Activity Report 2019

Project-Team GRAPHDECO

GRAPHics and DEsign with hEterogeneous COntent

RESEARCH CENTER
Sophia Antipolis - Méditerranée

THEME
Interaction and visualization

Table of contents

1. Team, Visitors, External Collaborators **2**
2. Overall Objectives **2**
3. Research Program **4**
3.1.1. Computer-Assisted Design with Heterogeneous Representations 4
3.1.2. Graphics with Uncertainty and Heterogeneous Content 5
4. Highlights of the Year **6**
5. New Software and Platforms **7**
5.1. SynDraw 7
5.2. DeepSketch 7
5.3. DPP 7
5.4. sibr-core 8
5.5. SGTDGP 8
5.6. Unity IBR 9
5.7. DeepRelighting 9
5.8. SingleDeepMat 9
5.9. MultiDeepMat 10
5.10. GuidedDeepMat 10
6. New Results **10**
6.1. Computer-Assisted Design with Heterogeneous Representations 10
6.1.1. Combining Voxel and Normal Predictions for Multi-View 3D Sketching 10
6.1.2. Video Motion Stylization by 2D Rigidification 11
6.1.3. Multi-Pose Interactive Linkage Design 12
6.1.4. Extracting Geometric Structures in Images with Delaunay Point Processes 12
6.1.5. Integer-Grid Sketch Vectorization 13
6.1.6. Surfacing Sparse Unorganized 3D Curves using Global Parametrization 13
6.1.7. OpenSketch: A Richly-Annotated Dataset of Product Design Sketches 13
6.1.8. Intersection vs. Occlusion: a Discrete Formulation of Line Drawing 3D Reconstruction 14
6.1.9. Data-driven sketch segmentation 15
6.1.10. Stroke-based concept sketch generation 15
6.1.11. Designing Programmable, Self-Actuated Structures 15
6.2. Graphics with Uncertainty and Heterogeneous Content 16
6.2.1. Multi-view relighting using a geometry-aware network 16
6.2.2. Flexible SVBRDF Capture with a Multi-Image Deep Network 16
6.2.3. Guided Acquisition of SVBRDFs 17
6.2.4. Mixed rendering and relighting for indoor scenes 17
6.2.5. DiCE: Dichoptic Contrast Enhancement for VR and Stereo Displays 18
6.2.6. Compositing Real Scenes using a relighting Network 18
6.2.7. Image-based Rendering of Urban Scenes based on Semantic Information 18
6.2.8. Synthetic Data for Image-based Rendering 19
6.2.9. Densified Surface Light Fields for Human Capture Video 19
6.2.10. Deep Bayesian Image-based Rendering 19
6.2.11. Path Guiding for Metropolis Light Transport 19
6.2.12. Improved Image-Based Rendering with Uncontrolled Capture 19
6.2.13. Practical video-based rendering of dynamic stationary environments 20
7. Bilateral Contracts and Grants with Industry **20**
7.1. Bilateral Contracts with Industry 20
7.2. Bilateral Grants with Industry 20
8. Partnerships and Cooperations **20**
8.1. National Initiatives 20

8.2. European Initiatives 21
8.2.1.1. D[3]: Drawing Interpretation for 3D Design 21
8.2.1.2. ERC FunGraph 21
8.2.1.3. Emotive 22
8.3. International Initiatives 22
8.3.1.1. Informal International Partners 22
8.3.1.2. Inria International Chairs 23
8.4. International Research Visitors 23
8.4.1. Visits of International Scientists 23
8.4.2. Visits to International Teams 23
9. Dissemination 23
9.1. Promoting Scientific Activities 23
9.1.1. Scientific Events: Organisation 23
9.1.2. Scientific Events: Selection 23
9.1.3. Journal 23
9.1.3.1. Member of the Editorial Boards 23
9.1.3.2. Reviewer - Reviewing Activities 23
9.1.4. Invited Talks 24
9.1.5. Leadership within the Scientific Community 24
9.1.6. Scientific Expertise 24
9.1.7. Research Administration 24
9.1.8. Interventions at Conferences 24
9.2. Teaching - Supervision - Juries 24
9.2.1. Teaching 24
9.2.2. Supervision 24
9.2.3. Juries 25
9.3. Popularization 25
9.3.1. Internal or external Inria responsibilities 25
9.3.2. Articles and contents 25
10. Bibliography 25

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

Rada Deeb [Inria, Post-Doctoral Fellow]
Thomas Leimkuhler [Inria, Post-Doctoral Fellow, from Sep 2019]
Tibor Stanko [Inria, Post-Doctoral Fellow]
Yulia Gryaditskaya [Inria, Post-Doctoral Fellow]

PhD Students

Johanna Delanoy [Inria, PhD Student, until Jun 2019]
Valentin Deschaintre [Optis/ANSYS, PhD Student, until Oct 2019, granted by CIFRE]
Stavros Diolatzis [Inria, PhD Student]
Felix Hahnlein [Inria, PhD Student, from Feb 2019]
David Jourdan [Inria, PhD Student]
Georgios Kopanas [Inria, PhD Student, from Dec 2019]
Julien Philip [Inria, PhD Student]
Siddhant Prakash [Inria, PhD Student, from Nov 2019]
Simon Rodriguez [Ministère de l'Enseignement Supérieur et de la Recherche, PhD Student]
Theo Thonat [Inria, PhD Student, until Jun 2019]

Technical staff

Sebastien Morgenthaler [Inria, Engineer]
Siddhant Prakash [Inria, Engineer, from Feb 2019 until Oct 2019]
Sebastian Vizcay [Inria, Engineer, until Oct 2019]
Bastien Wailly [Inria, Engineer]

Interns and Apprentices

Baptiste Nicolet [Inria, École polytechnique, Intern, Apr-Aug 2019]
Aly Bora Ulusoy [Inria, École polytechnique, Intern, Apr-Aug 2019]

Administrative Assistant

Sophie Honnorat [Inria, Administrative Assistant]

External Collaborator

Frédo 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.

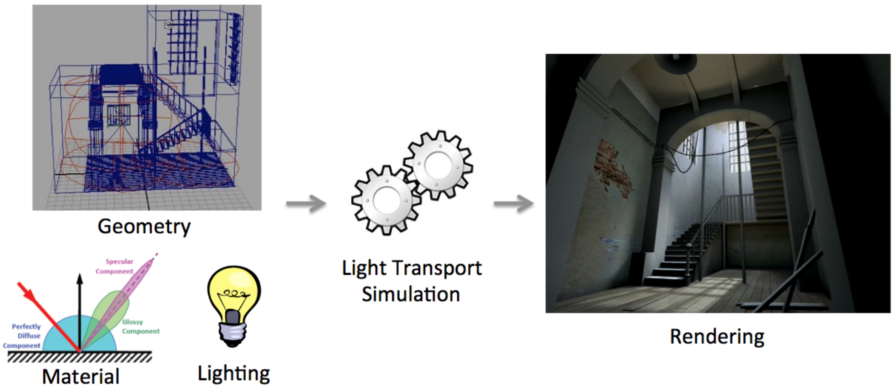

Figure 1. Traditional computer graphics pipeline. Rendering from www.thegnomonworkshop.com

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. Fig. 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

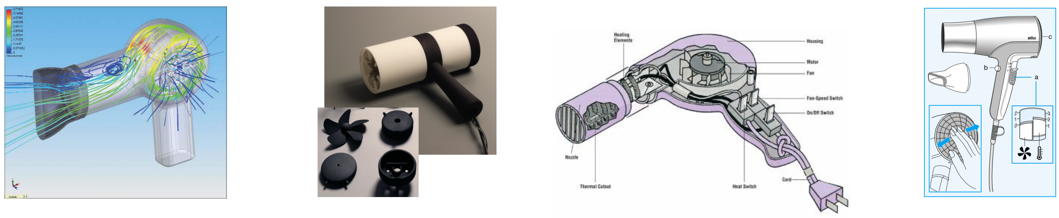

(f) Simulation (e) 3D Printing (g) Technical diagram (h) Instructions

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

The SIGGRAPH paper "Multi-view relighting using a geometry-aware network" by J. Philip et al. [19] was presented at the Adobe Max event in November 2019 in San Fransisco. The project was part of the 11 projects selected out of 200 to be presented at Adobe MAX under the name project #LightRightSneak (Link to the video of the event).

