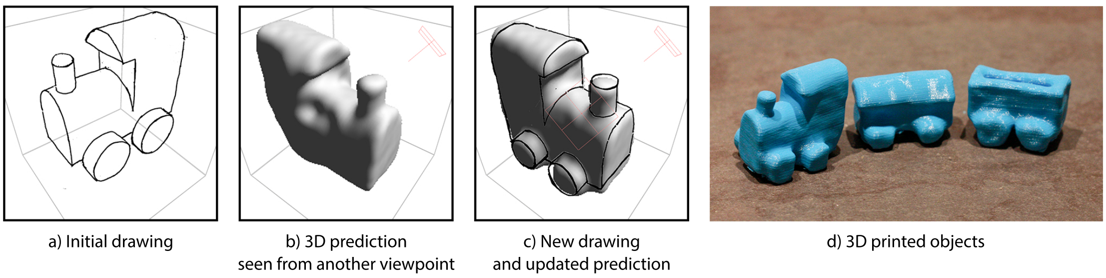

Figure 4. Our sketch-based modeling system can process as little as a single perspective drawing (a) to predict a volumetric object (b). Users can refine this prediction and complete it with novel parts by providing additional drawings from other viewpoints (c). This iterative sketching workflow allows quick 3D concept exploration and rapid prototyping (d).

6.1.2. Procedural Modeling of a Building from a Single Image

Participant: Adrien Bousseau.

Creating a virtual city is demanded for computer games, movies, and urban planning, but it takes a lot of time to create numerous 3D building models. Procedural modeling has become popular in recent years to overcome this issue, but creating a grammar to get a desired output is difficult and time consuming even for expert users. In this paper, we present an interactive tool that allows users to automatically generate such a grammar from a single image of a building. The user selects a photograph and highlights the silhouette of the target building as input to our method. Our pipeline automatically generates the building components, from large-scale building mass to fine-scale windows and doors geometry. Each stage of our pipeline combines convolutional neural networks (CNNs) and optimization to select and parameterize procedural grammars that reproduce the building elements of the picture. In the first stage, our method jointly estimates camera parameters and building mass shape. Once known, the building mass enables the rectification of the facades, which are given as input to the second stage that recovers the facade layout. This layout allows us to extract individual windows and doors that are subsequently fed to the last stage of the pipeline that selects procedural grammars for windows and doors. Finally, the grammars are combined to generate a complete procedural building as output. We devise a common methodology to make each stage of this pipeline tractable. This methodology consists in simplifying the input image to match the visual appearance of synthetic training data, and in using optimization to refine the parameters estimated by CNNs. We used our method to generate a variety of procedural models of buildings from existing photographs.

The work was published in Computer Graphics Forum, presented at Eurographics 2018 [15].

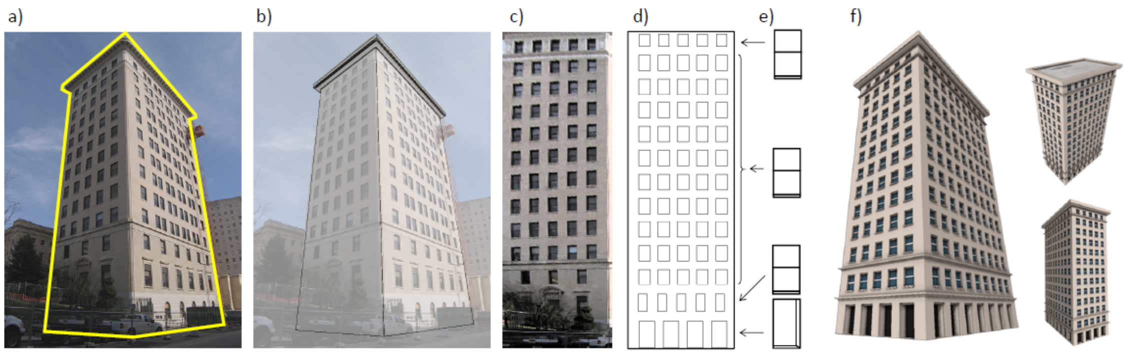

Figure 5. (a) Given an image and a silhouette of a building, (b) our approach automatically estimates the camera parameters and generates a building mass grammar as a first step. Then, (c) the facade image is rectified, and (d) the facade grammar is generated. (e) For each window non-terminal, the best window grammar is selected by maximum vote. (f) Finally the output grammar is constructed and a corresponding 3D geometry is generated.

6.1.3. *OpenSketch: A Richly-Annotated Dataset of Product Design Sketches*

Participants: Yulia Gryaditskaya, Frédéric Durand, Adrien Bousseau.

