

RESEARCH CENTRE

Inria Center
at **Université Côte d'Azur**

2022

ACTIVITY REPORT

Project-Team
GRAPHDECO

**GRAPHics and DEsign with
hEterogeneous COntent**

DOMAIN

Perception, Cognition and Interaction

THEME

Interaction and visualization

Inria

Contents

Project-Team GRAPHDECO	1
1 Team members, visitors, external collaborators	3
2 Overall objectives	4
2.1 General Presentation	4
3 Research program	5
3.1 Introduction	5
3.2 Computer-Assisted Design with Heterogeneous Representations	5
3.3 Graphics with Uncertainty and Heterogeneous Content	6
4 Application domains	7
5 Social and environmental responsibility	8
5.1 Footprint of research activities	8
5.2 Impact of research results	8
6 Highlights of the year	9
6.1 Awards	9
7 New software and platforms	9
7.1 New software	9
7.1.1 SynDraw	9
7.1.2 DeepSketch	9
7.1.3 sibr-core	10
7.1.4 IndoorRelighting	10
7.1.5 HybridIBR	10
7.1.6 FreeStyleGAN	11
7.1.7 PBNR	11
7.1.8 CASSIE	11
7.1.9 fabsim	12
7.1.10 activeExploration	12
8 New results	12
8.1 Computer-Assisted Design with Heterogeneous Representations	12
8.1.1 Free2CAD: Parsing Freehand Drawings into CAD Commands	12
8.1.2 Piecewise-Smooth Surface Fitting onto Unstructured 3D Sketches	13
8.1.3 Video doodles: Hand-Drawn Animations over Videos with Scene-Aware Keyframing	13
8.1.4 Symmetry-driven 3D Reconstruction from Concept Sketches	14
8.1.5 CAD2Sketch: Generating Concept Sketches from CAD Sequences	14
8.2 Graphics with Uncertainty and Heterogeneous Content	16
8.2.1 Active Exploration for Neural Global Illumination of Variable Scenes	16
8.2.2 Neural Catacaustics for Novel-View Synthesis of Reflections	16
8.2.3 Neural Precomputed Radiance Transfert	17
8.2.4 MesoGAN: Generative Neural Reflectance Shells	18
8.2.5 Deep scene-scale material estimation from multi-view indoor captures	18
8.2.6 Latent Relighting of 3D-Aware Images with Unconstrained Cameras	19
8.2.7 Synthetic dataset generation using neural rendering	19
8.2.8 NerfShop: Editing Neural Radiance Fields	20
8.2.9 3D Gaussians for Real-Time Neural Field Rendering	20
8.3 Physical Simulation for Graphics	20
8.3.1 Deep Reconstruction of 3D Smoke Densities from Artist Sketches	20
8.3.2 Computational Design of Self-Actuated Surfaces by Printing Plastic Ribbons on Stretched Fabric	21

8.3.3	Simulation of printed-on-fabric assemblies	22
8.3.4	Interactive simulation of plume and pyroclastic volcanic ejections	22
8.3.5	Neural Green's function for Laplacian systems	23
8.3.6	Glacial Erosion	23
8.3.7	Interactive design of 2D car profiles with aerodynamic feedback	24
8.3.8	ModalNerf: Neural Modal Analysis and Synthesis for Free-Viewpoint Navigation in Dynamically Vibrating Scenes	24
9	Bilateral contracts and grants with industry	24
9.1	Bilateral Grants with Industry	24
10	Partnerships and cooperations	25
10.1	International initiatives	25
10.1.1	International Collaborations	25
10.1.2	National Collaborations	25
10.2	International research visitors	25
10.2.1	Visits of international scientists	25
10.2.2	Visits to international teams	26
10.2.3	Participation in other international events	26
10.3	European initiatives	26
10.3.1	H2020 projects	26
10.4	National initiatives	28
10.4.1	Visits of French scientists	28
10.4.2	ANR projects	28
11	Dissemination	28
11.1	Promoting scientific activities	28
11.1.1	Journal	29
11.1.2	Invited talks	29
11.1.3	Leadership within the scientific community	29
11.1.4	Scientific expertise	29
11.1.5	Research administration	29
11.2	Teaching - Supervision - Juries	29
11.2.1	Teaching	29
11.2.2	Supervision	29
11.2.3	Juries	30
11.3	Popularization	30
12	Scientific production	30
12.1	Major publications	30
12.2	Publications of the year	31

Project-Team GRAPHDECO

Creation of the Project-Team: 2015 July 01

Keywords

Computer sciences and digital sciences

- A3.1.4. – Uncertain data
- A3.1.10. – Heterogeneous data
- A3.4.1. – Supervised learning
- A3.4.2. – Unsupervised learning
- A3.4.3. – Reinforcement learning
- A3.4.4. – Optimization and learning
- A3.4.5. – Bayesian methods
- 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
 - A5.1.5. – Body-based interfaces
 - A5.1.8. – 3D User Interfaces
 - A5.1.9. – User and perceptual studies
- A5.2. – Data visualization
- 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
 - A5.5.4. – Animation
- A5.6. – Virtual reality, augmented reality
 - A5.6.1. – Virtual reality
 - A5.6.2. – Augmented reality
 - A5.6.3. – Avatar simulation and embodiment
- A5.9.1. – Sampling, acquisition
- A5.9.3. – Reconstruction, enhancement
- A6.1. – Methods in mathematical modeling
- A6.2. – Scientific computing, Numerical Analysis & Optimization

A6.3.5. – Uncertainty Quantification

A8.3. – Geometry, Topology

A9.2. – Machine learning

A9.3. – Signal analysis

Other research topics and application domains

B3.2. – Climate and meteorology

B3.3.1. – Earth and subsoil

B3.3.2. – Water: sea & ocean, lake & river

B3.3.3. – Nearshore

B3.4.1. – Natural risks

B5. – Industry of the future

B5.2. – Design and manufacturing

B5.5. – Materials

B5.7. – 3D printing

B5.8. – Learning and training

B8. – Smart Cities and Territories

B8.3. – Urbanism and urban planning

B9. – Society and Knowledge

B9.1.2. – Serious games

B9.2. – Art

B9.2.2. – Cinema, Television

B9.2.3. – Video games

B9.5.1. – Computer science

B9.5.2. – Mathematics

B9.5.3. – Physics

B9.5.5. – Mechanics

B9.6. – Humanities

B9.6.6. – Archeology, History

B9.8. – Reproducibility

B9.11.1. – Environmental risks

1 Team members, visitors, external collaborators

Research Scientists

- George Drettakis [Team leader, INRIA, Senior Researcher, HDR]
- Adrien Bousseau [INRIA, Senior Researcher, HDR]
- Guillaume Cordonnier [INRIA, Researcher]

Post-Doctoral Fellows

- Stavros Diolatzis [INRIA, from Sep 2022]
- Bernhard Kerbl [INRIA, from Sep 2022]
- Changjian Li [Inria, until Jul 2022]

PhD Students

- Stavros Diolatzis [INRIA, until Jun 2022]
- Felix Hähnlein [INRIA, until Nov 2022]
- David Jourdan [Inria, until Mar 2022]
- Georgios Kopanas [INRIA]
- Siddhant Prakash [INRIA]
- Nicolas Rosset [INRIA]
- Nicolás Violante [INRIA, from Oct 2022]
- Emilie Yu [INRIA]

Technical Staff

- Felix Hahnlein [INRIA, Engineer, from Nov 2022]

Interns and Apprentices

- Roberto Felluga [INRIA, from Mar 2022 until Aug 2022]
- Clement Jambon [LIX, from Mar 2022 until Aug 2022]
- Automne Petitjean [DI-ENS, from Feb 2022]
- Yohan Poirier-Ginter [INRIA, from Aug 2022 until Nov 2022]
- Nicolás Violante [INRIA, from May 2022 until Sep 2022]

Administrative Assistant

- Sophie Honnorat [INRIA]

External Collaborators

- Yulia Gryaditskaya [University of Sussex, from Nov 2022]
- Thomas Leimkuhler [Max-Planck Institut Informatik]

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. 2.1(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. 2.1(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.

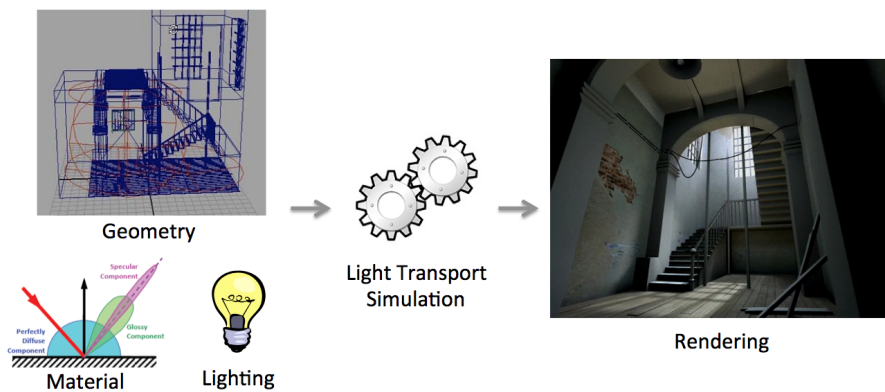


Figure 1: Traditional computer graphics pipeline. Rendering from www.thegnomonworkshop.com

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 un-**

certainty of a concept and its representations, from early sketches to 3D models and prototypes. Professional practice also needs to be adapted and facilitated to be accessible to all.

- Second, a new approach is required to **develop computer graphics models and algorithms capable of handling uncertain and heterogeneous data** as well as traditional synthetic content.

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

3 Research program

3.1 Introduction

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

3.2 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.