G. Drettakis presented at the French Academy of Sciences days at Sophia-Antipolis in June: video on the Academy of Sciences site.

4.1.1. Awards

Jean-Dominique Favreau (co-supervised with the TITANE team) received the best Ph.D. thesis award 2019 (assessit prize) from IGRV.

5. New Software and Platforms

5.1. SynDraw

KEYWORDS: Non-photorealistic rendering - Vector-based drawing - Geometry Processing

FUNCTIONAL DESCRIPTION: The SynDraw library extracts occluding contours and sharp features over a 3D shape, computes all their intersections using a binary space partitioning algorithm, and finally performs 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. Finally, the library embeds tools allowing one to add noise into an SVG drawing, in order to generate multiple images from a single sketch. SynthDraw is based on the geometry processing library libIGL.

RELEASE FUNCTIONAL DESCRIPTION: This first version extracts occluding contours, boundaries, creases, ridges, valleys, suggestive contours and demarcating curves. Visibility is computed with a view graph structure. Lines can be aggregated and/or filtered. Labels and outputs include: line type, visibility, depth and aligned normal map.

- Authors: Adrien Bousseau, Bastien Wailly and Adele Saint-Denis
- Contact: Bastien Wailly

5.2. 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.

- Authors: Adrien Bousseau and Bastien Wailly
- Contact: Adrien Bousseau

5.3. DPP

Delaunay Point Process for image analysis

KEYWORDS: Computer vision - Shape recognition - Delaunay triangulation - Stochastic process

FUNCTIONAL DESCRIPTION: The software extract 2D geometric structures (planar graphs, polygons...) from images

- Participants: Jean-Dominique Favreau, Florent Lafarge and Adrien Bousseau
- Contact: Florent Lafarge
- Publication: Extracting Geometric Structures in Images with Delaunay Point Processes

5.4. sibr-core

System for Image-Based Rendering

KEYWORD: Graphics

SCIENTIFIC DESCRIPTION: Core functionality to support Image-Based Rendering research. The core provides basic support for camera calibration, multi-view stereo meshes and basic image-based rendering functionality. Separate dependent repositories interface with the core for each research project. This library is an evolution of the previous SIBR software, but now is much more modular.

We plan to release the core module, as well as the code for several of our research papers, as well as papers from other authors for comparisons and benchmark purposes.

FUNCTIONAL DESCRIPTION: sibr-core 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.

- Authors: Sebastien Bonopera, Jérôme Esnault, Siddhant Prakash, Simon Rodriguez, Théo Thonat, Gaurav Chaurasia, Julien Philip and George Drettakis
- Contact: George Drettakis

5.5. 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-core 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-core, e.g., compute shadow maps of both exact and reconstructed representations etc.

- Authors: Laurent Boiron, Sébastien Morgenthaler, Georgios Kopanas, Julien Philip and George Drettakis
- Contact: George Drettakis

5.6. 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.

- Authors: Sebastian Vizcay and George Drettakis
- Contact: George Drettakis

5.7. DeepRelighting

Deep Geometry-Aware Multi-View Relighting

KEYWORD: Graphics

SCIENTIFIC DESCRIPTION: Implementation of the paper: Multi-view Relighting using a Geometry-Aware Network (https://hal.inria.fr/hal-02125095), based on the sibr-core library.

- Participants: Julien Philip and George Drettakis
- Contact: George Drettakis
- Publication: https://hal.inria.fr/hal-02125095

5.8. SingleDeepMat

Single-image deep material acquisition

KEYWORDS: Materials - 3D - Realistic rendering - Deep learning

SCIENTIFIC DESCRIPTION: Cook-Torrance SVBRDF parameter acquisition from a single Image using Deep learning

FUNCTIONAL DESCRIPTION: Allows material acquisition from a single picture, to then be rendered in a virtual environment. Implementation of the paper https://hal.inria.fr/hal-01793826/

RELEASE FUNCTIONAL DESCRIPTION: Based on Pix2Pix implementation by AffineLayer (Github)

- Participants: Valentin Deschaintre, Miika Aittala, Frédo Durand, George Drettakis and Adrien Bousseau
- Partner: CSAIL, MIT
- Contact: Adrien Bousseau
- Publication: Single-Image SVBRDF Capture with a Rendering-Aware Deep Network
- URL: https://team.inria.fr/graphdeco/projects/deep-materials/

5.9. MultiDeepMat

Multi-image deep material acquisition

KEYWORDS: 3D - Materials - Deep learning

SCIENTIFIC DESCRIPTION: Allows material acquisition from multiple pictures, to then be rendered in a virtual environment. Implementation of the paper https://hal.inria.fr/hal-02164993

RELEASE FUNCTIONAL DESCRIPTION: Code fully rewritten since the SingleDeepMat project, but some function are imported from it.

- Participants: Valentin Deschaintre, Miika Aittala, Frédo Durand, George Drettakis and Adrien Bousseau
- Contact: Adrien Bousseau
- Publication: Flexible SVBRDF Capture with a Multi-Image Deep Network
- URL: https://team.inria.fr/graphdeco/projects/multi-materials/

5.10. GuidedDeepMat

Guided deep material acquisition

KEYWORDS: Materials - 3D - Deep learning

SCIENTIFIC DESCRIPTION: Deep large scale HD material acquisition guided by an example small scale SVBRDF

RELEASE FUNCTIONAL DESCRIPTION: Code based on the MultiDeepMat project code.

- Participants: Valentin Deschaintre, George Drettakis and Adrien Bousseau
- Contact: Adrien Bousseau

6. New Results

6.1. Computer-Assisted Design with Heterogeneous Representations

6.1.1. Combining Voxel and Normal Predictions for Multi-View 3D Sketching

Participants: Johanna Delanoy, Adrien Bousseau.

Recent works on data-driven sketch-based modeling use either voxel grids or normal/depth maps as geometric representations compatible with convolutional neural networks. While voxel grids can represent complete objects – including parts not visible in the sketches – their memory consumption restricts them to low-resolution predictions. In contrast, a single normal or depth map can capture fine details, but multiple maps from different viewpoints need to be predicted and fused to produce a closed surface. We propose to combine these two representations to address their respective shortcomings in the context of a multi-view sketch-based modeling system. Our method predicts a voxel grid common to all the input sketches, along with one normal map per sketch. We then use the voxel grid as a support for normal map fusion by optimizing its extracted surface such that it is consistent with the re-projected normals, while being as piecewise-smooth as possible overall (Fig. 4). We compare our method with a recent voxel prediction system, demonstrating improved recovery of sharp features over a variety of man-made objects.