We collected a dataset of more than 400 product design sketches, representing 12 man-made objects drawn from two different view points by 7 to 15 product designers of varying expertise. Together with industrial design teachers, we distilled a taxonomy of the methods designers use to accurately sketch in perspective and used it to label each stroke of the 214 sketches drawn from one of the two viewpoints. We registered each sketch to its reference 3D model by annotating sparse correspondences. We made an analysis of our annotated sketches, which reveals systematic drawing strategies over time and shapes. We also developed several applications of our dataset for sketch-based modeling and sketch filtering. We will distribute our dataset under the Creative Commons CC0 license to foster research in digital sketching.

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

6.1.4. Line Drawing Vectorization using a Global Parameterization

Participants: Tibor Stanko, Adrien Bousseau.

Despite the progress made in recent years, automatic vectorization of line drawings remains a difficult task. For drawings containing noise, holes and oversketched strokes, the main challenges are the correct classification of curve junctions, filling the missing information, and clustering multiple strokes corresponding to a single curve. We propose a new line drawing vectorization method, which addresses the above challenges in a global manner. Inspired by the quad meshing literature, we compute a global parametrization of the input drawing, such that nearby strokes are mapped to a single straight line in the parametric domain, while junctions are mapped to straight line intersections. The vectorization is obtained by following the straight lines in the parametric domain, and mapping them back to the original space. This allows us to process both clean and sketchy drawings.

This work is an ongoing collaboration with David Bommes from University of Bern, Mikhail Bessmeltsev from University of Montreal, and Justin Solomon from MIT.

6.1.5. Image-Space Motion Rigidification for Video Stylization

Participants: Johanna Delanoy, Adrien Bousseau.

Existing video stylization methods often retain the 3D motion of the original video, making the result look like a 3D scene covered in paint rather than the 2D painting of a scene. In contrast, traditional hand-drawn animations often exhibit simplified in-plane motion, such as in the case of cut-out animations where the animator moves pieces of paper from frame to frame. Inspired by this technique, we propose to modify a video such that its content undergoes 2D rigid transforms. To achieve this goal, our approach applies motion segmentation and optimization to best approximate the input optical flow with piecewise-rigid transforms, and re-renders the video such that its content follows the simplified motion. The output of our method is a new video and its optical flow, which can be fed to any existing video stylization algorithm.

This work is a collaboration with Aaron Hertzmann from Adobe Research. It is currently under review.

6.1.6. Computational Design of Tensile Structures

Participants: David Jourdan, Adrien Bousseau.

Tensile structures are architectural shapes made of stretched elastic material that can be used to create large-span roofs. Their elastic properties make it quite challenging to obtain a specific shape, and the final shape of a tensile structure is usually found rather than imposed. We created a design tool for tensile structures that, unlike existing software, lets the user specify the shape they want and finds the closest fit.

This work is an ongoing collaboration with Melina Skouras from IMAGINE (Inria Rhone Alpes). A preliminary version was presented at JFIG (Journées Françaises d'Informatique Graphique) 2018.

6.2. Graphics with Uncertainty and Heterogeneous Content

6.2.1. Single-Image SVBRDF Capture with a Rendering-Aware Deep Network

Participants: Valentin Deschaintre, Aittala Miika, Frédéric Durand, George Drettakis, Adrien Bousseau.

Texture, highlights, and shading are some of many visual cues that allow humans to perceive material appearance in single pictures. Yet, recovering spatially-varying bi-directional reflectance distribution functions (SVBRDFs) from a single image based on such cues has challenged researchers in computer graphics for decades. We tackle lightweight appearance capture by training a deep neural network to automatically extract and make sense of these visual cues. Once trained, our network is capable of recovering per-pixel normal, diffuse albedo, specular albedo and specular roughness from a single picture of a flat surface lit by a hand-held flash. We achieve this goal by introducing several innovations on training data acquisition and network design. For training, we leverage a large dataset of artist-created, procedural SVBRDFs which we sample and render under multiple lighting directions. We further amplify the data by material mixing to cover a wide diversity of shading effects, which allows our network to work across many material classes. Motivated by the observation that distant regions of a material sample often offer complementary visual cues, we design a network that combines an encoder-decoder convolutional track for local feature extraction with a fully-connected track for *global feature* extraction and propagation. Many important material effects are view-dependent, and as such ambiguous when observed in a single image. We tackle this challenge by defining the loss as a differentiable SVBRDF similarity metric that compares the *renderings* of the predicted maps against renderings of the ground truth from several lighting and viewing directions. Combined together, these novel ingredients bring clear improvement over state of the art methods for single-shot capture of spatially varying BRDFs.

The work was published in ACM Transactions on Graphics and presented at SIGGRAPH 2018 [13], and was cited by several popular online ressources (https://venturebeat.com/2018/08/15/researchers-develop-ai-that-can-re-create-real-world-lighting-and-reflections/, https://www.youtube.com/watch?v=UkWnExEFADI).