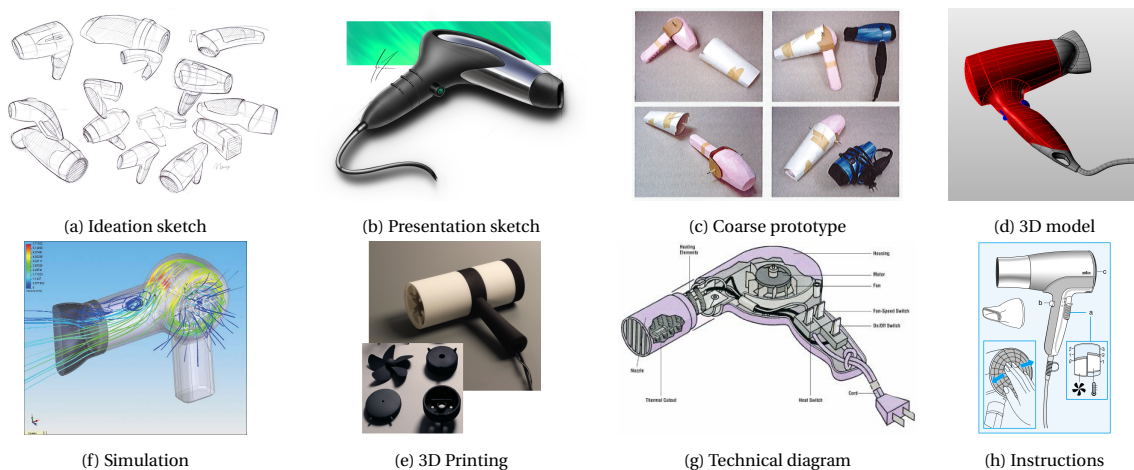


Figure 2: Various representations of a hair dryer at different stages of the design process. Image source, in order: c-maeng on deviantart.com, shauntur on deviantart.com, "Prototyping and Modelmaking for Product Design" Hallgrimsson, B., Laurence King Publishers, 2012, samsher511 on turbosquid.com,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

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

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

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

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

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

3.3 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 Application domains

Our research on design and computer graphics with heterogeneous data has the potential to change many different application domains. Such applications include:

Product design will be significantly accelerated and facilitated. Current industrial workflows separate 2D illustrators, 3D modelers and engineers who create physical prototypes, which results in a slow and complex process with frequent misunderstandings and corrective iterations between different people and different media. Our unified approach based on design principles could allow all processes to be done within a single framework, avoiding unnecessary iterations. This could significantly accelerate the design process (from months to weeks), result in much better communication between the different experts, or even create new types of experts who cross boundaries of disciplines today.

Mass customization will allow end customers to participate in the design of a product before buying it. In this context of “cloud-based design”, users of an e-commerce website will be provided with controls on the main variations of a product created by a professional designer. Intuitive modeling tools will also allow users to personalize the shape and appearance of the object while remaining within the bounds of the pre-defined design space.

Digital instructions for creating and repairing objects, in collaboration with other groups working in 3D fabrication, could have significant impact in sustainable development and allow anyone to be a creator of things, not just consumers, the motto of the makers movement.

Gaming experience individualization is an important emerging trend; using our results players will also be able to integrate personal objects or environments (e.g., their homes, neighborhoods) into any realistic 3D game. The success of creative games where the player constructs their world illustrates the potential of such solutions. This approach also applies to serious gaming, with applications in medicine, education/learning, training etc. Such interactive experiences with high-quality images of heterogeneous 3D content will be also applicable to archeology (e.g., realistic presentation of different reconstruction hypotheses), urban planning and renovation where new elements can be realistically used with captured imagery.

Virtual training, which today is restricted to pre-defined virtual environment(s) that are expensive and hard to create; with our solutions on-site data can be seamlessly and realistically used together with the actual virtual training environment. With our results, any real site can be captured, and the synthetic elements for the interventions rendered with high levels of realism, thus greatly enhancing the quality of the training experience.

Another interesting novel use of heterogeneous graphics could be for **news reports**. Using our interactive tool, a news reporter can take on-site footage, and combine it with 3D mapping data. The reporter can design the 3D presentation allowing the reader to zoom from a map or satellite imagery and better situate the geographic location of a news event. Subsequently, the reader will be able to zoom into a pre-existing street-level 3D online map to see the newly added footage presented in a highly realistic manner. A key aspect of these presentation is the ability of the reader to *interact* with the scene and the data, while maintaining a fully realistic and immersive experience. The realism of the presentation and the interactivity will greatly enhance the readers experience and improve comprehension of the news. The same advantages apply to enhanced **personal photography/videography**, resulting in much more engaging and lively memories. Such interactive experiences with high-quality images of heterogeneous 3D content will be also applicable to **archeology** (e.g., realistic presentation of different reconstruction hypotheses), **urban planning** and **renovation** where new elements can be realistically used with captured imagery.

Other applications may include scientific domains which use photogrammetric data (captured with various 3D scanners), such as geophysics and seismology. Note however that our goal is not to produce 3D data suitable for numerical simulations; our approaches can help in combining captured data with presentations and visualization of scientific information.

5 Social and environmental responsibility

5.1 Footprint of research activities

Deep learning algorithms use a significant amount of computing resources. We are attentive to this issue and plan to implement a more detailed policy for monitoring overall resource usage.

5.2 Impact of research results

G. Cordonnier collaborates with geologists and glaciologists on various projects, developing computationally efficient models that can have direct impact in climate related research. A. Bousseau regularly collaborates with designers; their needs serve as an inspiration for some of his research projects. Finally, the work in FUNGRAPH (G. Drettakis) has advanced research in visualization for reconstruction of real scenes that has potential application in planning and design.

6 Highlights of the year

6.1 Awards

This year marked the end of the D³ ERC Starting Grant. The end of the project resulted in numerous top publications, namely 3 ACM Transactions on Graphics papers and 1 SIGGRAPH Conference paper. In addition, the paper "Free2CAD: Parsing Freehand Drawings into CAD Commands" won an Honorable Mention best paper award at our flagship ACM SIGGRAPH conference. Our paper "Neural Precomputed Radiance Transfer" won the Gunter Enderle Best Paper Award (Honorable Mention) at Eurographics 22.

Emilie Yu was selected as one of the 10 [Wigraph Rising Stars 2022](#), which will allow her to benefit from an international mentorship program for women in computer graphics.

7 New software and platforms

7.1 New software

7.1.1 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 Contributions: 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, Adele Saint-Denis

Contact: Bastien Wailly

7.1.2 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 Contributions: 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, Bastien Wailly

Contact: Adrien Bousseau

7.1.3 sibr-core

Name: 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. sibr-core has been released as open source software, as well as the code for several of our research papers, as well as papers from other authors for comparisons and benchmark purposes.

The corresponding gitlab is: https://gitlab.inria.fr/sibr/sibr_core

The full documentation is at: <https://sibr.gitlabpages.inria.fr>

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, George Drettakis, Mahdi Benadel

Contact: George Drettakis

7.1.4 IndoorRelighting

Name: Free-viewpoint indoor neural relighting from multi-view stereo

Keywords: 3D rendering, Lighting simulation

Scientific Description: Implementation of the paper Free-viewpoint indoor neural relighting from multi-view stereo (currently under review), as a module for sibr.

Functional Description: Implementation of the paper Free-viewpoint indoor neural relighting from multi-view stereo (currently under review), as a module for sibr.

URL: <http://sibr.gitlabpages.inria.fr>

Contact: George Drettakis

Participants: George Drettakis, Julien Philip, Sébastien Morgenthaler, Michaël Gharbi

Partner: Adobe

7.1.5 HybridIBR

Name: Hybrid Image-Based Rendering

Keyword: 3D

Scientific Description: Implementation of the paper "Hybrid Image-based Rendering for Free-view Synthesis", (<https://hal.inria.fr/hal-03212598>) based on the sibr-core library. Code available: https://gitlab.inria.fr/sibr/projects/hybrid_ibr

Functional Description: Implementation of the paper "Hybrid Image-based Rendering for Free-view Synthesis", (<https://hal.inria.fr/hal-03212598>) based on the sibr-core library. Code available: https://gitlab.inria.fr/sibr/projects/hybrid_ibr

URL: <https://repo-sam.inria.fr/fungraph/hybrid-ibr/>

Publication: <https://hal.inria.fr/hal-03212598>

Contact: George Drettakis

Participants: Siddhant Prakash, Thomas Leimkuhler, Simon Rodriguez, George Drettakis

7.1.6 FreeStyleGAN

Keywords: Graphics, Generative Models

Functional Description: This codebase contains all applications and tools to replicate and experiment with the ideas brought forward in the publication "FreeStyleGAN: Free-view Editable Portrait Rendering with the Camera Manifold", Leimkühler & Drettakis, ACM Transactions on Graphics (SIGGRAPH Asia) 2021.

URL: <https://repo-sam.inria.fr/fungraph/freestylegan/>

Publication: [hal-03342414](https://hal.inria.fr/hal-03342414)

Authors: Thomas Leimkuhler, George Drettakis

Contact: George Drettakis

7.1.7 PBNR

Name: Point-Based Neural Rendering

Keyword: 3D

Scientific Description: Implementation of the method Point-Based Neural Rendering (<https://repo-sam.inria.fr/fungraph/differentiable-multi-view/>). This code includes the training module for a given scene provided as a multi-view set of calibrated images and a multi-view stereo mesh, as well as an interactive viewer built as part of SIBR.

Functional Description: Implementation of the method Point-Based Neural Rendering (<https://repo-sam.inria.fr/fungraph/differentiable-multi-view/>)

URL: <https://repo-sam.inria.fr/fungraph/differentiable-multi-view/>

Contact: George Drettakis

7.1.8 CASSIE

Name: CASSIE: Curve and Surface Sketching in Immersive Environments

Keywords: Virtual reality, 3D modeling

Scientific Description: This project is composed of: - a VR sketching interface in Unity - an optimisation method to enforce intersection and beautification constraints on an input 3D stroke - a method to construct and maintain a curve network data structure while the sketch is created - a method to locally look for intended surface patches in the curve network

Functional Description: This system is a 3D conceptual modeling user interface in VR that leverages free-hand mid-air sketching, and a novel 3D optimization framework to create connected curve network armatures, predictively surfaced using patches. Implementation of the paper <https://hal.inria.fr/hal-03149000>

Authors: Emilie Yu, Rahul Arora, Tibor Stanko, Adrien Bousseau, Karan Singh, J. Andreas Bærentzen

Contact: Emilie Yu

7.1.9 fabsim

Keywords: 3D, Graphics, Simulation

Scientific Description: Implemented models include: Discrete Elastic Rods (both for individual rods and rod networks) Discrete Shells Saint-Venant Kirchhoff, neo-hookean and incompressible neo-hookean membrane energies Mass-spring system

Functional Description: Static simulation of slender structures (rods, shells), implements known models from computer graphics

Contact: David Jourdan

Participants: Melina Skouras, Etienne Vouga

Partners: Etienne Vouga, Mélina Skouras

7.1.10 activeExploration

Name: Active Exploration for Neural Global Illumination of Variable Scenes

Keywords: 3D, Graphics, Active Learning, Deep learning

Scientific Description: Implementation of the method for the publication "Active Exploration for Neural Global Illumination of Variable Scenes", see <https://repo-sam.inria.fr/fungraph/active-exploration/>

Functional Description: Implementation of the method for the publication "Active Exploration for Neural Global Illumination of Variable Scenes", see <https://repo-sam.inria.fr/fungraph/active-exploration/>

URL: <https://repo-sam.inria.fr/fungraph/active-exploration/>

Authors: Stavros Diolatzis, Julien Philip, George Drettakis

Contact: George Drettakis

Partner: Adobe

8 New results

8.1 Computer-Assisted Design with Heterogeneous Representations

8.1.1 Free2CAD: Parsing Freehand Drawings into CAD Commands

Participants: Adrien Bousseau, Changjian Li.

