

RESEARCH CENTRE

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

2023

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
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
3.4 Physical Simulation of Natural Phenomena	7
4 Application domains	8
5 Social and environmental responsibility	9
5.1 Footprint of research activities	9
5.2 Impact of research results	9
6 Highlights of the year	9
6.1 Awards	9
6.2 Press release	9
7 New software, platforms, open data	9
7.1 New software	9
7.1.1 sibr-core	9
7.1.2 IndoorRelighting	10
7.1.3 PBNR	10
7.1.4 CASSIE	10
7.1.5 fabsim	11
7.1.6 activeExploration	11
7.1.7 3DGaussianSplats	11
7.1.8 NerfShop	12
7.1.9 pyLowStroke	12
7.1.10 pySBM	13
7.1.11 CAD2Sketch	13
7.1.12 AGE	13
7.2 Open data	14
7.2.1 OpenSketch	14
7.2.2 Multi-view datasets	14
8 New results	14
8.1 Computer-Assisted Design with Heterogeneous Representations	14
8.1.1 VideoDoodles: Hand-Drawn Animations on Videos with Scene-Aware Canvases	14
8.1.2 3D Layer Compositing for VR Painting	14
8.1.3 STIVi: Turning Perspective Sketching Videos into Interactive Tutorials	15
8.1.4 CADTalk: An Algorithm and Benchmark for Semantic Commenting of CAD Programs	15
8.2 Graphics with Uncertainty and Heterogeneous Content	16
8.2.1 MesoGAN: Generative Neural Reflectance Shells	16
8.2.2 Improving NeRF Quality by Progressive Camera Placement for Free-Viewpoint Navigation	16
8.2.3 NeRFshop: Interactive Editing of Neural Radiance Fields	17
8.2.4 ModalNerf: Neural Modal Analysis and Synthesis for Free-Viewpoint Navigation in Dynamically Vibrating Scenes	17
8.2.5 3D Gaussian Splatting for Real-Time Radiance Field Rendering	18
8.2.6 Trim Regions for Online Computation of From-Region Potentially Visible Sets	19

8.2.7	Level-of-Detail 3D Gaussian Splatting	20
8.2.8	Scaling Point Cloud Diffusion Models	20
8.2.9	Relighting Radiance Fields with Generative Multi-Illumination Data	21
8.2.10	Reducing the Memory Footprint of 3D Gaussian Splatting	21
8.2.11	Physically-Based Lighting for 3D Generative Models of Cars	21
8.3	Physical Simulation for Graphics	22
8.3.1	Forming Terrains by Glacial Erosion	22
8.3.2	Ice-flow model emulator based on physics-informed deep learning	22
8.3.3	Interactive design of 2D car profiles with aerodynamic feedback	23
8.3.4	Volcanic Skies: coupling explosive eruptions with atmospheric simulation to create consistent skylscapes	24
8.3.5	Modeling Debris Flow for Large Scale Terrain Generation	24
8.3.6	Fast terrain erosion through analytical solutions of the stream power law	25
8.3.7	Windblown sand around obstacles – simulation and validation of deposition patterns	25
9	Bilateral contracts and grants with industry	25
9.1	Bilateral Grants with Industry	25
10	Partnerships and cooperations	25
10.1	International initiatives	25
10.1.1	Visits of international scientists	25
10.1.2	Visits to international teams	27
10.2	European initiatives	28
10.2.1	H2020 projects	28
10.3	National initiatives	29
10.3.1	Visits of French scientists	29
10.3.2	ANR projects	30
11	Dissemination	31
11.1	Promoting scientific activities	31
11.1.1	Scientific events: selection	31
11.1.2	Journal	31
11.1.3	Invited talks	31
11.1.4	Leadership within the scientific community	32
11.1.5	Scientific expertise	32
11.1.6	Research administration	32
11.1.7	Supervision	32
11.1.8	Juries	32
12	Scientific production	33
12.1	Major publications	33
12.2	Publications of the year	33

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
 - A6.1.4. – Multiscale modeling

- A6.1.5. – Multiphysics modeling
- A6.2. – Scientific computing, Numerical Analysis & Optimization
- A6.2.6. – Optimization
- A6.2.8. – Computational geometry and meshes
- A6.3.1. – Inverse problems
- A6.3.2. – Data assimilation
- A6.3.5. – Uncertainty Quantification
- A6.5.2. – Fluid mechanics
- A6.5.3. – Transport
- A8.3. – Geometry, Topology
- A9.2. – Machine learning
- A9.3. – Signal analysis
- A9.10. – Hybrid approaches for AI

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
- B9.3. – Medias
- B9.5.1. – Computer science
- B9.5.2. – Mathematics
- B9.5.3. – Physics
- B9.5.5. – Mechanics
- B9.5.6. – Data science
- 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