Figure 6. From a single flash photograph of a material sample (insets), our deep learning approach predicts a spatially-varying BRDF. See supplemental materials for animations with a moving light.

6.2.2. Material Acquisition using an Arbitrary Number of Inputs

Participants: Valentin Deschaintre, Aittala Miika, Frédéric Durand, George Drettakis, Adrien Bousseau.

Single-image material acquisition methods try to solve the very ill-posed problem of appearance to parametric BRDF. We explore different acquisition configurations to solve the most important ambiguities while still focusing on convenience of acquisition. Our main exploration directions are multiple lights and view angles over multiple pictures. This is possible thanks to the use of deep learning and in-line input data rendering, allowing us to easily explore a wide variety of configurations simultaneously. We also specialize our network architecture to make the most of an arbitrary number of input, provided in any order.

6.2.3. Exploiting Repetitions for Image-Based Rendering of Facades

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

Street-level imagery is now abundant but does not have sufficient capture density to be usable for Image-Based Rendering (IBR) of facades. We presented a method that exploits repetitive elements in facades – such as windows – to perform data augmentation, in turn improving camera calibration, reconstructed geometry and overall rendering quality for IBR. The main intuition behind our approach is that a few views of several instances of an element provide similar information to many views of a single instance of that element. We first select similar instances of an element from 3-4 views of a facade and transform them into a common coordinate system (Fig. 7 (a)), creating a “platonic” element. We use this common space to refine the camera calibration of each view of each instance (Fig. 7 (b)) and to reconstruct a 3D mesh of the element with multi-view stereo, that we regularize to obtain a piecewise-planar mesh aligned with dominant image contours (Fig. 7 (c)). Observing the same element under multiple views also allows us to identify reflective areas – such as glass panels – (Fig. 7 (d)) which we use at rendering time to generate plausible reflections using an environment map. We also combine information from multiple viewpoints to augment our initial set of views of the elements (Fig. 7 (e)). Our detailed 3D mesh, augmented set of views, and reflection mask enable image-based rendering of much higher quality than results obtained using the input images directly(Fig. 7 (f)).

The work was published in Computer Graphics Forum, presented at the Eurographics Symposium on Rendering 2018 [16].

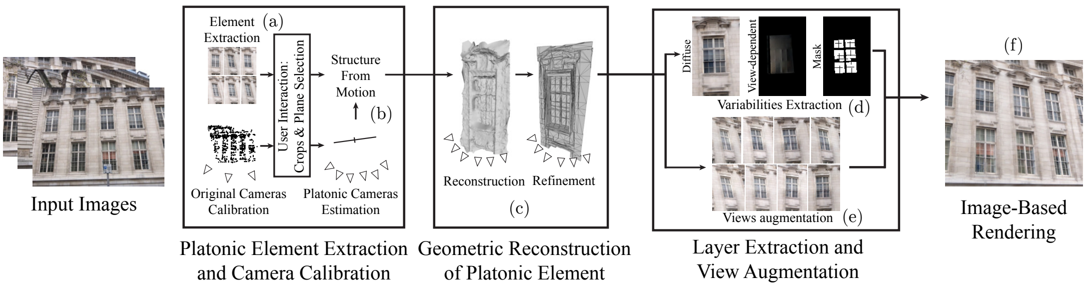

Figure 7. Overview of our technique for Image-Based Rendering of facades.

6.2.4. 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, allowing a form of scene editing for IBR. This requires multi-view inpainting of the input images. Previous methods suffer from several major limitations: they lack 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. We show results on large indoors and outdoors environments.

The work was presented at the ACM SIGGRAPH I3D Symposium on Interactive Computer Graphics and Games [19].

6.2.5. Deep Blending for Free-Viewpoint Image-Based Rendering

Participants: Julien Philip, George Drettakis.

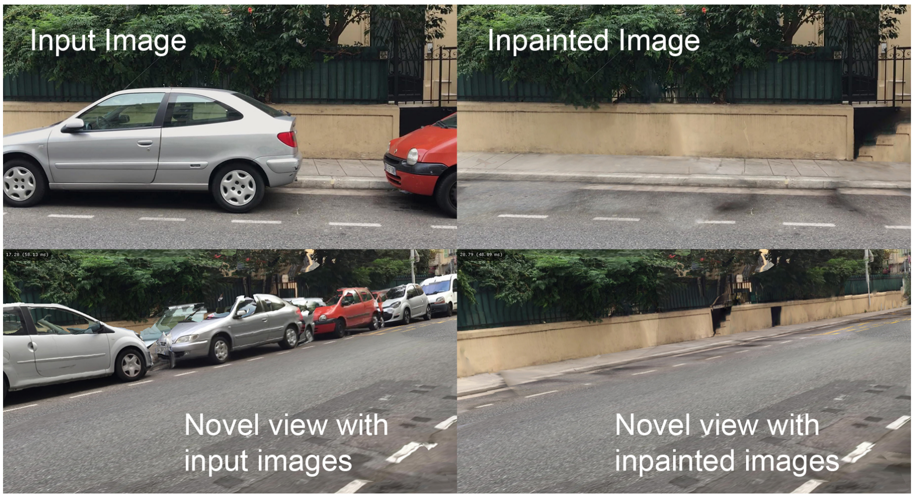

Figure 8. Our plane-based multi-view inpainting method allows us to remove cars in this large urban scene.