CAD modeling, despite being the industry-standard, remains restricted to usage by skilled practitioners due to two key barriers. First, the user must be able to mentally parse a final shape into a valid sequence of supported CAD commands; and second, the user must be sufficiently conversant with CAD software packages to be able to execute the corresponding CAD commands. As a step towards addressing both these challenges, we present Free2CAD wherein the user can simply sketch the final shape and our system parses the input strokes into a sequence of commands expressed in a simplified CAD language (Fig. 4). When executed, these commands reproduce the sketched object. Technically, we cast sketch-based CAD modeling as a sequence-to-sequence translation problem, for which we leverage the powerful Transformers neural network architecture. Given the sequence of pen strokes as input, we introduce the new task of grouping strokes that correspond to individual CAD operations. We combine stroke grouping with geometric fitting of the operation parameters, such that intermediate groups are geometrically corrected before being reused, as context, for subsequent steps in the sequence inference.

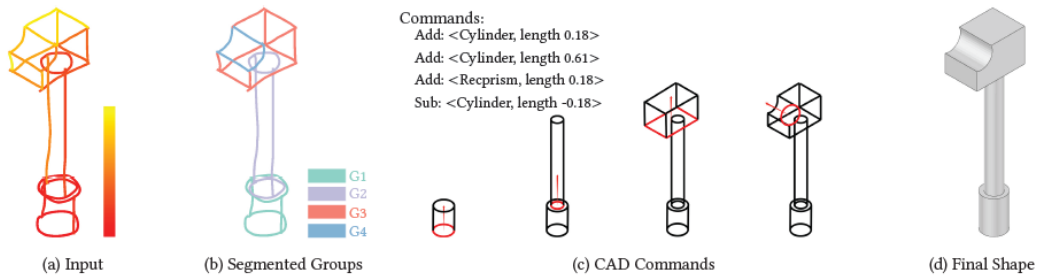


Figure 4: In this example, the user simply draws the final shape of the hammer and the anvil (a, strokes colored by drawing order as red to yellow). From this input, Free2CAD produces a corresponding set of CAD commands (c), thus freeing the user from having to mentally parse how to construct the shape with valid CAD operations. At the core of our system is a deep sequence-to-sequence neural network that automatically segments the input drawing into groups of strokes (b), each group representing a part of the shape that can be produced by a unique CAD command. Our system sequentially processes each group to recognize and fit the corresponding CAD commands, yielding a CAD program which, when executed, produces a 3D shape that is faithful to the input drawing (d).

Although trained on synthetically-generated data, we demonstrate that Free2CAD generalizes to sketches created from real-world CAD models as well as to sketches drawn by novice users.

This work is a collaboration with Niloy J. Mitra from UCL and Hao Pan from Microsoft Research Asia. It was published at ACM Transactions on Graphics [17], and presented at SIGGRAPH 2022 where it received an Honorable Mention Best Paper Award.

8.1.2 Piecewise-Smooth Surface Fitting onto Unstructured 3D Sketches

Participants: Emilie Yu, Adrien Bousseau.

We propose a method to transform unstructured 3D sketches into piecewise smooth surfaces that preserve sketched geometric features (Fig. 5d). Immersive 3D drawing and sketch-based 3D modeling applications increasingly produce imperfect and unstructured collections of 3D strokes as design output (Fig. 5a). These 3D sketches are readily perceived as piecewise smooth surfaces by viewers, but are poorly handled by existing 3D surface techniques tailored to well-connected curve networks or sparse point sets. Our algorithm is aligned with human tendency to imagine the strokes as a small set of simple smooth surfaces joined along stroke boundaries. Starting with an initial proxy surface (Fig. 5b), we iteratively segment the surface into smooth patches joined sharply along some strokes, and optimize these patches to fit surrounding strokes. Our evaluation is fourfold: we demonstrate the impact of various algorithmic parameters, we evaluate our method on synthetic sketches with known ground truth surfaces, we compare to prior art, and we show compelling results on more than 50 designs from a diverse set of 3D sketch sources.

This work is a collaboration with Rahul Arora and Karan Singh from University of Toronto, and J. Andreas Bærentzen from the Technical University of Denmark. It was published at ACM Transactions on Graphics, and presented at SIGGRAPH 2022 [21].

8.1.3 Video doodles: Hand-Drawn Animations over Videos with Scene-Aware Keyframing

Participants: Emilie Yu, Adrien Bousseau.

Video doodles are an emerging mixed media art, where an artist combines video content with hand-drawn animations to create unique and memorable video clips. The key challenge faced by artists who

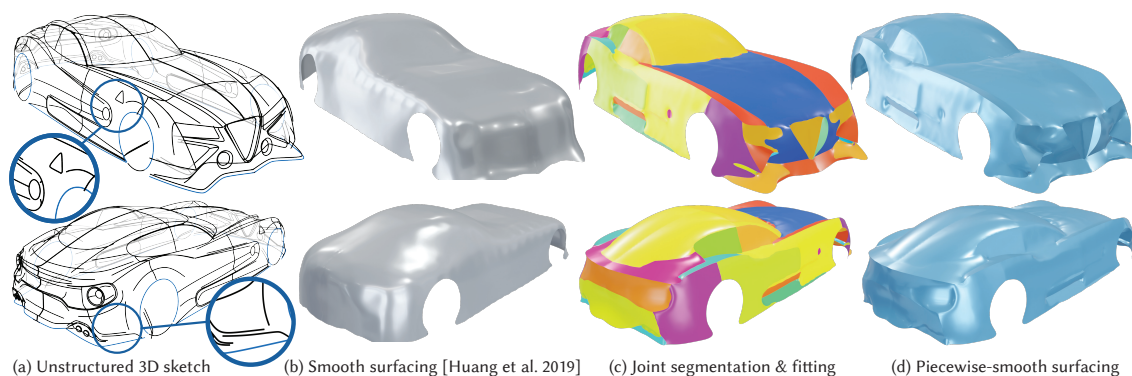


Figure 5: 3D sketches created in Virtual Reality (VR) or with sketch-based modeling systems often depict piecewise-smooth surfaces (a), but lack proper inter-stroke connectivity to detect the individual surface patches. State-of-the-art surfacing algorithms only produce smooth surfaces from such sparse and unstructured 3D data (b). Our algorithm segments such an initial smooth surface into regions aligned with the pen strokes to produce a piecewise-smooth surface that better captures the intended shape. ©James Robbins, used with permission.

create video doodles is to design the hand-drawn animation such that it interacts convincingly with the video content. We propose a novel scene-aware keyframe interpolation technique, that helps artists insert animated doodles into videos. Our contribution enables even novices to turn casually-captured videos into convincing *video doodles* for a variety of applications, ranging from fun posts on social media to engaging sport or craft tutorials.

This work is a collaboration with Kevin Blackburn-Matzen, Cuong Nguyen, Oliver Wang, Rubaiat Habib Kazi and Wilmot Li from Adobe Research San Francisco.

8.1.4 Symmetry-driven 3D Reconstruction from Concept Sketches

Participants: Felix Hähnlein, Yulia Gryaditskaya, Alla Sheffer, Adrien Bousseau.

Concept sketches, ubiquitously used in industrial design, are inherently imprecise yet highly effective at communicating 3D shape to human observers. In this project, we developed a new symmetry-driven algorithm for recovering designer-intended 3D geometry from concept sketches (Fig. 6). We observe that most concept sketches of human-made shapes are structured around locally symmetric building blocks, defined by triplets of orthogonal symmetry planes. We identify potential building blocks using a combination of 2D symmetries and drawing order. We reconstruct each such building block by leveraging a combination of perceptual cues and observations about designer drawing choices. We cast this reconstruction as an integer programming problem where we seek to identify, among the large set of candidate symmetry correspondences formed by approximate pen strokes, the subset that results in the most symmetric and well-connected shape. We demonstrate the robustness of our approach by reconstructing 82 sketches, which exhibit significant over-sketching, inaccurate perspective, partial symmetry, and other imperfections. In a comparative study, participants judged our results as superior to the state-of-the-art by a ratio of 2:1.

This work is a collaboration with Yulia Graditskaya from the Surrey Institute for People Centred AI and Centre for Vision, Speech and Signal Processing (CVSSP) and Alla Sheffer from the University of British Columbia (UBC). This work was published and presented at the ACM SIGGRAPH conferece 2022 [22].

8.1.5 CAD2Sketch: Generating Concept Sketches from CAD Sequences

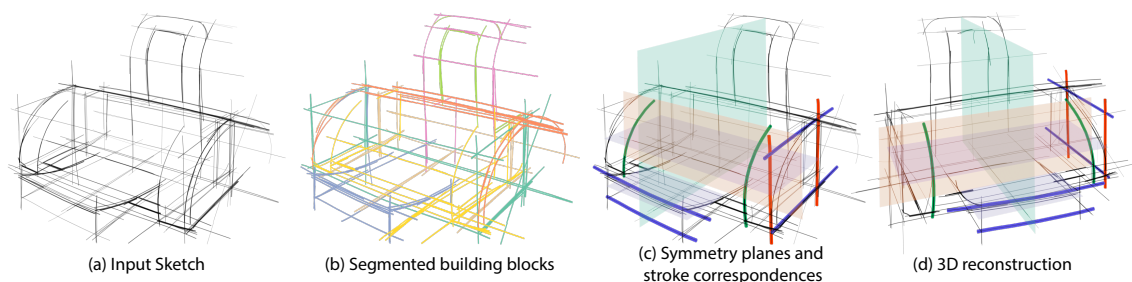


Figure 6: Concept sketches are dominated by symmetric strokes both in the shapes they depict, and in the construction lines employed to draw these shapes in perspective (a). Our algorithm first decomposes the sketch into locally-symmetric groups of strokes (b), and proceeds to identify pairs of strokes that are symmetric with respect to triplets of axis-aligned planes (c). At its core, our method selects the stroke correspondences that result in the most symmetric and well connected shape (d). We only show a subset of the symmetry planes and correspondences for illustration purpose.