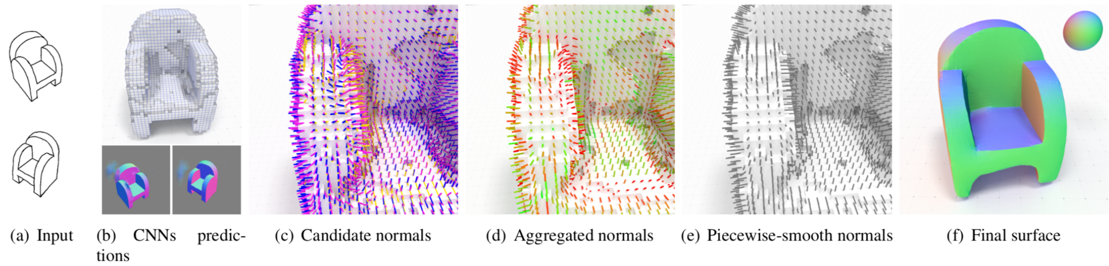

Figure 4. Our method takes as input multiple sketches of an object (a). We first apply existing deep neural networks to predict a volumetric reconstruction of the shape as well as one normal map per sketch (b). We re-project the normal maps on the voxel grid in complement to the surface normal computed from the volumetric prediction (c). We aggregate these different normals into a distribution represented by a mean vector and a standard deviation (d). We optimize this normal field to make it piecewise smooth (e) and use it to regularize the surface (f). The final surface preserves the overall shape of the predicted voxel grid as well as the sharp features of the predicted normal maps.

This work is a collaboration with David Coeurjolly from Université de Lyon and Jacques-Olivier Lachaud from Université Savoie Mont Blanc. The work was published in the journal Computer & Graphics and presented at the SMI conference [14].

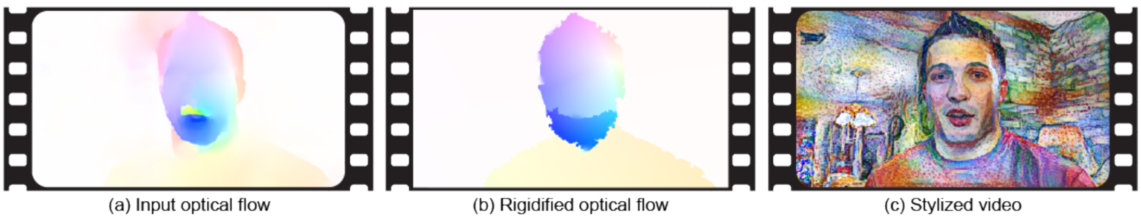

Figure 5. Our method takes as input a video and its optical flow (a). We segment the video and optimize its pixel trajectories to produce a new video that exhibits piecewise-rigid motion (b). The resulting rigidified video can be stylized with existing algorithms (c) to produce animations where the style elements (brush strokes, paper texture) produce a strong sense of 2D motion.

6.1.2. Video Motion Stylization by 2D Rigidification

Participants: Johanna Delanoy, Adrien Bousseau.

We introduce a video stylization method that increases the apparent rigidity of motion. Existing stylization methods often retain the 3D motion of the original video, making the result look like a 3D scene covered in paint rather than a 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 (Fig. 5). 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 was presented at the ACM/EG Expressive Symposium [21].

6.1.3. Multi-Pose Interactive Linkage Design

Participant: Adrien Bousseau.

We introduce an interactive tool for novice users to design mechanical objects made of 2.5D linkages. Users simply draw the shape of the object and a few key poses of its multiple moving parts. Our approach automatically generates a one-degree-of-freedom linkage that connects the fixed and moving parts, such that the moving parts traverse all input poses in order without any collision with the fixed and other moving parts. In addition, our approach avoids common linkage defects and favors compact linkages and smooth motion trajectories. Finally, our system automatically generates the 3D geometry of the object and its links, allowing the rapid creation of a physical mockup of the designed object (Fig. 6).

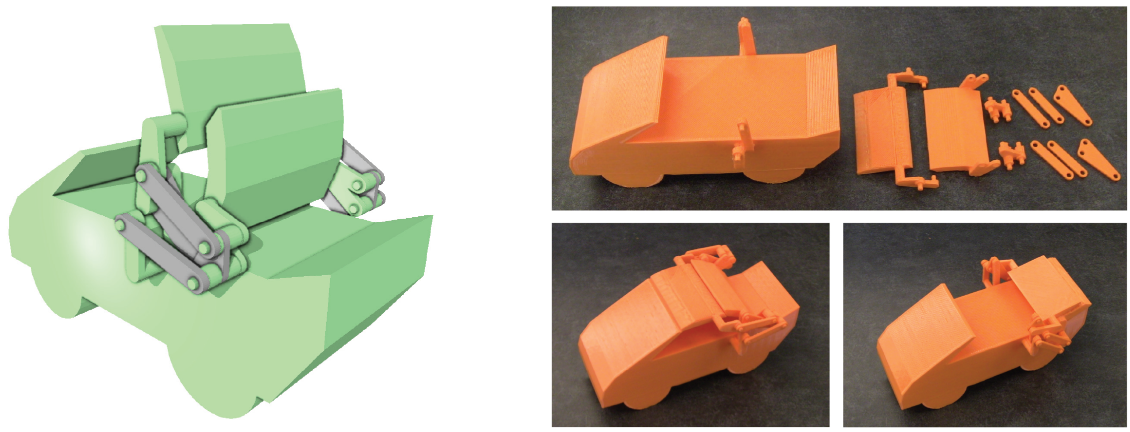

Figure 6. Our interactive system facilitates the creation (left) and fabrication (right) of mechanical objects.

This work was conducted in collaboration with Gen Nishida and Daniel G. Aliaga from Purdue University, was published in Computer Graphics Forum and presented at the Eurographics conference [18].

6.1.4. Extracting Geometric Structures in Images with Delaunay Point Processes

Participant: Adrien Bousseau.

We introduce Delaunay Point Processes, a framework for the extraction of geometric structures from images. Our approach simultaneously locates and groups geometric primitives (line segments, triangles) to form extended structures (line networks, polygons) for a variety of image analysis tasks. Similarly to traditional point processes, our approach uses Markov Chain Monte Carlo to minimize an energy that balances fidelity to the input image data with geometric priors on the output structures. However, while existing point processes struggle to model structures composed of inter-connected components, we propose to embed the point process into a Delaunay triangulation, which provides high-quality connectivity by construction. We further leverage key properties of the Delaunay triangulation to devise a fast Markov Chain Monte Carlo sampler. We demonstrate the flexibility of our approach on a variety of applications, including line network extraction, object contouring, and mesh-based image compression (see Fig. 7).

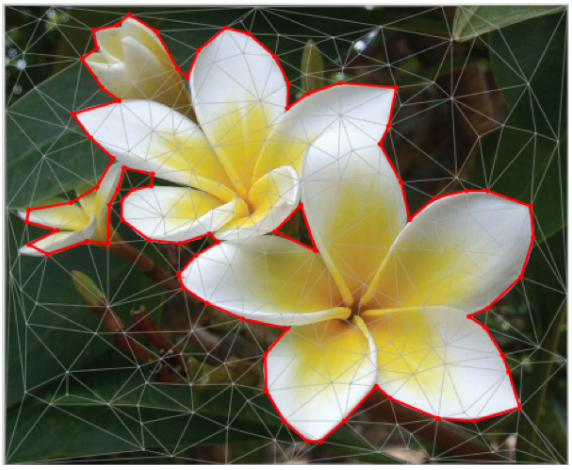

Figure 7. Our method extracts geometric structures like the countour of these flowers by optimizing a dynamic Delaunay triangulation.

This work was conducted in collaboration with Jean-Dominique Favreau and Florent Lafarge (TITANE group), and published in IEEE PAMI [16].

6.1.5. Integer-Grid Sketch Vectorization

Participants: Tibor Stanko, Adrien Bousseau.