Free-viewpoint image-based rendering (IBR) is a standing challenge. IBR methods combine warped versions of input photos to synthesize a novel view. The image quality of this combination is directly affected by geometric inaccuracies of multi-view stereo (MVS) reconstruction and by view- and image-dependent effects that produce artifacts when contributions from different input views are blended. We present a new deep learning approach to blending for IBR, in which we use held-out real image data to learn blending weights to combine input photo contributions. Our Deep Blending method requires us to address several challenges to achieve our goal of interactive free-viewpoint IBR navigation. We first need to provide sufficiently accurate geometry so the Convolutional Neural Network (CNN) can succeed in finding correct blending weights. We do this by combining two different MVS reconstructions with complementary accuracy vs. completeness tradeoffs. To tightly integrate learning in an interactive IBR system, we need to adapt our rendering algorithm to produce a fixed number of input layers that can then be blended by the CNN. We generate training data with a variety of captured scenes, using each input photo as ground truth in a held-out approach. We also design the network architecture and the training loss to provide high quality novel view synthesis, while reducing temporal flickering artifacts. Our results demonstrate free-viewpoint IBR in a wide variety of scenes, clearly surpassing previous methods in visual quality, especially when moving far from the input cameras.

This work is a collaboration with Peter Hedman and Gabriel Brostow from University College London and True Price and Jan-Michael Frahm from University of North Carolina at Chapel Hill. It was published in ACM Transactions on Graphics and presented at SIGGRAPH Asia 2018 [14].

6.2.6. Thin Structures in Image Based Rendering

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

This work proposes a novel method to handle thin structures in Image-Based Rendering (IBR), and specifically structures supported by simple geometric shapes such as planes, cylinders, etc. These structures, e.g. railings, fences, oven grills etc, are present in many man-made environments and are extremely challenging for multi-view 3D reconstruction, representing a major limitation of existing IBR methods. Our key insight is to exploit multi-view information to compute multi-layer alpha mattes to extract the thin structures. We use two multi-view terms in a graph-cut segmentation, the first based on multi-view foreground color prediction and the second ensuring multi-view consistency of labels. Occlusion of the background can challenge reprojection error calculation and we use multi-view median images and variance, with multiple layers of thin structures.

Figure 9. Deep Blending for Free-Viewpoint Image-Based Rendering

Our end-to-end solution uses the multi-layer segmentation to create per-view mattes and the median colors and variance to extract a clean background. We introduce a new multi-pass IBR algorithm based on depth-peeling to allow free-viewpoint navigation of multi-layer semi-transparent thin structures. Our results show significant improvement in rendering quality for thin structures compared to previous image-based rendering solutions.

The work was published in the journal Computer Graphics Forum, and was presented at the Eurographics Symposium on Rendering (EGSR) 2018 [17].

Figure 10. Thin structures are present in many environments, both indoors and outdoors (far left). Our solution extracts multi-view mattes together with clean background images and geometry (center). These elements are used by our multi-layer rendering algorithm that allows free-viewpoint navigation, with significantly improved quality compared to previous solutions (right).

6.2.7. Multi-Scale Simulation of Nonlinear Thin-Shell Sound with Wave Turbulence

Participants: Gabriel Cirio, George Drettakis.

Thin shells – solids that are thin in one dimension compared to the other two – often emit rich nonlinear sounds when struck. Strong excitations can even cause chaotic thin-shell vibrations, producing sounds whose energy spectrum diffuses from low to high frequencies over time – a phenomenon known as wave turbulence. It is all these nonlinearities that grant shells such as cymbals and gongs their characteristic "glinting" sound. Yet, simulation models that efficiently capture these sound effects remain elusive. In this project, we proposed a physically based, multi-scale reduced simulation method to synthesize nonlinear thin-shell sounds. We first split nonlinear vibrations into two scales, with a small low-frequency part simulated in a fully nonlinear way, and a high-frequency part containing many more modes approximated through time-varying linearization. This allows us to capture interesting nonlinearities in the shells' deformation, tens of times faster than previous approaches. Furthermore, we propose a method that enriches simulated sounds with wave turbulent sound details through a phenomenological diffusion model in the frequency domain, and thereby sidestep the expensive simulation of chaotic high-frequency dynamics. We show several examples of our simulations, illustrating the efficiency and realism of our model, see Fig. 11.