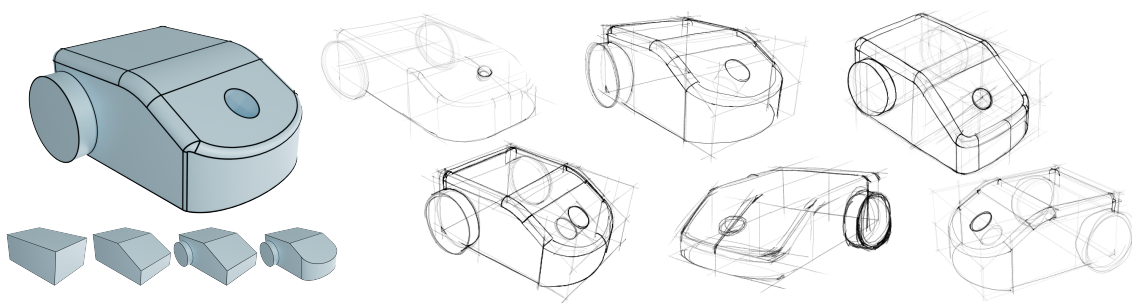


Figure 7: We developed CAD2Sketch to synthesize concept sketches directly from CAD sequences. Our work focuses on reducing the domain gap between freehand concept sketches and synthetically generated ones. We achieve this by focusing on generating intermediate lines (i.e., scaffold and auxiliary construction lines used by artists to reduce errors in ‘sighting’ perspective effects) and by recreating drawing style to reflect importance of lines. Here we show three real-world sketches and three synthetic sketches generated with our approach.

Participants: Felix Hähnlein, Changjian Li, Niloy Mitra, , Adrien Bousseau.

Concept sketches are ubiquitous in industrial design, as they allow designers to quickly depict imaginary 3D objects. To construct their sketches with accurate perspective, designers rely on longstanding drawing techniques, including the use of auxiliary construction lines to identify midpoints of perspective planes, to align points vertically and horizontally, and to project planar curves from one perspective plane to another. In this project, we developed a method to synthesize such construction lines from CAD sequences (Fig. 7). Importantly, our method balances the presence of construction lines with overall clutter, such that the resulting sketch is both well-constructed and readable, as professional designers are trained to do. In addition to generating sketches that are visually similar to real ones, we apply our method to synthesize a large quantity of paired sketches and normal maps, and show that the resulting dataset can be used to train a neural network to infer normals from concept sketches.

This work is a collaboration with Changjian Li from the University of Edinburgh and Niloy Mitra from the University of College London (UCL) and Adobe Research. This work was published in ACM Transactions on Graphics and presented at SIGGRAPH Asia 2022 [12].

8.2 Graphics with Uncertainty and Heterogeneous Content

8.2.1 Active Exploration for Neural Global Illumination of Variable Scenes

Participants: Stavros Diolatzis, George Drettakis.

Neural rendering algorithms introduce a fundamentally new approach for photorealistic rendering, typically by learning a neural representation of illumination on large numbers of ground truth images. When training for a given *variable* scene, i.e., changing objects (e.g. moving, scaling etc=, materials, lights and viewpoint, the space D of possible training data instances quickly becomes unmanageable as the dimensions of variable parameters increase. We introduce a novel *Active Exploration* method using Markov Chain Monte Carlo, which *explores* D , generating samples (i.e., ground truth renderings) that best help training and interleaves training and on-the-fly sample data generation. We introduce a self-tuning sample reuse strategy to minimize the expensive step of rendering training samples. We apply our approach on a neural generator that learns to render novel scene instances given an explicit parameterization of the scene configuration. Our results show that Active Exploration trains our network much more efficiently than uniformly sampling, and together with our resolution enhancement approach, achieves better quality than uniform sampling at convergence. Our method allows interactive rendering of hard light transport paths (e.g., complex caustics) – that require very high samples counts to be captured – and provides dynamic scene navigation and manipulation, after training for 5-18 hours depending on required quality and variations.

This work is a collaboration with Julien Philip from Adobe Research. It was published in ACM Transactions on Graphics and was presented at SIGGRAPH 2022 [11].




Figure 8: We introduce a neural rendering method that allows interactive navigation in a scene with dynamically changing properties, i.e., viewpoint, materials and geometry position and full global illumination effects. With our Active Exploration we can train a neural network efficiently to learn global illumination for all the configurations of these variable properties, allowing interactive rendering at runtime.

8.2.2 Neural Catacaustics for Novel-View Synthesis of Reflections

Participants: Georgios Kopanas, Clement Jambon, Thomas Leimkuhler, George Drettakis.

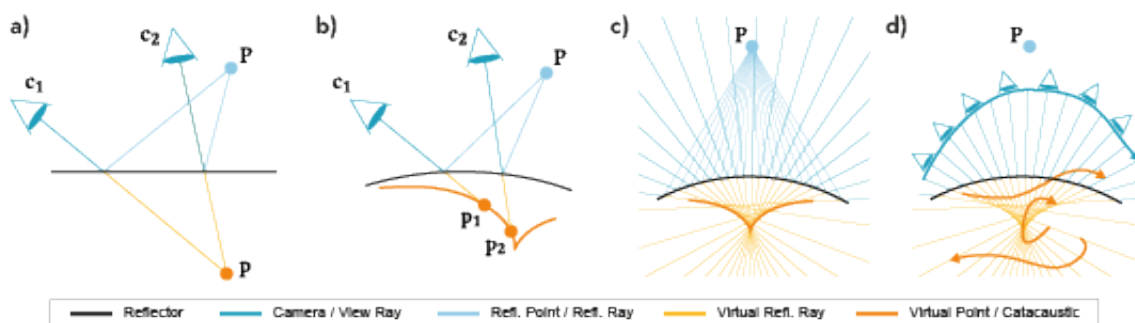


Figure 9: The geometry of catacaustics: (a) In the case of a planar reflector, a reflected point P results in a static virtual point p , independent of camera position c . (b) For curved reflectors (here a convex example), camera motion leads to p tracing a surface, called catacaustic. (c) The catacaustic for a single point P is defined by the envelope of virtual reflected rays, depicted as the bold orange curve, which at each point is tangent to one of the virtual reflected rays. Virtual images of P are formed only on the catacaustic trajectory. (d) In a point-based rendering setup, where the tangent constraint is not satisfied, the trajectory of virtual points is not unique. An infinite number of trajectories (three examples are shown) lead to the same apparent reflection. Reflection rays are omitted for clarity.

View-dependent effects such as reflections pose a substantial challenge for image-based and neural rendering algorithms. Above all, curved reflectors are particularly hard, as they lead to highly non-linear reflection flows as the camera moves.

We introduce a new point-based representation to compute Neural Point Catacaustics allowing novel-view synthesis of scenes with curved reflectors, from a set of casually-captured input photos. At the core of our method is a neural warp field that models catacaustic trajectories of reflections, so complex specular effects can be rendered using efficient point splatting in conjunction with a neural renderer. One of our key contributions is the explicit representation of reflections with a reflection point cloud which is displaced by the neural warp field, and a primary point cloud which is optimized to represent the rest of the scene. After a short manual annotation step, our approach allows interactive high-quality renderings of novel views with accurate reflection flow.

Additionally, the explicit representation of reflection flow supports several forms of scene manipulation in captured scenes, such as reflection editing, cloning of specular objects, reflection tracking across views, and comfortable stereo viewing.

This work was in collaboration with G. Rainer, previously a postdoc in our group. It was published in ACM Transactions on Graphics, and presented at SIGGRAPH Asia 2022 [15].

8.2.3 Neural Precomputed Radiance Transfer

Participants: Adrien Bousseau, George Drettakis.

Recent advances in neural rendering indicate immense promise for architectures that learn light transport, allowing efficient rendering of global illumination effects once such methods are trained. The training phase of these methods can be seen as a form of pre-computation, which has a long standing history in Computer Graphics. In particular, Pre-computed Radiance Transfer (PRT) achieves real-time rendering by freezing some variables of the scene (geometry, materials) and encoding the distribution of others, allowing interactive rendering at runtime. We adopt the same configuration as PRT for global illumination of static scenes under dynamic environment lighting and investigate different neural network

architectures (see Fig. 10), inspired by the design principles and theoretical analysis of PRT. We introduce four different architectures, and show that those based on knowledge of light transport models and PRT-inspired principles improve the quality of global illumination predictions at equal training time and network size, without the need for high-end ray-tracing hardware.

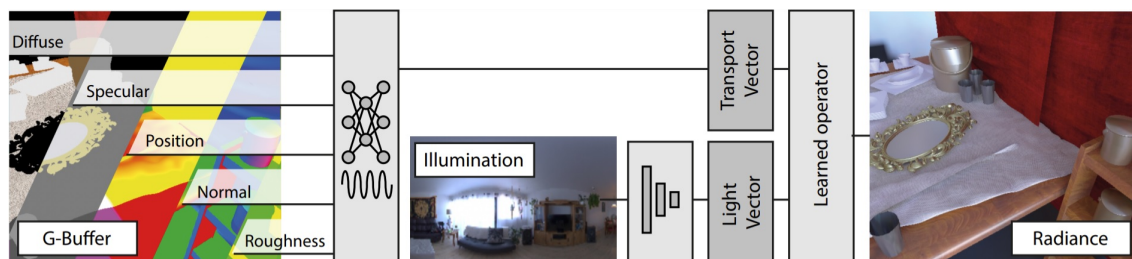


Figure 10: We investigate the design space of scene-specific neural renderers that take environment map lighting and G-buffer pixels as input. We show that architectures based on design principles from the PRT literature can greatly improve quality at no extra cost.

The main author was Gilles Rainer, previously a postdoc at GRAPHDECO. This work is a collaboration with Tobias Ritschel from University College London. This work was published in Computer Graphics Forum, and presented at Eurographics 2022 [19].

8.2.4 MesoGAN: Generative Neural Reflectance Shells

Participants: George Drettakis, Stavros Diolatzis.