A major challenge in line drawing vectorization is segmenting the input bitmap into separate curves. This segmentation is especially problematic for rough sketches, where curves are depicted using multiple overdrawn strokes. Inspired by feature-aligned mesh quadrangulation methods in geometry processing, we propose to extract vector curve networks by parametrizing the image with local drawing-aligned integer grids. The regular structure of the grid facilitates the extraction of clean line junctions; due to the grid's discrete nature, nearby strokes are implicitly grouped together. Our method successfully vectorizes both clean and rough line drawings, whereas previous methods focused on only one of those drawing types.

This work is an ongoing collaboration with David Bommes from University of Bern and Mikhail Bessmeltsev from University of Montreal. It is currently under review.

6.1.6. Surfacing Sparse Unorganized 3D Curves using Global Parametrization

Participants: Tibor Stanko, Adrien Bousseau.

Designers use sketching to quickly externalize ideas, often using a handful of curves to express complex shapes. Recent years have brought a plethora of new tools for creating designs directly in 3D. The output of these tools is often a set of sparse, unorganized curves. We propose a novel method for automatic conversion of such unorganized curves into clean curve networks ready for surfacing. The core of our method is a global curve-aligned parametrization, which allows us to automatically aggregate information from neighboring curves and produce an output with valid topology.

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

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

Participants: Yulia Gryaditskaya, Adrien Bousseau, Fredo Durand.

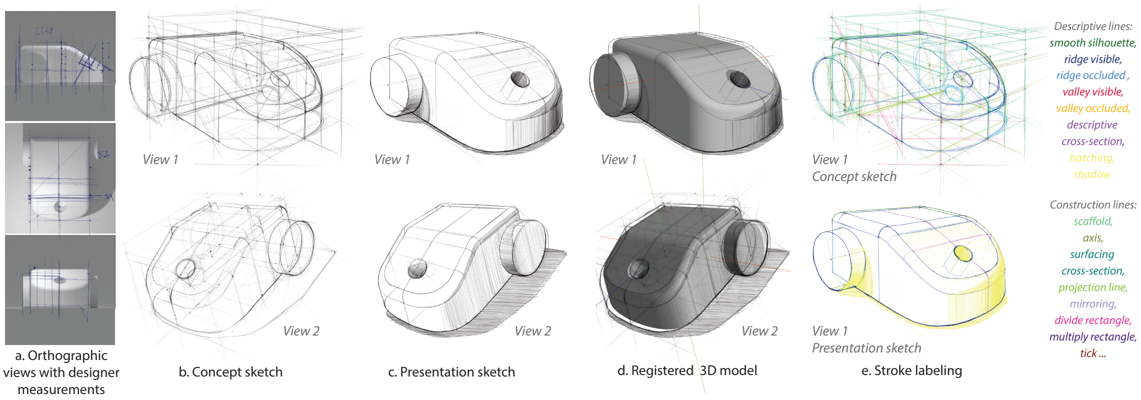

Figure 8. We showed designers three orthographic views (a) of the object and asked them to draw it from two different perspective views (b). We also asked to replicate each of their sketches as a clean presentation drawing (c). We semi-automatically registered 3D models to each sketch (d), and we manually labeled different types of lines in all concept sketches and presentation drawings from the first viewpoint (e).

Product designers extensively use sketches to create and communicate 3D shapes and thus form an ideal audience for sketch-based modeling, non-photorealistic rendering and sketch filtering. However, sketching requires significant expertise and time, making design sketches a scarce resource for the research community. We introduce *OpenSketch*, a dataset of product design sketches aimed at offering a rich source of information for a variety of computer-aided design tasks. *OpenSketch* contains more than 400 sketches representing 12 man-made objects drawn by 7 to 15 product designers of varying expertise. We provided participants with front, side and top views of these objects (Fig. 8a), and instructed them to draw from two *novel* perspective viewpoints(Fig. 8b). This drawing task forces designers to *construct the shape* from their mental vision rather than directly copy what they see. They achieve this task by employing a variety of sketching techniques and methods not observed in prior datasets. Together with industrial design teachers, we distilled a taxonomy of line types and used it to label each stroke of the 214 sketches drawn from one of the two viewpoints (Fig. 8e). While some of these lines have long been known in computer graphics, others remain to be reproduced algorithmically or exploited for shape inference. In addition, we also asked participants to produce clean presentation drawings from each of their sketches, resulting in aligned pairs of drawings of different styles (Fig. 8c). Finally, we registered each sketch to its reference 3D model by annotating sparse correspondences (Fig. 8d). We provide an analysis of our annotated sketches, which reveals systematic drawing strategies over time and shapes, as well as a positive correlation between presence of construction lines and accuracy. Our sketches, in combination with provided annotations, form challenging benchmarks for existing algorithms as well as a great source of inspiration for future developments. We illustrate the versatility of our data by using it to test a 3D reconstruction deep network trained on synthetic drawings, as well as to train a filtering network to convert concept sketches into presentation drawings. We distribute our dataset under the Creative Commons CC0 license: https://ns.inria.fr/d3/OpenSketch.

This work is a collaboration with Mark Sypesteyn, Jan Willem Hoftijzer and Sylvia Pont from TU Delft, Netherlands. This work was published at ACM Transactions on Graphics, and presented at SIGGRAPH Asia 2019 [17].

6.1.8. Intersection vs. Occlusion: a Discrete Formulation of Line Drawing 3D Reconstruction

Participants: Yulia Gryaditskaya, Adrien Bousseau, Felix Hähnlein.

The popularity of sketches in design stems from their ability to communicate complex 3D shapes with a handful of lines. Yet, this economy of means also makes sketch interpretation a challenging task, as global 3D understanding needs to emerge from scattered pen strokes. To tackle this challenge, many prior methods cast 3D reconstruction of line drawings as a global optimization that seeks to satisfy a number of geometric criteria, including orthogonality, planarity, symmetry. However, all of these methods require users to distinguish line intersections that exist in 3D from the ones that are only due to occlusions. These user annotations are critical to the success of existing algorithms, since mistakenly treating an occlusion as a true intersection would connect distant parts of the shape, with dramatic consequences on the overall optimization procedure. We propose a line drawing 3D reconstruction method that automatically discriminates 3D intersections from occlusions. This automation not only reduces user burden, it also allows our method to scale to real-world sketches composed of hundreds of pen strokes, for which the number of intersections is too high to make existing user-assisted methods practical. Our key idea is to associate each 2D intersection with a binary variable that indicates if the intersection should be preserved in 3D. Our algorithm then searches for the assignment of binary values that yields the best 3D shape, as measured with similar criteria as the ones used by prior work for 3D reconstruction. However, the combinatorial nature of this binary assignment problem prevents trying all possible configurations. Our main technical contribution is an efficient search algorithm that leverages principles of how product designers draw to reconstruct complex 3D drawings within minutes.

This work is a collaboration with Alla Sheffer (Professor at University of British Columbia) and Chenxi Liu (PhD student at University of British Columbia).

6.1.9. Data-driven sketch segmentation

Participants: Yulia Gryaditskaya, Felix Hähnlein, Adrien Bousseau.

Deep learning achieves impressive performance on image segmentation, which has motivated the recent development of deep neural networks for the related task of sketch segmentation, where the goal is to assign labels to the different strokes that compose a line drawing. However, while natural images are well represented as bitmaps, line drawings can also be represented as vector graphics, such as point sequences and point clouds. In addition to offering different trade-offs on resolution and storage, vector representations often come with additional information, such as stroke ordering and speed.

In this project, we evaluate three crucial design choices for sketch segmentation using deep-learning: which sketch representation to use, which information to encode in this representation, and which loss function to optimize. Our findings suggest that point clouds represent a competitive alternative to bitmaps for sketch segmentation, and that providing extra-geometric information improves performance.

6.1.10. Stroke-based concept sketch generation

Participants: Felix Hähnlein, Yulia Gryaditskaya, Adrien Bousseau.