This work is a collaboration with Ante Qu from Stanford, Eitan Grinspun and Changzi Zheng from Columbia. This work was published at ACM Transactions on Graphics, and presented at SIGGRAPH 2018 [11].

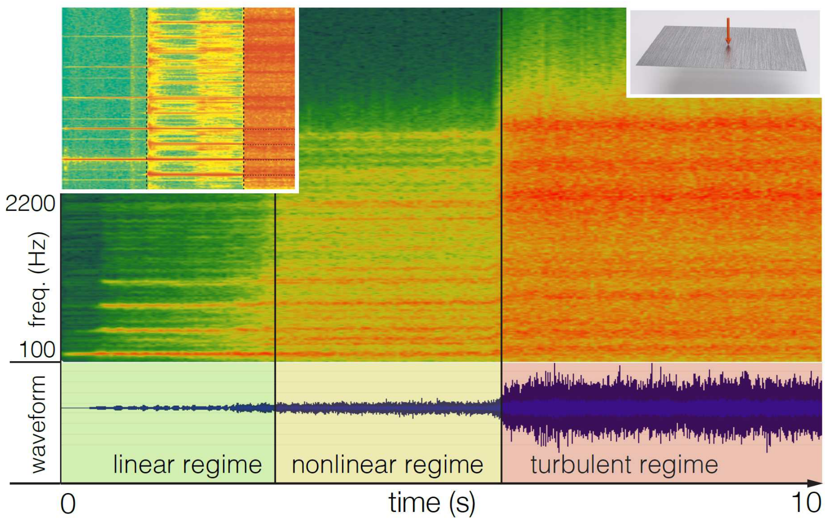

Figure 11. Thin-shell bifurcation. We excite a thin plate with increasing forces (the red arrow in the top-right inset) and simulate its dynamical responses). As the force increases, its vibration bifurcates, changing from linear vibration (left) to nonlinear (middle), and finally moving into a turbulent regime (right). This spectrogram is generated without any wave turbulence enrichment, indicating that model is able to capture chaos, albeit in low frequencies. We note that this spectrogram is qualitatively close to spectrograms from physical experiments, shown in the top-left inset (Image courtesy of Cyril Touzé).

6.2.8. Learning to Relight Multi-View Photographs from Synthetic Data

Participants: Julien Philip, George Drettakis.

We introduce an image relighting method that allows users to alter the lighting in their photos given multiple views of the same scene. Our method uses a deep convolutional network trained on synthetic photorealistic images. The use of a 3D reconstruction of the surroundings allows to guide the relighting process.

This ongoing project is a collaboration with Tinghui Zhou and Alexei A. Efros from UC Berkeley, and Michael Gharbi from Adobe research.

6.2.9. Exploiting Semantic Information for Street-level Image-Based Rendering

Participants: Simon Rodriguez, George Drettakis.

Following our work on facade rendering (Sec. 6.2.3), this ongoing project explores the use of semantic segmentation to inform Image-Based Rendering algorithms. In particular, we plan to devise algorithms that adapt to different types of objects in the scene (cars, buildings, trees).

6.2.10. Casual Video Based Rendering of Stochastic Phenomena

Participants: Theo Thonat, Miika Aittala, Frédéric Durand, George Drettakis.

The goal of this work is to extend traditional Image Based Rendering to capture subtle motions in real scenes. We want to allow free-viewpoint navigation with casual capture, such as a user taking photos and videos with a single smartphone or DSLR camera, and a tripod. We focus on stochastic time-dependent textures such as leaves in the wind, water or fire, to cope with the challenge of using unsynchronized videos.

This ongoing work is a collaboration with Sylvain Paris from Adobe Research.

6.2.11. Cutting-Edge VR/AR Display Technologies

Participant: Koulieris Georgios.

Near-eye (VR/AR) displays suffer from technical, interaction as well as visual quality issues which hinder their commercial potential. We presented a tutorial that delivered an overview of cutting-edge VR/AR display technologies, focusing on technical, interaction and perceptual issues which, if solved, will drive the next generation of display technologies. The most recent advancements in near-eye displays were presented providing (i) correct accommodation cues, (ii) near-eye varifocal AR, (iii) high dynamic range rendition, (iv) gaze-aware capabilities, either predictive or based on eye-tracking as well as (v) motion-awareness (Fig. 12). Future avenues for academic and industrial research related to the next generation of AR/VR display technologies were analyzed.

This work is a collaboration with Kaan Akşit (NVIDIA), Christian Richardt (University of Bath), Rafal Mantiuk (University of Cambridge) and Katerina Mania (Technical University of Crete). The work was presented at IEEE VR 2018, 18-22 March, Reutlingen, Germany [18].