We introduce MesoGAN, a model for generative 3D neural textures. This new graphics primitive represents mesoscale appearance by combining the strengths of generative adversarial networks (StyleGAN) and volumetric neural field rendering. The primitive can be applied to surfaces as a *neural reflectance shell*; a thin volumetric layer above the surface with appearance parameters defined by a neural network. To construct the neural shell, we first generate a 2D feature texture using StyleGAN with carefully randomized Fourier features to support arbitrarily sized textures without repeating artifacts. We augment the 2D feature texture with a learned height feature, which aids the neural field renderer in producing volumetric parameters from the 2D texture. To facilitate filtering, and to enable end-to-end training within memory constraints of current hardware, we utilize a hierarchical texturing approach and train our model on multi-scale synthetic datasets of 3D mesoscale structures. We propose one possible approach for conditioning MesoGAN on artistic parameters (e.g., fiber length, density of strands, lighting direction) and demonstrate and discuss integration into physically based renderers.

This project is a collaboration with Jonathan Granskog, Fabrice Rouselle and Jan Novak from Nvidia, and Ravi Ramamoorthi from the University of California, San Diego and is submitted for publication.

8.2.5 Deep scene-scale material estimation from multi-view indoor captures

Participants: Siddhant Prakash, Adrien Bousseau, George Drettakis.

The movie and video game industries have adopted photogrammetry as a way to create digital 3D assets from multiple photographs of a real-world scene. However, photogrammetry algorithms typically output an RGB texture atlas of the scene that only serves as visual guidance for skilled artists to create material maps suitable for physically-based rendering. We present a learning-based approach that automatically produces digital assets ready for physically-based rendering, by estimating approximate material maps from multi-view captures of indoor scenes that are used with retopologized geometry. We base our approach on a material estimation Convolutional Neural Network (CNN) that we execute on

each input image. We leverage the view-dependent visual cues provided by the multiple observations of the scene by gathering, for each pixel of a given image, the color of the corresponding point in other images. This image-space CNN provides us with an ensemble of predictions, which we merge in texture space as the last step of our approach. Our results demonstrate that the recovered assets can be directly used for physically-based rendering and editing of real indoor scenes from any viewpoint and novel lighting - see Fig. 11. Our method generates approximate material maps in a fraction of time compared to the closest previous solutions.

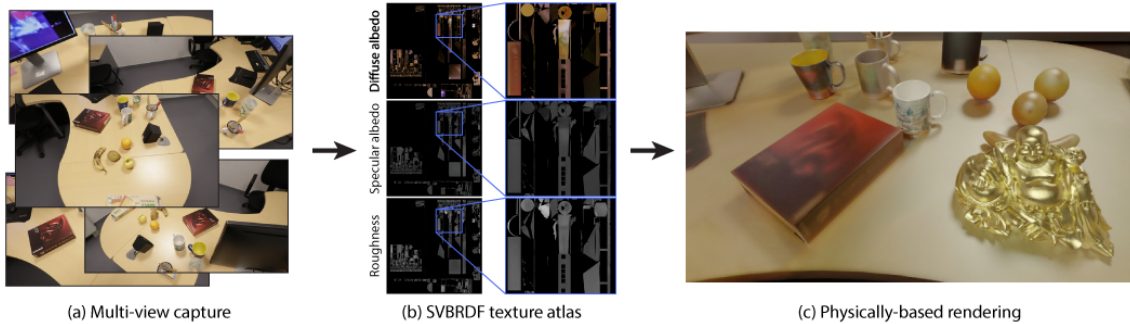


Figure 11: We present a method which takes multiple photographs of a scene as input (a) and predicts surface materials in the form of material maps corresponding to each input view. Merging these image-space predictions in texture space yields a texture atlas (b) that can be mapped onto retopologized geometry to produce *digital 3D assets* ready for full physically-based rendering, i.e., rendering from any viewpoint, changing/adding lights and objects, for example in (c) the golden statuette and the white mug have been added to the scene that has modified lighting and is rendered from a novel viewpoint not in the input.

This work was published in the journal *Computers & Graphics* in October 2022 [18] and was in collaboration with G. Rainer, previously a postdoc in our group.

8.2.6 Latent Relighting of 3D-Aware Images with Unconstrained Cameras

Participants: Nicolás Violante, Thomas Leimkuhler, Stavros Diolatzis, George Dretakis.

We propose a relightable 3D-aware image synthesis model that decouples the generation of an object of interest and an environment map, enabling free-viewpoint rendering with illumination control in latent space. Our model is trained end-to-end on 2D images with corresponding camera poses. We show the advantages of our method by training it on a challenging dataset of cars, which includes glossy surfaces, transparent windows, visible backgrounds, and unrestricted cameras. By modeling the environment map with a separate generator, we are able to perform relighting via latent space manipulations without modifying other important attributes of the scene. In addition, we show that our model can project and relight real images using both latent and real environment maps. This is work in progress.

8.2.7 Synthetic dataset generation using neural rendering

Participants: Siddhant Prakash, Stavros Diolatzis, Adrien Bousseau, George Dretakis.

Traditional rendering with path tracer is slow but is used to create large datasets for training neural networks making the data generation process highly time-consuming. Additionally, traditional datasets require substantial storage capacity and introduce transfer overhead making the use of large rendered datasets impractical. To speed up data generation and increase variability in data generated, we explore

neural rendering as an alternative to path tracing. Neural rendering generates images with full global illumination at a rate significantly faster than path tracing, making it highly suitable for on-the-fly data generation for training tasks. We study the pros and cons of using neural rendering for data generation highlighting different scenarios where neural rendering has clear advantage over path tracing. This is work in progress.

8.2.8 NerfShop: Editing Neural Radiance Fields

Participants: Clement Jambon, Bernhard Kerbl, Georgios Kopanas, Stavros Diolatzis, Thomas Leimkuhler, George Drettakis.

We present a method that allows interactive editing of Neural Radiance Fields. We build on the interactive display system of Instant Neural Graphics Primitives, which allows fluid user interaction. The user can interact with the scene with scribbles, that are then projected onto the neural field density lifting them to 3D. This defines a region that is enclosed in a “cage” mesh that is subsequently tetrahedralized. The tetrahedra can then be manipulated easily using rigid or non-rigid deformations, allowing interactive editing of reconstructed scenes with non-rigid deformations. This work is submitted for publication.

8.2.9 3D Gaussians for Real-Time Neural Field Rendering

Participants: Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuhler, George Drettakis.

We present the first method that allows real-time rendering of radiance fields. This is achieved with three components: first we represent the scene with 3D Gaussians, initialized with a sparse point cloud for Structure-from-Motion calibration a set of photos of a scene. Then we optimize the parameters of the Gaussians, interleaving optimization and Gaussian density control. Finally our fast CUDA-based rasterizer allows anisotropic splatting, full alpha-blending and fast backpropagation. Our results are better than all previous methods in image quality and as fast as the recent solutions such as Instant Neural Graphics Primitives for the same quality. This work is submitted for publication.

8.3 Physical Simulation for Graphics

8.3.1 Deep Reconstruction of 3D Smoke Densities from Artist Sketches

Participants: Guillaume Cordonnier.

Creative processes of artists often start with hand-drawn sketches illustrating an object. Pre-visualizing these keyframes is especially challenging when applied to volumetric materials such as smoke. The authored 3D density volumes must capture realistic flow details and turbulent structures, which is highly non-trivial and remains a manual and time-consuming process. We therefore present a method to compute a 3D smoke density field directly from 2D artist sketches (see Fig. 12), bridging the gap between early-stage prototyping of smoke keyframes and pre-visualization. From the sketch inputs, we compute an initial volume estimate and optimize the density iteratively with an updater CNN. Our differentiable sketcher is embedded into the end-to-end training, which results in robust reconstructions. Our training data set and sketch augmentation strategy are designed such that it enables general applicability. We evaluate the method on synthetic inputs and sketches from artists depicting both realistic smoke volumes and highly non-physical smoke shapes. The high computational performance and robustness of our method at test time allows interactive authoring sessions of volumetric density fields for rapid prototyping of ideas by novice users.



Figure 12: Our method reconstructs volumetric density fields from sketch inputs using an updater CNN. The efficient reconstruction at test time enables interactive authoring and editing of arbitrary smoke shapes.

This work is a collaboration with Xingchang Huang from the Max Planck Institute for Informatics, Byungsoo Kim, Laura Wuelfroth, Markus Gross, and Barbara Solenthaler from ETH Zurich. This work was published at Computer Graphics Forum and presented at Eurographics 2022 [14].

8.3.2 Computational Design of Self-Actuated Surfaces by Printing Plastic Ribbons on Stretched Fabric

Participants: David Jourdan, Adrien Bousseau.

We introduce a new mechanism for self-actuating deployable structures, based on printing a dense pattern of closely-spaced plastic ribbons on sheets of pre-stretched elastic fabric (Fig. 13). We leverage two shape-changing effects that occur when such an assembly is printed and allowed to relax: first, the incompressible plastic ribbons frustrate the contraction of the fabric back to its rest state, forcing residual strain in the fabric and creating intrinsic curvature. Second, the differential compression at the interface between the plastic and fabric layers yields a bilayer effect in the direction of the ribbons, making each ribbon buckle into an arc at equilibrium state and creating extrinsic curvature. We describe an inverse design tool to fabricate low-cost, lightweight prototypes of freeform surfaces using the controllable directional distortion and curvature offered by this mechanism. The core of our method is a parameterization algorithm that bounds surface distortions along and across principal curvature directions, along with a pattern synthesis algorithm that covers a surface with ribbons to match the target distortions and curvature given by the aforementioned parameterization. We demonstrate the flexibility and accuracy of our method by fabricating and measuring a variety of surfaces, including nearly-developable surfaces as well as surfaces with positive and negative mean curvature, which we achieve thanks to a simple hardware setup that allows printing on both sides of the fabric.

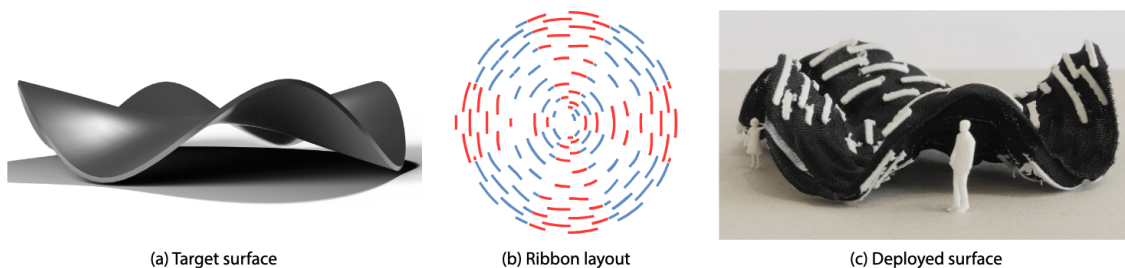


Figure 13: Given a target 3D surface (a), our method computes a flat layout of plastic ribbons (b) which, when printed on pre-stretched fabric, causes the surface to deploy into 3D when the fabric is released (c). We designed a custom frame to easily print ribbons on both sides of the fabric (b, red and blue), allowing us to reproduce surfaces with positive and negative extrinsic curvature.