State-of-the-art non-photorealisting rendering algorithms can generate lines representing salient visual features on objects. However, very few methods exist for generating lines outside of an object, as is the case for most construction lines, used in technical drawings and design sketches. Furthermore, most methods do not generate human-like strokes and do not consider the drawing order of a sketch.

In this project, we address these issues by proposing a reinforcement learning framework, where a virtual agent tries to generate a construction sketch of a given 3D model. One key element of our approach is the study and the mathematical formalization of drawing strategies used by industrial designers.

6.1.11. Designing Programmable, Self-Actuated Structures

Participants: David Jourdan, Adrien Bousseau.

Self-actuated structures are material assemblies that can deform from an initially simpler state to a more complex, curved one, by automatically deforming to shape. Most relevant to applications in manufacturing are self-actuated shapes that are fabricated flat, considerably reducing the cost and complexity of manufacturing curved 3D surfaces. While there are many ways to design self-actuated materials (e.g. using heat or water as actuation mechanisms), we use 3D printing to embed rigid patterns into prestressed fabric, which is then released and assumes a shape matching a given target when reaching static equilibrium.

While using a 3D printer to embed plastic curves into prestressed fabric is a technique that has been experimented on before, it has been mostly restricted to piecewise minimal surfaces, making it impossible to reproduce most shapes. By using a dense packing of 3-pointed stars, we are able to create convex shapes and positive gaussian curvature, moreover we found a direct link between the stars dimensions and the induced curvature, allowing us to build an inverse design tool that can faithfully reproduce some target shapes.

This is a collaboration with Mélina Skouras of Inria Rhône Alpes and Etienne Vouga of the University of Texas at Austin.

6.2. Graphics with Uncertainty and Heterogeneous Content

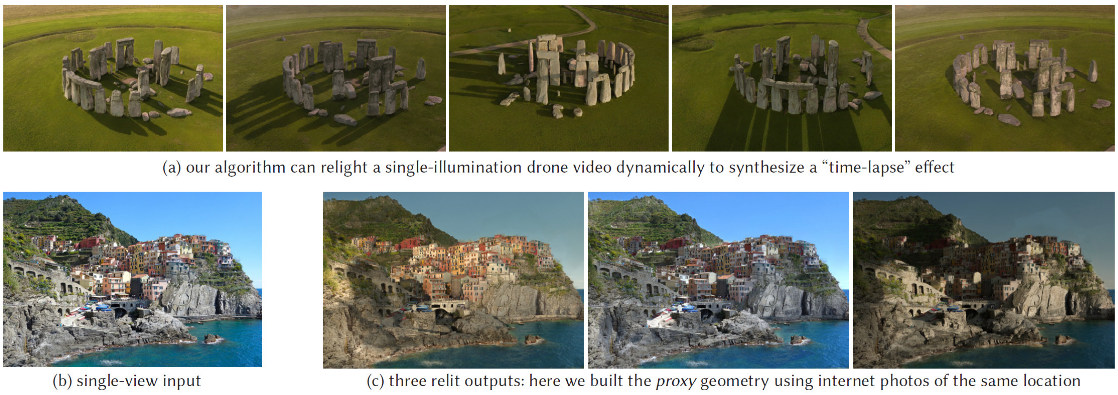

Figure 9. Results of our method: multi-view relighting using a geometry-aware network.

6.2.1. Multi-view relighting using a geometry-aware network

Participants: Julien Philip, George Drettakis.

We propose the first learning-based algorithm that can relight images in a plausible and controllable manner given multiple views of an outdoor scene. In particular, we introduce a geometry-aware neural network that utilizes multiple geometry cues (normal maps, specular direction, etc.) and source and target shadow masks computed from a noisy proxy geometry obtained by multi-view stereo. Our model is a three-stage pipeline: two subnetworks refine the source and target shadow masks, and a third performs the final relighting. Furthermore, we introduce a novel representation for the shadow masks, which we call RGB shadow images. They reproject the colors from all views into the shadowed pixels and enable our network to cope with inacuraccies in the proxy and the non-locality of the shadow casting interactions. Acquiring large-scale multi-view relighting datasets for real scenes is challenging, so we train our network on photorealistic synthetic data. At train time, we also compute a noisy stereo-based geometric proxy, this time from the synthetic renderings. This allows us to bridge the gap between the real and synthetic domains. Our model generalizes well to real scenes. It can alter the illumination of drone footage, image-based renderings, textured mesh reconstructions, and even internet photo collections (see Fig. 9).

This work was in collaboration with M. Gharbi of Adobe Research and A. Efros and T. Zhang of UC Berkeley, and was published in ACM Transactions on Graphics and presented at SIGGRAPH 2019 [19].

6.2.2. Flexible SVBRDF Capture with a Multi-Image Deep Network

Participants: Valentin Deschaintre, Frédo Durand, George Drettakis, Adrien Bousseau.

Empowered by deep learning, recent methods for material capture can estimate a spatially-varying reflectance from a single photograph. Such lightweight capture is in stark contrast with the tens or hundreds of pictures required by traditional optimization-based approaches. However, a single image is often simply not enough to observe the rich appearance of real-world materials. We present a deep-learning method capable of estimating material appearance from a variable number of uncalibrated and unordered pictures captured with a handheld camera and flash. Thanks to an order-independent fusing layer, this architecture extracts the most useful information from each picture, while benefiting from strong priors learned from data. The method can handle both view and light direction variation without calibration. We show how our method improves its prediction with the number of input pictures, and reaches high quality reconstructions with as little as 1 to 10 images – a sweet spot between existing single-image and complex multi-image approaches – see Fig. 10.

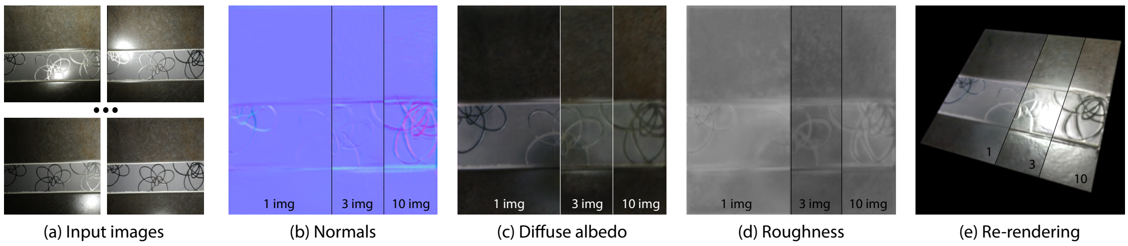

Figure 10. Our deep learning method for SVBRDF capture supports a variable number of input photographs taken with uncalibrated light-view directions (a, rectified). While a single image is enough to obtain a first plausible estimate of the SVBRDF maps, more images provide new cues to our method, improving its prediction. In this example, adding images reveals fine normal variations (b), removes highlight residuals in the diffuse albedo (c), and reveals the difference of roughness between the stone, the stripe, and the thin pattern (d).

This work is a collaboration with Miika Aittala from MIT CSAIL. This work was published in Computer Graphics Forum, and presented at EGSR 2019 [15].

A short paper and poster summarizing this work together with our 2018 "Single-Image SVBRDF Capture with a Rendering-Aware Deep Network" was published in the Siggraph Asia doctoral consortium 2019 [22].

6.2.3. Guided Acquisition of SVBRDFs

Participants: Valentin Deschaintre, George Drettakis, Adrien Bousseau.

Another project is under development to capture a large-scale SVBRDF from a few pictures of a planar surface. Many existing lightweight methods for SVBRDF capture take as input flash pictures, which need to be acquired close to the surface of interest restricting the scale of capture. We complement such small-scale inputs with a picture of the entire surface, taken under ambient lighting. Our method then fuses these two sources of information to propagate the SVBRDFs estimated from each close-up flash picture to all pixels of the large image. Thanks to our two-scale approach, we can capture surfaces several meters wide, such as walls, doors and furniture. In addition, our method can also be used to create large SVBRDFs from internet pictures, where we use artist-designed SVBRDFs as exemplars of the small-scale behavior of the surface.