Figure 12. We presented novel display technologies, including but not limited to (left-to-right) varifocal augmented reality displays, body-tracking displays and focus-tunable displays.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

- Valentin Deschaintre has a CIFRE PhD fellowship on Material Acquisition using Machine Learning, in collaboration with Optis - Ansys, a company specialized in material acquisition and rendering.
- As part of a long standing collaboration with Adobe, Theo Thonnat interned with Sylvain Paris (Boston), Julien Philip works with Michael Gharbi (San Francisco) and J. Delanoy with Aaron Hertzmann (San Francisco).
- Adrien Bousseau and Bastien Wailly worked with the InriaTech engineers to implement a sketch recognition engine in the context of a collaboration with the start-up EpicNPoc.

8. Partnerships and Cooperations

8.1. National Initiatives

8.1.1. 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. 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. PicPlay involved the development of several tools for converting and processing datasets. During this year we established contacts with industrial partners in the automobile industry and in the construction/public works industry who expressed interest in using the technology in their projects for visualization and navigation of captured environments.

8.2. European Initiatives

8.2.1. FP7 & H2020 Projects

8.2.1.1. ERC D3

Participants: Yulia Gryaditskaya, Tibor Stanko, Bastien Wailly, David Jourdan, Adrien Bousseau.

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.

8.2.1.2. ERC FunGraph

Participants: Sébastien Morgenthaler, George Drettakis, Rada Deeb, Diolatzis Stavros.

The ERC Advanced Grant FunGraph proposes a new methodology by introducing the concepts of rendering and input uncertainty. We define output or rendering uncertainty as the expected error of a rendering solution over the parameters and algorithmic components used with respect to an ideal image, and input uncertainty as the expected error of the content over the different parameters involved in its generation, compared to an ideal scene being represented. Here the ideal scene is a perfectly accurate model of the real world, i.e., its geometry, materials and lights; the ideal image is an infinite resolution, high-dynamic range image of this scene.

By introducing methods to estimate rendering uncertainty we will quantify the expected error of previously incompatible rendering components with a unique methodology for accurate, approximate and image-based renderers. This will allow FunGraph to define unified rendering algorithms that can exploit the advantages of these very different approaches in a single algorithmic framework, providing a fundamentally different approach to rendering. A key component of these solutions is the use of captured content: we will develop methods to estimate input uncertainty and to propagate it to the unified rendering algorithms, allowing this content to be exploited by all rendering approaches.

The goal of FunGraph is to fundamentally transform computer graphics rendering, by providing a solid theoretical framework based on uncertainty to develop a new generation of rendering algorithms. These algorithms will fully exploit the spectacular – but previously disparate and disjoint – advances in rendering, and benefit from the enormous wealth offered by constantly improving captured input content.

8.2.1.3. Emotive

Participants: Julien Philip, Sebastiàn Vizcay, George Drettakis.

https://emotiveproject.eu/

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 will do 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. In 2018, we developed a mixed reality plugin for Unity that allows the use of IBR in a VR experience used in one of the EMOTIVE user experiences using a VIVE HMD.

8.3. International Initiatives

8.3.1. Inria International Partners

8.3.1.1. Informal International Partners

We maintain close collaborations with international experts, including

- University College London (G. Brostow, P. Hedman)
- UC Berkeley (A. Efros)
- Purdue University (D. Aliaga)
- George Mason University (Y. Gingold)
- Tu Delft (M. Sypesteyn, J. W. Hoftijzer and S. Pont)

8.4. International Research Visitors

8.4.1. Visits of International Scientists

- Carol O'Sullivan, Trinity College Dublin, visited the group for one week in August.
- Peter Hedman, University College London, visited us for a few days in July.
- Miika Aittala, MIT, visited the group for one month in July.
- Yotam Gingold, George Mason University, visited the group for one month in June.

8.4.2. Visits to International Teams

8.4.2.1. Research Stays Abroad

Several students and postdocs visited our international collaborators:

- Yulia Gryaditskaya and Valentin Deschaintre visited the research group of Fredo Durand at MIT for two weeks. They presented their work to several groups (HCI, geometry, computer graphics).
- Tibor Stanko spent two weeks at RWTH Aachen University, Germany, to collaborate with David Bommes.
- Johanna Delanoy did a 3-months internship at Adobe Research (San Francisco) to collaborate with Aaron Hertzmann.
- Julien Philip spent a week at University College London to visit Gabriel Brostow and five weeks at University of California, Berkeley, to visit Alexei A. Efros. During this visit, he presented his work to the computer graphics groups of Stanford and UC Berkeley.

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Scientific Events Organisation

9.1.1.1. General Chair, Scientific Chair

G. Drettakis chairs the Eurographics (EG) working group on Rendering, and the steering committee of EG Symposium on Rendering.