This work is a collaboration with Méлина Skouras from Inria, Université Grenoble Alpes (Imagine team) and Etienne Vouga from University of Texas at Austin. It was published in Computer Graphics Forum, and presented at Eurographics 2022 [13].

8.3.3 Simulation of printed-on-fabric assemblies

Participants: David Jourdan, Adrien Bousseau.

Printing-on-fabric is an affordable and practical method for creating self-actuated deployable surfaces: thin strips of plastic are deposited on top of a pre-stretched piece of fabric using a commodity 3D printer; the structure, once released, morphs to a programmed 3D shape. Several physics-aware modeling tools have recently been proposed to help designing such surfaces. However, existing simulators do not capture well all the deformations these structures can exhibit. In this work, we propose a new model for simulating printed-on-fabric composites based on a tailored bilayer formulation for modeling plastic-on-top-of-fabric strips, and an extended Saint-Venant–Kirchhoff material law for modeling the surrounding stretchy fabric (Fig. 14). We show how to calibrate our model through a series of standard experiments. Finally, we demonstrate the improved accuracy of our simulator by conducting various tests.

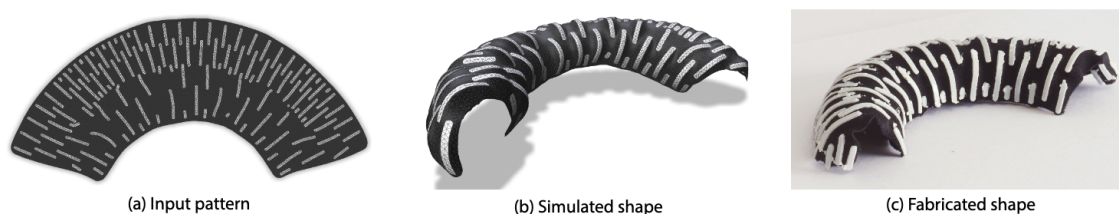


Figure 14: Printing-on-fabric consists in depositing strips of plastic over a pre-stretched piece of fabric. On release, the forces exerted by the fabric make the overall structure buckle into a 3D shape. Our simulator predicts the 3D shape adopted by a given pattern of plastic strips by accounting for the physical properties of the fabric and its interactions with the plastic layer. Starting with this radial pattern to be printed on pre-stretched fabric (a), our simulator produces a toric shape (b) that closely matches the fabricated result (c).

This work is a collaboration with Méлина Skouras and Victor Romero from Inria, Université Grenoble Alpes (Imagine and Elan teams) and Etienne Vouga from University of Texas at Austin. It was presented at the ACM Symposium on Computational Fabrication [23].

8.3.4 Interactive simulation of plume and pyroclastic volcanic ejections

Participants: Guillaume Cordonnier.

The goal of this project was to design an interactive animation method for the ejection of gas and ashes mixtures in volcano eruption (Fig. 15). Our novel, layered solution combines a coarse-grain, physically-based simulation of the ejection dynamics with a consistent, procedural animation of multi-resolution details. We show that this layered model can be used to capture the two main types of ejection, namely ascending plume columns composed of rapidly rising gas carrying ash which progressively entrains more air, and pyroclastic flows which descend the slopes of the volcano depositing ash, ultimately leading to smaller plumes along their way. We validate the large-scale consistency of our model through comparison with geoscience data, and discuss both real-time visualization and off-line, realistic rendering.

This work is a collaboration with Maud Lastic, Damien Rohmer, Marie-Paule Cani from Ecole Polytechnique, Chaude Jaupart from the Institut de Physique du Globe de Paris, and Fabrice Neyret from



Figure 15: Interactive simulation (left) and off-line rendering of ejection phenomena during volcanic eruption: Rising column under the wind (middle); Pyroclastic flow (right), where secondary columns are generated along the way.

CNRS and Université Grenoble Alpes. This work was presented at the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games 2022 [16].

8.3.5 Neural Green's function for Laplacian systems

Participants: Guillaume Cordonnier.

Solving linear systems of equations stemming from Laplacian operators is at the heart of a wide range of applications. Due to the sparsity of the linear systems, iterative solvers such as Conjugate Gradient and Multigrid are usually employed when the solution has a large number of degrees of freedom. These iterative solvers can be seen as sparse approximations of the Green's function for the Laplacian operator. The goal of this project is to develop a machine learning approach that regresses a Green's function from boundary conditions. This is enabled by a Green's function that can be effectively represented in a multi-scale fashion, drastically reducing the cost associated with a dense matrix representation. Additionally, since the Green's function is solely dependent on boundary conditions, training the proposed neural network does not require sampling the right-hand side of the linear system. We show results that our method outperforms state of the art Conjugate Gradient and Multigrid methods (Fig. 16).

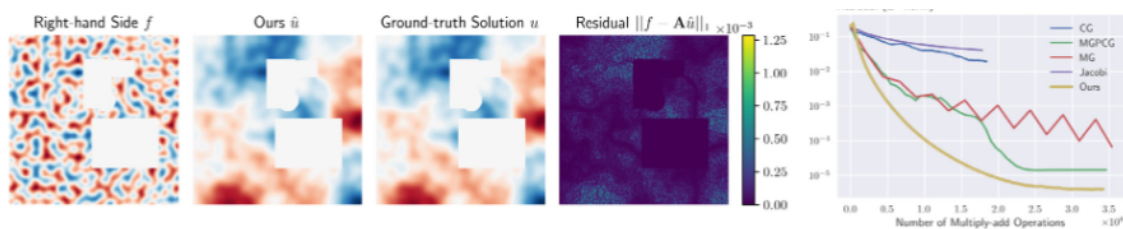


Figure 16: Solving the Poisson Equation using multi-level Green's function optimized on single scenes in a resolution of 257×257 . The interior objects shown in white all have Neumann boundaries. Our method outperforms all competing solvers when complex boundary conditions are set.

This work is a collaboration with Jingwei Tang and Barbara Solenthaler from ETH Zurich and Vinicius C. Azevedo From Disney Research. This work was published in Computer and Graphics [20].

8.3.6 Glacial Erosion

Participants: Guillaume Cordonnier.

Glacial erosion has a significant impact on the appearance of mountains and yet has not been considered in computer graphics. In this project, we tackle the numerous challenges of the modeling of

ice dynamics and the interleaved interactions with the abraded bedrock. Our goal is to propose a fast, physically accurate simulation method to generate glacier-carved landscapes.

This work is a collaboration with Guillaume Jouvet from the University of Lausanne, Jean Braun from GFZ and Potsdam University, Bedrich Benes from Purdue University, James Gain from Cape Town University, Marie-Paule Cani from Ecole Polytechnique, Eric Galin, Eric Guérin and Adrien Peytavie from LIRIS.

8.3.7 Interactive design of 2D car profiles with aerodynamic feedback

Participants: Nicolas Rosset, Adrien Bousseau, Guillaume Cordonnier.

In this project, we propose an interactive system to assist designers in creating aerodynamic car profiles. It relies on a neural surrogate model to predict fluid flow around car shapes, providing fluid visualization and shape optimization feedback to designers as soon as they sketch a car profile. Instead of focusing on time-averaged fluid flows, we describe how to train our model on instantaneous, synchronized observations extracted from multiple pre-computed simulations, such that we can visualize and optimize for dynamic flow features, such as vortices. We architected our model to support gradient-based shape optimization within a learned latent space of car profiles. Furthermore, we designed our model to support pointwise queries of fluid properties around car shapes, allowing us to adapt computational cost to application needs.

This work is in collaboration with Regis Duvigneau from the Acumes team of Inria Sophia Antipolis.

8.3.8 ModalNerf: Neural Modal Analysis and Synthesis for Free-Viewpoint Navigation in Dynamically Vibrating Scenes

Participants: Yohan Poirier-Ginter, Guillaume Cordonnier, George Drettakis.

ModalNerf is the first method that allows the capture of geometrical and physical properties of a scene using a single camera. Our method augments a radiance field by training a particle-based deformation field that allows easy physical analysis and synthesis of new motion using modal analysis. This is the first method that allows manipulation of the motion of a captured 3D scene while allowing free viewpoint navigation.

This work is a collaboration with Ayush Tewari from MIT and is submitted for publication.

9 Bilateral contracts and grants with industry

9.1 Bilateral Grants with Industry

Participants: George Drettakis, Adrien Bousseau.

We have received regular donations from Adobe Research, thanks to our collaborations with J. Philip and with N. Mitra.

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 International Collaborations

A. Bousseau collaborates regularly with A. Sheffer at the University of British Columbia, and N. Mitra from UCL (UK)/Adobe.

G. Drettakis collaborates regularly with J. Philip at Adobe Research London, and collaborated with NVIDIA Research (J. Novak, F. Rousselle, M. Aittala) as part of the followup from S. Diolatzis internship. Together with A. Bousseau they collaborated with T. Ritschel from UCL (UK). He also collaborates with the Université de Laval in Quebec with J-F Lalonde.

G. Cordonnier collaborates regularly with B. Benes from Perdue (US), J. Gain from Cape Town University (South Africa), J. Braun from GFZ and University of Potsdam (Germany).

E. Yu and Y. Poirier-Ginter both received MITACS Globalink Research Award: E. Yu will visit U. Toronto and for Y. Poirier-Ginter the award financed his visit to Inria.

10.1.2 National Collaborations

LIX

Participants: Guillaume Cordonnier.

G. Cordonnier collaborates with Marie-Paule Cani and Damien Rohmer of LIX, Ecole polytechnique.

LIRIS

Participants: Guillaume Cordonnier.

G. Cordonnier collaborates with Eric Galin, Eric Guérin and Adrien Peytavie of LIRIS.

EXSITU

Participants: Adrien Bousseau.

A. Bousseau collaborates with F. Tsandilas of EXSITU Inria project team as part of the co-supervision of the Ph.D. thesis of Capucine Nghien.