6.2.4. Mixed rendering and relighting for indoor scenes

Participants: Julien Philip, Michaël Gharbi, George Drettakis.

We are investigating a mixed image rendering and relighting method that allows a user to move freely in a multi-view interior scene while altering its lighting. Our method uses a deep convolutional network trained on synthetic photo-realistic images. We adapt classical path tracing techniques to approximate complex lighting effects such as color bleeding and reflections.

6.2.5. DiCE: Dichoptic Contrast Enhancement for VR and Stereo Displays

Participant: George Drettakis.

In stereoscopic displays, such as those used in VR/AR headsets, our eyes are presented with two different views. The disparity between the views is typically used to convey depth cues, but it could be also used to enhance image appearance. We devise a novel technique that takes advantage of binocular fusion to boost perceived local contrast and visual quality of images. Since the technique is based on fixed tone curves, it has negligible computational cost and it is well suited for real-time applications, such as VR rendering. To control the trade-of between contrast gain and binocular rivalry, we conducted a series of experiments to explain the factors that dominate rivalry perception in a dichoptic presentation where two images of different contrasts are displayed (see Fig. 11). With this new finding, we can effectively enhance contrast and control rivalry in mono- and stereoscopic images, and in VR rendering, as conirmed in validation experiments.

Figure 11. Comparison of standard stereo images and the images with enhanced perceived contrast using our DiCE method. They can be cross-fused with the assistance of the dots above the images. Notice the enhanced contrast in the shadows and highlights of the scene. The stereo images are from Big Buck Bunny by Blender Foundation.

This work was in collaboration with Durham University (G. Koulieris, past postdoc of the group), Cambridge (F. Zhong, R. Mantiuk), UC Berkeley (M. Banks) and ENS Renne (M. Chambe), and was published in ACM Transactions on Graphics and presented at SIGGRAPH Asia 2019 [20].

6.2.6. Compositing Real Scenes using a relighting Network

Participants: Baptiste Nicolet, Julien Philip, George Drettakis.

Image-Based Rendering (IBR) allows for fast rendering of photorealistic novel viewpoints of real-world scenes captured by photographs. While it facilitates the very tedious traditional content creation process, it lacks user control over the appearance of the scene. We propose a novel approach to create novel scenes from a composition of multiple IBR scenes. This method relies on the use of a relighting network, which we first use to match the lighting conditions of each scene, and then to synthesize shadows between scenes in the final composition. This work has been submitted for publication.

6.2.7. Image-based Rendering of Urban Scenes based on Semantic Information

Participants: Simon Rodriguez, Siddhant Prakash, George Drettakis.

Cityscapes exhibit many hard cases for image-based rendering techniques, such as reflective and transparent surfaces. Pre-existing information about the scene can be leveraged to tackle these difficult cases. By relying on semantic information, it is possible to address those regions with tailored algorithms to improve reconstruction and rendering. This project is a collaboration with Peter Hedman from University College of London. This work has been submitted for publication.

6.2.8. Synthetic Data for Image-based Rendering

Participants: Simon Rodriguez, Thomas Leimkühler, George Drettakis.

This project explores the potential of Image-based rendering techniques in the context of real-time rendering for synthetic scenes. Accurate information can be precomputed from the input synthetic scene and used at run-time to improve the quality of approximate global illumination effects while preserving performance. This project is a collaboration with Chris Wyman and Peter Shirley from NVIDIA Research.

6.2.9. Densified Surface Light Fields for Human Capture Video

Participants: Rada Deeb, George Drettakis.

In this project, we focus on video-based rendering for mid-scale platforms. Having a mid-scale platform introduces one important problem for image-based rendering techniques due to low angular resolution. This leads to unrealistic view-dependent effects. We propose to use the temporal domain in a multidimensional surface light field approach in order to enhance the angular resolution. In addition, our approach provides a compact representation essential to dealing with the large amount of data introduced by videos compared to image-based techniques. In addition, we evaluate the use of deep encoder-decoder networks to learn a more compact representation of our multidimensional surface light field. This work is in collaboration with Edmond Boyer, MORPHEO team, Inria Grenoble.

6.2.10. Deep Bayesian Image-based Rendering

Participants: Thomas Leimkühler, George Drettakis.

Deep learning has permeated the field of computer graphics and continues to be instrumental in producing state-of-the-art research results. In the context of image-based rendering, deep architectures are now routinely used for tasks such as blending weight prediction, view extrapolation, or re-lighting. Current algorithms, however, do not take into account the different sources of uncertainty arising from the several stages of the image-based rendering pipeline. In this project, we investigate the use of Bayesian deep learning models to estimate and exploit these uncertainties. We are interested in devising principled methods which combine the expressive power of modern deep learning with the well-groundedness of classical Bayesian models.

6.2.11. Path Guiding for Metropolis Light Transport

Participants: Stavros Diolatzis, George Drettakis.

Path guiding has been proven to be an effective way to achieve faster convergence in Monte Carlo renderings by learning the incident radiance field. However, current path guiding techniques could be beaten by unguided path tracing due to their overhead or inability to incorporate the BSDF distribution factor. In our work, we improve path guiding and Metropolis light transport algorithms with low overhead product sampling between the incoming radiance and BSDF values. We demonstrate that our method has better convergence compared to the previous state-of-the-art techniques. Moreover, combining path guiding with MLT solves the global exploration issues ensuring convergence to the stationary distribution.

This work is an ongoing collaboration with Wenzel Jakob from Ecole Polytechnique Fédérale de Lausanne and Adrien Gruson from McGill University.

6.2.12. Improved Image-Based Rendering with Uncontrolled Capture

Participants: Siddhant Prakash, George Drettakis.

Current state-of-the-art Image Based Rendering (IBR) algorithms, such as Deep Blending, use per-view geometry to render candidate views and machine learning to improve rendering of novel views. The casual capture process employed introduces in visible color artifacts during rendering due to automated camera settings, and incur significant computational overhead when using per-view meshes. We aim to find a global solution to harmonize color inconsistency across the entire set of images in a given dataset, and also improve the performance of IBR algorithms by limiting the use of more advanced techniques only to regions where they are required.

6.2.13. Practical video-based rendering of dynamic stationary environments

Participants: Théo Thonat, 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 and a tripod. We focus on stochastic time-dependent textures such as waves, flames or waterfalls. We have developed a video representation able to tackle the challenge of blending unsynchronized videos.

This work is a collaboration with Sylvain Paris from Adobe Research, Miika Aittala from MIT CSAIL, and Yagiz Aksoy from ETH Zurich, and has been submitted for publication.

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.

7.2. Bilateral Grants with Industry

- As part of a long standing collaboration with Adobe, this year Julien Philip interned with Michael Gharbi (San Francisco). This follows previous internships of J. Delanoy with Aaron Hertzmann (San Francisco) and Theo Thonnat with Sylvain Paris (Boston),
- 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. EpicNPoc

Participants: Bastien Wailly, Adrien Bousseau.

EpicNPoc is a startup working on user interface design for the car industry. Together with two InriaTech engineers, we developed a small proof-of-concept that adapts our drawing recognition technology [9] to their needs. We first adapted our drawing synthesis algorithms to generate artificial sketches of user interface widgets, which include typical distortions and inaccuracies present in real sketches. The two engineers from InriaTech then used this technology to generate a large dataset of drawings, and to train a deep neural network to recognize the widgets in real drawings. The two engineers also integrated the trained network into a real-time system that recognizes widgets as they are drawn on a white board. We advised the engineers in their choice of a deep network architecture and on how to train this network to work on drawings. The result of this collaboration helped EpicNPoc appreciate the robustness of this technology, as well as to evaluate remaining challenges, such as convert the recognized widgets into working user-interface source code.