9.1.2. Scientific Events Selection

9.1.2.1. Member of the Conference Program Committees

- George Drettakis participated in the program committee of Eurographics Symposium on Rendering (EGSR) 2018 and Eurographics 2018 STAR.
- Adrien Bousseau participated in the program committee of Eurographics 2018 and SIGGRAPH 2018.

9.1.3. Journal

9.1.3.1. Member of the Editorial Boards

- George Drettakis was an Associate Editor of ACM Transactions on Graphics until May 2018, and of Computational Visual Media (CVM). He also chaired the ACM TOG Editor-in-Chief (EiC) search committee, which appointed Marc Alexa as ACM TOG EiC.
- Adrien Bousseau is an Associate Editor of The Visual Computer Journal.

9.1.3.2. Reviewer - Reviewing Activities

- Yulia Gryaditskaya has been a reviewer for Eurographics.
- Tibor Stanko has been a reviewer for DGCI 2019 – International Conference on Discrete Geometry for Computer Imagery, and the Computers and Graphics journal.
- Adrien Bousseau was reviewer for SIGGRAPH Asia 2018, ACM Transactions on Graphics, ACM CHI.
- George Drettakis was reviewer for SIGGRAPH.

9.1.4. Invited Talks

- George Drettakis gave a keynote talk at Driving Simulation Conference (DSC '18) (http://dsc2018.org/)
- Adrien Bousseau gave an invited talk at the SophIA summit event on artificial intelligence (http://sophia-summit.fr/) and a keynote talk at the Skin of Things workshop on material perception and depiction (https://theskinofthings.github.io/).
- Tibor Stanko was invited to present his PhD work at Journées Informatique et Géométrie (JIG) 2018 in Lyon.

9.1.5. Scientific Expertise

Adrien Bousseau reviewed a grant proposal for the ANR.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

Licence : Simon Rodriguez, theoretical computer science (L1) and software engineering (L3), 40h eq. TD, Polytech Nice-Sophia - Université Côte d'Azur (France)

Licence : Julien Philip, Introduction au Web et Application du Web, L2, 64h, Polytech Nice-Sophia - Université Côte d'Azur (France)

Master: George Drettakis and Adrien Bousseau, Introduction to Computer Graphics, M1, 20h eq. TD, Université Côte d'Azur (France)

9.2.2. Supervision

PhD : Jean-Dominique Favreau, Compact image vectorization by stochastic approaches, Université Côte d'Azur, March 15th 2018, Florent Lafarge (Titane) and Adrien Bousseau

PhD in progress : Johanna Delanoy, Data-driven sketch-based modeling, since October 2015, Adrien Bousseau

PhD in progress : Valentin Deschaintre, Data-driven material capture, since November 2016, Adrien Bousseau and George Drettakis

PhD in progress : David Jourdan, Interactive architectural design, since October 2018, Adrien Bousseau and Melina Skouras (Imagine)

PhD in progress : Julien Philip, Data-driven image-based rendering and relighting, since November 2016, George Drettakis

PhD in progress : Simon Rodriguez, Leveraging semantic information in image-based rendering, since November 2016, George Drettakis

PhD in progress : Théo Thonat, Image-based rendering of thin and stochastic structures, since November 2015, George Drettakis

9.2.3. Juries

- George Drettakis was a Ph.D. Reviewer for Fabian Langguth (TU Darmstadt), and member “commité suivi de thèse” of J-P. Bauchet.
- Adrien Bousseau was a Ph.D. Reviewer for James Hennessey (University College London) and Dorian Nogeng (Ecole Polytechnique).

9.3. Popularization

9.3.1. Internal or external Inria responsibilities

- George Drettakis chairs the local “Jacques Morgenstern” Colloquium organizing committee.
- Adrien Bousseau is a member of “commité du centre” and “commité du suivi doctoral”.

9.3.2. Articles and contents

- Nice Matin published an article about the ERC project of George Drettakis.

9.3.3. Interventions

- Bastien Wailly gave a demo of our sketch-based modeling tool during *Fete de la Science*.
- Valentin Deschaintre presented his work to high school student and advised them on a small scientific project for *Math C2+*.

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] J.-D. Favreau, F. Lafarge, A. Bousseau. *Fidelity vs. Simplicity: a Global Approach to Line Drawing Vectorization*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", 2016, http://www-sop.inria.fr/reves/Basilic/2016/FLB16

[4] J.-D. FAVREAU, F. LAFARGE, A. BOUSSEAU. *Photo2ClipArt: Image Abstraction and Vectorization Using Layered Linear Gradients*, in "ACM Transactions on Graphics (SIGGRAPH Asia Conference Proceedings)", November 2017, vol. 36, n^o 6, http://www-sop.inria.fr/reves/Basilic/2017/FLB17