10.2 International research visitors

10.2.1 Visits of international scientists

Other international visits to the team For the Ph.D. defense of S. Diolatzis, T. Ritschel from UCL visited the team, and gave a presentation.

For the Ph.D. defense of F. Hahnlein, Marc Alexa (TU Berlin), Pedro Company (Univ. Jaume I) and Yulia Graditskaya (Univ. Surrey) visited the team and gave talks in a mini-workshop "Creating and Processing 3D Shapes".

For the Ph.D. defense of S. Prakash, Marc Stamminger (TU Erlangen) visited the team and gave a presentation in the mini-workshop "Recent Advances in Appearance Capture and Rendering".

James Gain (Cape Town University) visited the team in June and gave a presentation.

10.2.2 Visits to international teams

F. Hahnlein visited UC London (Niloy Mitra), University of Washington (Adriana Schulz), TU Delft and KAIST (South Korea, Seok-Hyung Bae).

Adrien Bousseau is on sabbatical at TU Delft (Netherlands) from Sep. 2022 to Jul. 2023.

E. Yu visited TU Delft and Technical University of Denmark (J. Andreas Bærentzen).

10.2.3 Participation in other international events

E. Yu participated in the Wigraph Rising Stars Workshop 2022, co-located with the SIGGRAPH conference.

10.3 European initiatives

10.3.1 H2020 projects

D3

Participants: Adrien Bousseau, Emilie Yu, David Jourdan, Felix Hähnlein, Changjian Li.

For more information see: [project website with description and main results](#).

The official EC site: [project on cordis.europa.eu](#).

Title: Interpreting Drawings for 3D Design

Duration: From February 1, 2017 to July 31, 2022

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France

Inria contact: Adrien BOUSSEAU

Coordinator: Adrien BOUSSEAU

Summary: 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.

During the final year of the project, we published several key contributions towards this goal. Specifically, we published an automatic algorithm to lift real-world design drawings to 3D by leveraging symmetry [22], and an algorithm that reconstructs a surface mesh from 3D drawings [21]. We also worked on a system that uses machine learning to parse freehand drawings and translate them into editable CAD models [17], along with a system to generate synthetic drawings from CAD models, which can serve to create synthetic datasets for training machine learning algorithms [12]. Finally, we presented a simulator and an inverse design tool to rapidly prototype physical models of freeform surfaces using 3D printing on fabric [13, 23].

FUNGRAPH

Participants: George Drettakis, Stavros Diolatzis, Georgios Kopanas, Sidhant Prakash, Bernhard Kerbl, Nicolás Violante.

[FUNGRAPH project on cordis.europa.eu](https://cordis.europa.eu/project/FUNGRAPH)

Title: A New Foundation for Computer Graphics with Inherent Uncertainty

Duration: From October 1, 2018 to September 30, 2024

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France

Inria contact: George Drettakis

Coordinator: George Drettakis

Summary: The use of Computer Graphics (CG) is constantly expanding, e.g., in Virtual and Augmented Reality, requiring realistic interactive renderings of complex virtual environments at a much wider scale than available today. CG has many limitations we must overcome to satisfy these demands. High-quality accurate rendering needs expensive simulation, while fast approximate rendering algorithms have no guarantee on accuracy; both need manually-designed expensive-to-create content. Capture (e.g., reconstruction from photos) can provide content, but it is uncertain (i.e., inaccurate and incomplete). Image-based rendering (IBR) can display such content, but lacks flexibility to modify the scene. These different rendering algorithms have incompatible but complementary tradeoffs in quality, speed and flexibility; they cannot currently be used together, and only IBR can directly use captured content. To address these problems

FunGraph will revisit the foundations of Computer Graphics, so these disparate methods can be used together, introducing the treatment of uncertainty to achieve this goal.

FunGraph introduces estimation of rendering uncertainty, quantifying the expected error of rendering components, and propagation of input uncertainty of captured content to the renderer. The ultimate goal is to define a unified renderer exploiting the advantages of each approach in a single algorithm. Our methodology builds on the use of extensive synthetic (and captured) “ground truth” data, the domain of Uncertainty Quantification adapted to our problems and recent advances in machine learning – Bayesian Deep Learning in particular.

FunGraph will fundamentally transform computer graphics, and rendering in particular, by proposing a principled methodology based on uncertainty to develop a new generation of algorithms that fully exploit the spectacular (but previously incompatible) advances in rendering, and fully benefit from the wealth offered by constantly improving captured content.

10.4 National initiatives

10.4.1 Visits of French scientists

For the Ph.D. defense of S. Prakash, Tamy Boubekeur (Adobe Research) and Romain Packanowski (Inria Bordeaux) visited the team and gave presentations in the mini-workshop "Recent Advances in Appearance Capture and Rendering".

For the Ph.D. defense of F. Hahnlein, Raphaëlle Chaine (Univ. Claude Bernard Lyon 1) visited the team and gave talks in a mini-workshop "Creating and Processing 3D Shapes".

10.4.2 ANR projects

ANR GLACIS [Project description on anr.fr](#)

Title: Graphical Languages for Creating Infographics

Duration: From April 1, 2022 to March 31, 2026

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- LIRIS
- University of Toronto

Inria contact: Theophanis TSANDILAS

Coordinator: Theophanis TSANDILAS

Summary: Visualizations are commonly used to summarize complex data, illustrate problems and solutions, tell stories over data, or shape public attitudes. Unfortunately, dominant visualization systems largely target scientists and data-analysis tasks and thus fail to support communication purposes. This project looks at visualization design practices. It investigates tools and techniques that can help graphic designers, illustrators, data journalists, and infographic artists, produce creative and effective visualizations for communication. The project aims to address the more ambitious goal of computer-aided design tools, where visualization creation is driven by the graphics, starting from sketches, moving to flexible graphical structures that embed constraints, and ending to data and generative parametric instructions, which can then re-feed the designer's sketches and graphics. The partners bring expertise from Human-Computer Interaction, Information Visualization, and Computer Graphics.

11 Dissemination

11.1 Promoting scientific activities

Member of the conference program committees G. Drettakis was member of the technical papers committee of SIGGRAPH 22.

Reviewer

- Adrien Bousseau was a reviewer for ACM Siggraph, Siggraph Asia, Eurographics.
- Guillaume Cordonnier was a reviewer for Eurographics 2022 short papers, Siggraph and ANR.
- Siddhant Prakash was a reviewer for the International Conference for 3D Vision (3DV22).
- Stavros Diolatzis was a reviewer for Eurographics'22.
- George Drettakis was a reviewer for SIGGRAPH Asia'22.
- Bernhard Kerbl was a reviewer for SIGGRAPH Asia Technical Communications and Posters and Computer Graphics Forum.

11.1.1 Journal

Member of the editorial boards Adrien Bousseau is associate editor for Computer Graphics Forum

11.1.2 Invited talks

Adrien Bousseau gave an invited talk at the Workshop on Sketching for Human Expressivity (SHE) at ECCV 2022. He also gave a talk at the Design for AI Symposium organized by TU Delft Industrial Design Engineering.

Guillaume Cordonnier presented at the workshop: Artificial Intelligence for Waves - [AI4W](#) and at the Academie des Sciences, Intersection des applications des sciences, joint talk on "Reality and simulation in the applications of geomorphology"; with Jean Braun (GFZ and Potsdam University).

George Drettakis gave invited talks at ETH Zurich and EPFL on Neural Rendering in November 2022.

11.1.3 Leadership within the scientific community

- George Drettakis chairs the Eurographics (EG) working group on Rendering, and the steering committee of EG Symposium on Rendering.
- George Drettakis serves as chair of the [ACM SIGGRAPH Papers Advisory Group](#) which chooses the technical papers chairs of ACM SIGGRAPH conferences and is responsible for all issues related to publication policy of our flagship conferences SIGGRAPH and SIGGRAPH Asia.

11.1.4 Scientific expertise

Adrien Bousseau participated in the selection committee for the [Gilles Khan Ph.D. award](#) given by the French Society in Computer Science (SIF).

11.1.5 Research administration

Guillaume Cordonnier is a member of the direction committee of the [GT IGRV CNRS working group](#) and chairs the GDR Ph.D. award.

George Drettakis is a member (suppleant) of the Inria Scientific Council, co-leads the Working Group on Rendering of the [GT IGRV](#) with R. Vergne and is a member of the Administrative Council of the [Eurographics French Chapter](#).

G. Drettakis chairs the organizing committee for the Morgenstern Colloquium in Sophia-Antipolis and he is a member of the platforms committee at Inria Sophia-Antipolis.

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- Adrien Bousseau, Guillaume Cordonnier and George Drettakis taught the course "Graphics and Learning", 32h, M2 at the Ecole Normale Supérieure de Lyon.
- Guillaume Cordonnier taught Machine Learning, with Pierre Alliez and Gaetan Bahl (NXP), Polytech Nice Sophia
- E. Yu taught at UCA: Object Oriented Conception and Programming, 54h (TD, L3)
- Bernhard Kerbl taught at FH Salzburg, for a total of 39 hours.

11.2.2 Supervision

- Ph.D. defended Jun. 22: Stavros Diolatzis, Guiding and Learning for Illumination Algorithms, since April 2019, George Drettakis. Thesis [\[24\]](#).
- Ph.D. in progress: Georgios Kopanas, Neural Rendering with Uncertainty for Full Scenes, since September 2020, George Drettakis.

- Ph.D. defended Dec. 22: Siddhant Prakash, Rendering with uncertain data with application to augmented reality, since November 2019, George Drettakis. Thesis [27].
- Ph.D. defended March 2022: David Jourdan, Computational design of self-shaping textiles, since October 2018, Adrien Bousseau. Thesis [26].
- Ph.D. defended Nov. 22: Felix Hähnlein, Binary Optimization for the Analysis and Synthesis of Concept Sketches, since September 2019, Adrien Bousseau. Thesis [25].
- Ph.D. in progress: Emilie Yu, 3D drawing in Virtual Reality, since October 2020, Adrien Bousseau.
- Ph.D. in progress: Nicolas Rosset, sketch-based design of aerodynamic shapes, since October 2021, Adrien Bousseau and Guillaume Cordonnier.
- Ph.D. in progress: Nicolas Violante, Generative Models for Neural Rendering, since October 2022, George Drettakis.

11.2.3 Juries

Adrien Bousseau was examiner for the Ph.D. thesis of Robin Kips (Institut Polytechnique de Paris) and a member of the committee for Elie Michel (Institut Polytechnique de Paris).