8.2. European Initiatives

8.2.1. FP7 & H2020 Projects

8.2.1.1. D^3: Drawing Interpretation for 3D Design

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

Line drawing is a fundamental tool for designers to quickly visualize 3D concepts. The goal of this ERC project is to develop algorithms capable of understanding design drawings. The first 30 months of the project allowed us to make significant progress in our understanding of how designers draw, and to propose preliminary solutions to the challenge of reconstructing 3D shapes from design drawings.

To better understand design sketching, we have collected a dataset of more than 400 professional design sketches [17]. We manually labeled the drawing techniques used in each sketch, and we registered all sketches to reference 3D models. Analyzing this data revealed systematic strategies employed by designers to convey 3D shapes, which will inspire the development of novel algorithms for drawing interpretation. In addition, our annotated sketches and associated 3D models form a challenging benchmark to test existing methods.

We proposed several methods to recover 3D information from drawings. A first family of method employs deep learning to predict what 3D shape is represented in a drawing. We applied this strategy in the context of architectural design, where we reconstruct 3D building by recognizing their constituent components (building mass, facade, window). We also presented an interactive system that allows users to create 3D objects by drawing from multiple viewpoints [14]. The second family of methods leverages geometric properties of the lines drawn to optimize the 3D reconstruction. In particular, we exploited properties of developable surfaces to reconstruct sketches of fashion items.

A long-term goal of our research is to evaluate the physical validity of a concept directly from a drawing. We obtained promising results towards this goal for the particular case of mechanical objects. We proposed an interactive system where users design the shape and motion of an articulated object, and our method automatically synthesize a mechanism that animates the object while avoiding collisions [18]. The geometry synthesized by our method is ready to be fabricated for rapid prototyping.

8.2.1.2. ERC FunGraph

Participants: George Drettakis, Thomas Leimkühler, Sébastien Morgenthaler, Rada Deeb, Stavros Diolatzis, Siddhant Prakash, Simon Rodriguez, Julien Philip.

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 did this by providing the means to authors of cultural products to create high-quality, interactive, personalized digital stories. The project was evaluated in December with very positive initial feedback.

GRAPHDECO contributed by developing novel image-based rendering techniques to help museum curators and archeologists provide more engaging experiences. We developed a mixed reality plugin for Unity that allows the use of IBR and we developed, in collaboration with ATHENA, a VR experience used in one of the EMOTIVE user experiences using a VIVE HMD. This demo was presented at a public event in November in Glasgow, and used by over 25 museum professionals with very positive feedback.

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

- McGill (Canada) (A. Gruson)
- UBC (Canada), (A. Sheffer)
- TU Delft (NL) (M. Sypesteyn, J. W. Hoftijzer and S. Pont)
- EPFL (Switzerland) (W. Jakob)
- U Bern (Switzerland) (D. Bommes)
- University College London (UK) (G. Brostow, P. Hedman)
- NVIDIA Research (USA, Finland), (C. Wyman, P. Shirley, M. Aittala)
- Adobe Research (USA), (A. Hertzmann, S. Paris, M. Gharbi)
- UC Berkeley (USA) (A. Efros)
- Purdue University (USA) (D. Aliaga, G. Nishida)
- U Texas, Austin (USA), (E. Vouga)
- George Mason University (USA) (Y. Gingold)

8.3.1.2. Inria International Chairs

Fredo Durand, Massachusetts Institute of Technology (United States)

Duration: 2016 - 2020

8.4. International Research Visitors

8.4.1. Visits of International Scientists

- Justin Solomon (MIT) in March.
- Mikhail Bessmeltsev (University of Montreal) in June.
- Aaron Hertzmann (Adobe Research) in June.
- Pierre Benard (U. Bordeaux), Daniel Sykora (U. Prague) and TT Wong (Hong Kong Polytechnic) in June.
- Tobias Ritschel (MPI Saarbrucken), Hendrik Lensch (U. Tuebingen) and Yann Gousseau (Telecom Paris) in June.
- Peter Hedman (UCL), September and October.
- Guillaume Coordonnier (ETH Zurich) in October.
- Alyosha Efros (U. Berkeley)
- Holly Rushmeier (Yale) and Abhijeet Ghosh (Imperial College London) in November.
- Niloy Mitra (UCL), Adrien Gruson (McGill) and Michael Gharbi (Adobe) in November.

8.4.1.1. Internships

J. Philip at Adobe Research, June 1st- September 28th, 2019. San Francisco.

8.4.2. Visits to International Teams

8.4.2.1. Research Stays Abroad

T. Stanko spent two weeks at University of Montreal to collaborate with Mikhail Bessmeltsev, and S. Rodriguez spent 5 weeks at NVIDIA research in Seattle (host C. Wyman).

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Scientific Events: Organisation

9.1.1.1. Member of the Organizing Committees

V. Deschaintre is the Web chair of EGSR 2020.

9.1.2. Scientific Events: Selection

9.1.2.1. Member of the Conference Program Committees (PC)

Y. Gryaditskaya was a PC member of Computational Visual Media Conference (CVM) 2020. A. Bousseau was a PC member for SIGGRAPH '19, Eurographics'19 and SMI'19. G. Drettakis was a member of the PC of Pacific Graphics'19 and EGSR'19 and participated in the papers sort for SIGGRAPH'19 in January in Zurich.

9.1.3. Journal

9.1.3.1. Member of the Editorial Boards

- George Drettakis is an Associate Editor of Computational Visual Media (CVM).

9.1.3.2. Reviewer - Reviewing Activities

Y. Gryaditskaya was a reviewer for IEEE Transactions on Image Processing (TIP), T. Stanko for Computer Aided Geometric Design journal, V. Deschaintre for SIGGRAPH Asia, Eurographics papers and STARs, R. Deeb for JOSA A, Optics express and CVIU and T. Thonat for Computer and Graphics.

9.1.4. Invited Talks

- Y. Gryaditskaya and V. Deschaintre gave invited talks at Ecole Polytechnique in July 2019.
- V. Deschaintre gave an invited talk at the Materials and Appearance Modelling workshop in Strasbourg in July.
- J. Philip and G. Drettakis gave an invited talk on multi-view relighting using a geometry-aware network at NASA Jet Propulsion Laboratory, Pasadena in July 28th, 2019.
- A. Bousseau gave a Thinkshell Architectural Geometry Lesson, Navier Laboratory (Paris, France), "Interpreting Drawings for 3D Design".
- G. Drettakis presented at the French Academy of Sciences days at Sophia-Antipolis (June 21st, 2019) on the topic "Computer Graphics and Machine Learning". Video on the Academy of Sciences site.
- G. Drettakis presented IBR research at the following companies: Mikros Image in Paris in September and Airbus Systems in Sophia in October.

9.1.5. Leadership within the Scientific Community

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

9.1.6. Scientific Expertise

G. Drettakis was an evaluator for the French ANR, ERC consolidator grants and Swiss National Research foundation, and is a member of the jury of the Ph.D. thesis award of the IG-RV (https://prixigrv2018.sciencesconf.org/). A. Bousseau was an evaluator for the IdEX University of Strasbourg (postdoc and PhD fellowships) and a member of the Eurographics Ph.D. award committee.

9.1.7. Research Administration

Adrien Bousseau is a member of "comité du centre" and "comité du suivi doctoral".