[5] P. HEDMAN, T. RITSCHEL, G. DRETTAKIS, G. BROSTOW. *Scalable Inside-Out Image-Based Rendering*, in "ACM Transactions on Graphics (SIGGRAPH Asia Conference Proceedings)", December 2016, vol. 35, n^o 6, http://www-sop.inria.fr/reves/Basilic/2016/HRDB16

[6] E. IARUSSI, D. BOMMES, A. BOUSSEAU. *BendFields: Regularized Curvature Fields from Rough Concept Sketches*, in "ACM Transactions on Graphics", 2015, http://www-sop.inria.fr/reves/Basilic/2015/IBB15

[7] G.-A. KOULIERIS, B. BUI, M. S. BANKS, G. DRETTAKIS. *Accommodation and Comfort in Head-Mounted Displays*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", July 2017, vol. 36, n^o 4, 11 p. , http://www-sop.inria.fr/reves/Basilic/2017/KBBD17

[8] G. NISHIDA, I. GARCIA-DORADO, D. G. ALIAGA, B. BENES, A. BOUSSEAU. *Interactive Sketching of Urban Procedural Models*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", 2016, http://www-sop.inria.fr/reves/Basilic/2016/NGGBB16

[9] R. ORTIZ-CAYON, A. DJELOUAH, F. MASSA, M. AUBRY, G. DRETTAKIS. *Automatic 3D Car Model Alignment for Mixed Image-Based Rendering*, in "International Conference on 3D Vision (3DV)", 2016, http://www-sop.inria.fr/reves/Basilic/2016/ODMAD16

[10] 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

Articles in International Peer-Reviewed Journals

[11] G. CIRIO, A. QU, G. DRETTAKIS, E. GRINSPUN, C. ZHENG. *Multi-Scale Simulation of Nonlinear Thin-Shell Sound with Wave Turbulence*, in "ACM Transactions on Graphics", 2018, vol. 37, n^o 4, 14 p. [*DOI :* 10.1145/3197517.3201361], https://hal.inria.fr/hal-01797920

[12] J. DELANOY, M. AUBRY, P. ISOLA, A. A. EFROS, A. BOUSSEAU. *3D Sketching using Multi-View Deep Volumetric Prediction*, in "Proceedings of the ACM on Computer Graphics and Interactive Techniques", July 2018, vol. 1, n^o 1, pp. 1-22, https://arxiv.org/abs/1707.08390 [*DOI :* 10.1145/3203197], https://hal.inria.fr/hal-01799600

[13] V. DESCHAINTRE, M. AITTALA, F. DURAND, G. DRETTAKIS, A. BOUSSEAU. *Single-Image SVBRDF Capture with a Rendering-Aware Deep Network*, in "ACM Transactions on Graphics", 2018, vol. 37, pp. 128 - 143 [*DOI :* 10.1145/3197517.3201378], https://hal.inria.fr/hal-01793826

[14] P. HEDMAN, J. PHILIP, T. PRICE, J.-M. FRAHM, G. DRETTAKIS, G. J. BROSTOW. *Deep Blending for Free-Viewpoint Image-Based Rendering*, in "ACM Transactions on Graphics", 2018, vol. 37, n^o 6 [*DOI :* 10.1145/3272127.3275084], https://hal.inria.fr/hal-01886237

[15] G. NISHIDA, A. BOUSSEAU, D. G. ALIAGA. *Procedural Modeling of a Building from a Single Image*, in "Computer Graphics Forum", 2018, vol. 37, n^{o} 2 [*DOI :* 10.1111/CGF.13372], https://hal.inria.fr/hal-01810207

[16] S. RODRIGUEZ, A. BOUSSEAU, F. DURAND, G. DRETTAKIS. *Exploiting Repetitions for Image-Based Rendering of Facades*, in "Computer Graphics Forum", July 2018, vol. 37, https://hal.inria.fr/hal-01814058

[17] T. THONAT, A. DJELOUAH, F. DURAND, G. DRETTAKIS. *Thin Structures in Image Based Rendering*, in "Computer Graphics Forum", 2018, vol. 37, https://hal.inria.fr/hal-01817948

International Conferences with Proceedings

[18] G.-A. KOULIERIS, K. AKSIT, C. RICHARDT, R. MANTIUK, K. MANIA. *Cutting-Edge VR/AR Display Technologies (Gaze-, Accommodation-, Motion-aware and HDR-enabled)*, in "IEEE VR 2018 - 25th IEEE Conference on Virtual Reality and 3D User Interfaces", Reutlingen, Germany, March 2018, https://hal.inria.fr/hal-01695622

[19] J. PHILIP, G. DRETTAKIS. *Plane-Based Multi-View Inpainting for Image-Based Rendering in Large Scenes*, in "I3D 2018 - ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games", Montréal, Canada, May 2018, pp. 1-11 [*DOI :* 10.1145/3190834.3190846], https://hal.inria.fr/hal-01745568