G. Drettakis was examiner for the Ph.D. thesis of D. Thul (ETH Zurich) and P. Mezières (Toulouse) and member of the committee for X. Yao (Telecom Paris).

11.3 Popularization

Adrien Bousseau gave a talk on the environmental impact of machine learning to high school students during the MATHC2+ event.

12 Scientific production

12.1 Major publications

- [1] V. Deschaintre, M. Aittala, F. Durand, G. Drettakis and A. Bousseau. ‘Single-Image SVBRDF Capture with a Rendering-Aware Deep Network’. In: *ACM Transactions on Graphics* 37 (2018), pp. 128–143. DOI: [10.1145/3197517.3201378](https://doi.org/10.1145/3197517.3201378). URL: <https://hal.inria.fr/hal-01793826>.
- [2] P. Ecormier-Nocca, G. Cordonnier, P. Carrez, A.-m. Moigne, P. Memari, B. Benes and M.-P. Cani. ‘Authoring Consistent Landscapes with Flora and Fauna’. In: *ACM Transactions on Graphics* (Aug. 2021). DOI: [10.1145/3450626.3459952](https://doi.org/10.1145/3450626.3459952). URL: <https://hal.inria.fr/hal-03245206>.
- [3] Y. Gryaditskaya, F. Hähnlein, C. Liu, A. Sheffer and A. Bousseau. ‘Lifting Freehand Concept Sketches into 3D’. In: *ACM Transactions on Graphics*. Proceedings of SIGGRAPH Asia (Nov. 2020). DOI: [10.1145/3414685.3417851](https://doi.org/10.1145/3414685.3417851). URL: <https://hal.inria.fr/hal-02952509>.
- [4] Y. Gryaditskaya, M. Sypsteyn, J. W. Hoftijzer, S. Pont, F. Durand and A. Bousseau. ‘OpenSketch: A Richly-Annotated Dataset of Product Design Sketches’. In: *ACM Transactions on Graphics* (2019). DOI: [10.1145/3355089.3356533](https://doi.org/10.1145/3355089.3356533). URL: <https://hal.inria.fr/hal-02284134>.
- [5] P. Hedman, J. Philip, T. Price, J.-M. Frahm, G. Drettakis and G. Brostow. ‘Deep Blending for Free-Viewpoint Image-Based Rendering’. In: *ACM Transactions on Graphics (SIGGRAPH Asia Conference Proceedings)* 37.6 (Nov. 2018). URL: <http://www-sop.inria.fr/revs/Basilic/2018/HPPFDB18>.
- [6] G.-A. Koulieris, B. Bui, M. S. Banks and G. Drettakis. ‘Accommodation and Comfort in Head-Mounted Displays’. In: *ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)* 36.4 (July 2017), p. 11. URL: <http://www-sop.inria.fr/revs/Basilic/2017/KBBD17>.
- [7] C. Li, H. Pan, A. Bousseau and N. J. Mitra. ‘Sketch2CAD: Sequential CAD Modeling by Sketching in Context’. In: *ACM Transactions on Graphics*. Proceedings SIGGRAPH Asia (2020). URL: <https://hal.inria.fr/hal-02937921>.

- [8] J. Philip, M. Gharbi, T. Zhou, A. A. Efros and G. Drettakis. ‘Multi-view Relighting using a Geometry-Aware Network’. In: *ACM Transactions on Graphics* 38 (2019). DOI: [10.1145/3306346.3323013](https://doi.org/10.1145/3306346.3323013). URL: <https://hal.inria.fr/hal-02125095>.
- [9] J. Philip, S. Morgenthaler, M. Gharbi and G. Drettakis. ‘Free-viewpoint Indoor Neural Relighting from Multi-view Stereo’. In: *ACM Transactions on Graphics* (2021). DOI: [10.1145/3469842](https://doi.org/10.1145/3469842). URL: <https://hal.inria.fr/hal-03265780>.
- [10] E. Yu, R. Arora, T. Stanko, J. A. Baerentzen, K. Singh and A. Bousseau. ‘CASSIE: Curve and Surface Sketching in Immersive Environments’. In: *CHI 2021 - ACM Conference on Human Factors in Computing Systems*. Yokohama, Japan, 8th May 2021. DOI: [10.1145/3411764.3445158](https://doi.org/10.1145/3411764.3445158). URL: <https://hal.inria.fr/hal-03149000>.

12.2 Publications of the year

International journals

- [11] S. Diolatzis, J. Philip and G. Drettakis. ‘Active Exploration for Neural Global Illumination of Variable Scenes’. In: *ACM Transactions on Graphics* (2022). DOI: [10.1145/3522735](https://doi.org/10.1145/3522735). URL: <https://hal.inria.fr/hal-03594357>.
- [12] F. Hähnlein, C. Li, N. J. M. Mitra and A. Bousseau. ‘CAD2Sketch: Generating Concept Sketches from CAD Sequences’. In: *ACM Transactions on Graphics* (2022). DOI: [10.1145/3550454.3555488](https://doi.org/10.1145/3550454.3555488). URL: <https://hal.inria.fr/hal-03842632>.
- [13] D. Jourdan, M. Skouras, E. Vouga and A. Bousseau. ‘Computational Design of Self-Actuated Surfaces by Printing Plastic Ribbons on Stretched Fabric’. In: *Computer Graphics Forum* 41.2 (Apr. 2022), pp. 1–14. DOI: [10.1111/cgf.14489](https://doi.org/10.1111/cgf.14489). URL: <https://hal.inria.fr/hal-03588434>.
- [14] B. Kim, X. Huang, L. Wuelfroth, J. Tang, G. Cordonnier, M. Gross and B. Solenthaler. ‘Deep Reconstruction of 3D Smoke Densities from Artist Sketches’. In: *Computer Graphics Forum* 41.2 (May 2022), pp. 97–110. DOI: [10.1111/cgf.14461](https://doi.org/10.1111/cgf.14461). URL: <https://hal.archives-ouvertes.fr/hal-03711264>.
- [15] G. Kopanas, T. Leimkühler, G. Rainer, C. Jambon and G. Drettakis. ‘Neural Point Catacaustics for Novel-View Synthesis of Reflections’. In: *ACM Transactions on Graphics* 41.6 (Dec. 2022), Article 201. DOI: [10.1145/3550454.3555497](https://doi.org/10.1145/3550454.3555497). URL: <https://hal.inria.fr/hal-03799441>.
- [16] M. Lastic, D. Rohmer, G. Cordonnier, C. Jaupart, F. Neyret and M.-P. Cani. ‘Interactive simulation of plume and pyroclastic volcanic ejections’. In: *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 5.1 (May 2022), pp. 1–15. DOI: [10.1145/3522609](https://doi.org/10.1145/3522609). URL: <https://hal.science/hal-03639288>.
- [17] C. Li, H. Pan, A. Bousseau and N. J. Mitra. ‘Free2CAD: Parsing Freehand Drawings into CAD Commands’. In: *ACM Transactions on Graphics* (24th July 2022). DOI: [10.1145/3528223.3530133](https://doi.org/10.1145/3528223.3530133). URL: <https://hal.inria.fr/hal-03683482>.
- [18] S. Prakash, G. Rainer, A. Bousseau and G. Drettakis. ‘Deep scene-scale material estimation from multi-view indoor captures’. In: *Computers & Graphics: X* (10th Oct. 2022). DOI: [10.1016/j.cag.2022.09.010](https://doi.org/10.1016/j.cag.2022.09.010). URL: <https://hal.inria.fr/hal-03812554>.
- [19] G. Rainer, A. Bousseau, T. Ritschel and G. Drettakis. ‘Neural Precomputed Radiance Transfer’. In: *Computer Graphics Forum* (Apr. 2022). DOI: [10.1111/cgf.14480](https://doi.org/10.1111/cgf.14480). URL: <https://hal.inria.fr/hal-03581439>.
- [20] J. Tang, V. Azevedo, G. Cordonnier and B. Solenthaler. ‘Neural Green’s function for Laplacian systems’. In: *Computers and Graphics* 107 (Oct. 2022), pp. 186–196. DOI: [10.1016/j.cag.2022.07.016](https://doi.org/10.1016/j.cag.2022.07.016). URL: <https://hal.science/hal-04000012>.
- [21] E. Yu, R. Arora, J. A. Bærentzen, K. Singh and A. Bousseau. ‘Piecewise-Smooth Surface Fitting onto Unstructured 3D Sketches’. In: *ACM Transactions on Graphics* 41.4 (July 2022), pp. 1–16. DOI: [10.1145/3528223.3530100](https://doi.org/10.1145/3528223.3530100). URL: <https://hal.inria.fr/hal-03691151>.

International peer-reviewed conferences

- [22] F. Hähnlein, Y. Gryaditskaya, A. Sheffer and A. Bousseau. 'Symmetry-driven 3D Reconstruction from Concept Sketches'. In: SIGGRAPH 2022 - International Conference & Exhibition On Computer Graphics & Interactive Techniques. 19. Vancouver, Canada: ACM, 7th Aug. 2022, pp. 1–8. DOI: [10.1145/3528233.3530723](https://doi.org/10.1145/3528233.3530723). URL: <https://hal.inria.fr/hal-03760673>.
- [23] D. Jourdan, V. Romero, E. Vouga, A. Bousseau and M. Skouras. 'Simulation of printed-on-fabric assemblies'. In: SCF 2022 - 7th annual ACM Symposium on Computational Fabrication. Seattle, United States: ACM, 26th Oct. 2022, pp. 1–11. DOI: [10.1145/3559400.3562001](https://doi.org/10.1145/3559400.3562001). URL: <https://hal.science/hal-03815213>.

Doctoral dissertations and habilitation theses

- [24] S. Diolatzis. 'Learning radiance fields : from global illumination to generative models'. Université Côte d'Azur, 24th June 2022. URL: <https://theses.hal.science/tel-03814937>.
- [25] F. Hähnlein. 'Binary Optimization for the Analysis and Synthesis of Concept Sketches'. Université Cote d'Azur, 29th Nov. 2022. URL: <https://hal.inria.fr/tel-03971826>.
- [26] D. Jourdan. 'Computational design of self-shaping textiles'. Université Côte d'Azur, 29th Mar. 2022. URL: <https://theses.hal.science/tel-03662318>.
- [27] S. Prakash. 'Using Appearance to Efficiently Render and Edit Captured Scenes'. Université Côte d'Azur, 20th Dec. 2022. URL: <https://hal.inria.fr/tel-03968751>.