9.1.8. Interventions at Conferences

- J. Delanoy presented her work at the 8th ACM/EG Expressive Symposium, 5-6 May 2019, Genoa, Italy, and at SMI 2019 (Geometry Summit), 17-21 June 2019, Vancouver, Canada.
- F. Hähnlein presented a talk about "Data-driven sketch segmentation" project at JFIGRV2019 in Marseille, and also attended NeuroSTIC2019 and the UCA Deep Learning Summer School 2019.
- J. Philip presented his paper "Multi-view relighting using a geometry network" at SIGGRAPH 2019 in Los Angeles in August.
- V. Deschaintre presented his paper "Flexible SVBRDF Capture with a Multi-Image Deep Network" at EGSR 2019 in Strasbourg in July.
- Y. Gryaditskaya presented her paper "OpenSketch: A Richly-Annotated Dataset of Product Design Sketches" at SIGGRAPH Asia 2019 in Brisbane in November.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

Licence: T. Stanko Fondements mathématiques 1 (L1), 40h eq. TD, Université Côte d'Azur (France).

Master: G. Drettakis, A. Bousseau, Data visualization (M1), 12h eq. TD. Université Côte d'Azur (France).

9.2.2. Supervision

Ph.D.: Théo Thonat, Image-based rendering of thin and stochastic structures, defended June 2019, George Drettakis [13].

Ph.D.: Johanna Delanoy, Data-driven sketch-based modeling, defended June 2019, Adrien Bousseau [11].

Ph.D.: Valentin Deschaintre, Data-driven material capture, defended November 2019, Adrien Bousseau and George Drettakis [12].

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

Ph.D. in progress: Felix Hahnlein, Line Drawing Generation and Interpretation, since February 2019, Adrien Bousseau.

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

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

Ph.D. in progress: Stavros Diolatzis, Guiding and Learning for Illumination Algorithms, since April 2019, George Drettakis.

9.2.3. Juries

G. Drettakis was a member of the Ph.D. jury J-P. Bauchet (UCA) and Ph.D. reviewer of T. Leimkhueler (U. Saarbruecken). A. Bousseau was Ph.D. Reviewer for Marek Dvoroznak (TU Prague), Li Changjian (The University of Hong Kong) and Geoffrey Guingo (Grenoble University).

9.3. Popularization

9.3.1. Internal or external Inria responsibilities

- George Drettakis chairs the local “Jacques Morgenstern” Colloquium organizing committee and was elected a member of the Inria Scientific Council.

9.3.2. Articles and contents

J. Philip participated in the Youtube video interview: Intelligences artificielles: créatives mais perverties ? - Macroscopie La chaine - Video link - July 2019.

The work of J. Philip was the topic of a press release: L’algorithme d’apprentissage profond qui permet de changer l’éclairage des photos et de vidéos. Link to inria Press release and other media coverage Link to Pix Fan press release

The work Multi-view relighting using a geometry-aware network was presented at Adobe Max - November 2019. The project was part of the 11 projects selected out of 200 to be presented at Adobe MAX under the name project #LightRightSneak Link to the event.

Some press coverage of this event regarding the project : Link to Tech cafe listen starting at 32’00". Link to slrlouge. Link to petapixel.

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, J. PHILIP, T. PRICE, J.-M. FRAHM, G. DRETTAKIS, G. BROSTOW. *Deep Blending for Free-Viewpoint Image-Based Rendering*, in "ACM Transactions on Graphics (SIGGRAPH Asia Conference Proceedings)", November 2018, vol. 37, n^o 6, http://www-sop.inria.fr/reves/Basilic/2018/HPPFDB18

[6] 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

[7] 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

[8] 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

[9] 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

[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

Doctoral Dissertations and Habilitation Theses

[11] J. DELANOY. *Interpreting and Generating Artistic Depictions: Applications to Sketch-Based Modeling and Video Stylization*, Université Côte d'Azur, June 2019, https://hal.inria.fr/tel-02284941

[12] V. DESCHAINTRE. *Lightweight material acquisition using deep learning*, Université côte d'azur ; Inria - Sophia Antipolis, November 2019, https://hal.inria.fr/tel-02418445

[13] T. THONAT. *Multi-view inpainting, segmentation and video blending, for more versatile Image Based Rendering*, Université Côte d'Azur, June 2019, https://hal.inria.fr/tel-02417599

Articles in International Peer-Reviewed Journals

[14] J. DELANOY, D. COEURJOLLY, J.-O. LACHAUD, A. BOUSSEAU. *Combining Voxel and Normal Predictions for Multi-View 3D Sketching*, in "Computers and Graphics", June 2019, vol. 82, pp. 65–72 [*DOI :* 10.1016/J.CAG.2019.05.024], https://hal.archives-ouvertes.fr/hal-02141469

[15] V. DESCHAINTRE, M. AITTALA, F. DURAND, G. DRETTAKIS, A. BOUSSEAU. *Flexible SVBRDF Capture with a Multi-Image Deep Network*, in "Computer Graphics Forum", July 2019, vol. 38 [*DOI :* 10.1111/CGF.13765], https://hal.archives-ouvertes.fr/hal-02164993

[16] J.-D. FAVREAU, F. LAFARGE, A. BOUSSEAU, A. AUVOLAT. *Extracting Geometric Structures in Images with Delaunay Point Processes*, in "IEEE Transactions on Pattern Analysis and Machine Intelligence", 2019 [*DOI :* 10.1109/TPAMI.2018.2890586], https://hal.inria.fr/hal-01950791

[17] Y. GRYADITSKAYA, M. SYPESTEYN, J. W. HOFTIJZER, S. PONT, F. DURAND, A. BOUSSEAU. *OpenSketch: A Richly-Annotated Dataset of Product Design Sketches*, in "ACM Transactions on Graphics", 2019 [*DOI :* 10.1145/3355089.3356533], https://hal.inria.fr/hal-02284134

[18] G. NISHIDA, A. BOUSSEAU, D. G. ALIAGA. *Multi-Pose Interactive Linkage Design*, in "Computer Graphics Forum", 2019, vol. 38, n^o 2 [*DOI :* 10.1111/CGF.13637], https://hal.inria.fr/hal-02091234

[19] J. PHILIP, M. GHARBI, T. ZHOU, A. A. EFROS, G. DRETTAKIS. *Multi-view Relighting using a Geometry-Aware Network*, in "ACM Transactions on Graphics", 2019, vol. 38, forthcoming [*DOI :* 10.1145/3306346.3323013], https://hal.inria.fr/hal-02125095

[20] F. ZHONG, G.-A. KOULIERIS, G. DRETTAKIS, M. S. BANKS, M. CHAMBE, F. DURAND, R. MANTIUK. *DiCE: Dichoptic Contrast Enhancement for VR and Stereo Displays*, in "ACM Transactions on Graphics", November 2019, vol. 38, n^o 6 [*DOI :* 10.1145/3355089.3356552], https://hal.inria.fr/hal-02281073

International Conferences with Proceedings

[21] J. DELANOY, A. BOUSSEAU, A. HERTZMANN. *Video Motion Stylization by 2D Rigidification*, in "Expressive 2019 - 8th ACM/ Eurographics Proceedings of the Symposium", Genoa, Italy, May 2019, https://hal.inria.fr/hal-02106047

Conferences without Proceedings

[22] V. DESCHAINTRE. *Material acquisition using deep learning*, in "SIGGRAPH 2019 - Asia Doctoral Consortium", Brisbane, Australia, November 2019, https://hal.archives-ouvertes.fr/hal-02388382

[23] B. WAILLY, A. BOUSSEAU. *Line rendering of 3D meshes for data-driven sketch-based modeling*, in "Journées Francaises d'Informatique Graphique et de Réalité virtuelle", Marseille, France, AFIG, October 2019, https://hal.inria.fr/hal-02391655