- Alban Gauthier [Inria, Post-Doctoral Fellow, from Mar 2023]
- Bernhard Kerbl [UNIV VIENNE, from Jul 2023]
- Georgios Kopanas [Inria, Post-Doctoral Fellow, from Dec 2023]
- Andreas Meuleman [Inria, Post-Doctoral Fellow, from Oct 2023]
- Anran Qi [Inria, Post-Doctoral Fellow, from Dec 2023]

PhD Students

- Berend Baas [INRIA, from Dec 2023]
- Aryamaan Jain [INRIA, from Oct 2023]
- Panagiotis Papantonakis [INRIA, from Sep 2023]
- Yohan Poirier-Ginter [UNIV LAVAL QUEBEC, from Sep 2023]
- Nicolas Rosset [INRIA]
- Petros Tzathas [INRIA, from Nov 2023]
- Nicolas Violante [INRIA]
- Emilie Yu [INRIA]

Technical Staff

- Alexandre Lanvin [INRIA, Engineer, from Dec 2023]

Interns and Apprentices

- Yannis Kedadry [Inria, Intern, from May 2023]
- Nazar Misyats [Inria, Intern, from Apr 2023 until Jul 2023]
- Vishal Pani [INRIA, Intern, from Aug 2023]
- Automne Petitjean [Inria, Intern, until Jan 2023]

Administrative Assistant

- Sophie Honnorat [INRIA, from Nov 2023]

Visiting Scientists

- Frederic Durand [MIT, from Sep 2023]
- Gilda Manfredi [UNIV SALENTO, from Oct 2023]

2 Overall objectives

In traditional Computer Graphics (CG), input is *accurately modeled* by artists. Artists first create the 3D geometry – i.e., the surfaces used to represent the 3D scene. This task can be achieved using tools akin to sculpting for human-made objects, or using physical simulation for objects formed by natural phenomena. Artists 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.

Creating all this 3D content by hand is a notoriously tedious process, both for novice users who do not have the skills to use complex modeling software, and for creative professionals who are primarily interested in obtaining a diversity of imagery and prototypes rather than in accurately specifying all the ingredients listed above. While physical simulation can alleviate some of this work for certain classes of objects (landscapes, fluids, plants), simulation algorithms are often costly and difficult to control.

Once all 3D elements of a scene are in place, a *rendering* algorithm is employed to generate a shaded, realistic image. Rendering algorithms typically involve the accurate simulation of light transport, accounting for the complex interactions between light and materials as light bounces over the surfaces of the scene to reach the camera. Similarly to the simulation of natural phenomena, the simulation of light transport is computationally expensive, and only provides meaningful results if the input is accurate and complete.

A major recent development is that many alternative sources of 3D content are becoming available. Cheap depth sensors but also video and photos allow anyone to capture real objects. However, the resulting 3D models are often inaccurate and incomplete due to limitations of these sensors and acquisition setups. There have also been significant advances in casual content creation, e.g., sketch-based modeling tools. But the resulting models are often approximate since people rarely draw accurate perspective and proportions, nor fine details. Unfortunately, the traditional Computer Graphics pipeline outlined above is unable to directly handle the *uncertainty* present in cheap sources of 3D content. 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 and simulation all the way to rendering and prototyping.

The technological shifts we mention above, together with developments in computer vision and machine learning, and the availability of large repositories of images, videos and 3D models represent a great opportunity for new imaging methods. In GraphDeco, we have identified three major scientific challenges that we strive to address to make such visual content widely accessible:

- First, the design pipeline needs to be revisited to **explicitly account for the variability and uncertainty of a concept and its representations**, from early sketches to 3D models and prototypes. Professional practice also needs to be adapted to be accessible to all.
- Second, a new approach is required to **develop computer graphics models and rendering algorithms capable of handling uncertain and heterogeneous data** as well as traditional synthetic content.
- Third, physical simulation needs to be combined with approximate user inputs to **produce content that is realistic and controllable**.

We have developed a common thread that unifies these three axes: the combination of machine learning with optimization and simulation, allowing the treatment of uncertain data for the synthesis of visual content. This common methodology – which falls under the umbrella term of *machine learning for visual computing* – provides a shared language and toolbox for the three research axes in our group, allowing frequent and in-depth collaborations between all three permanent researchers of the group, and a strong cohesive dynamic for Ph.D. students and postdocs.

As a result of this approach, GRAPHDECO is one of the few groups worldwide with in-depth expertise of both computer graphics techniques and deep learning approaches, in all three “traditional pillars” of CG: modeling, animation and rendering.

3 Research program

3.1 Introduction

Our research program is oriented around three main axes: 1) Computer-Assisted Design with Heterogeneous Representations, 2) Graphics with Uncertainty and Heterogeneous Content, and 3) Physical Simulation of Natural Phenomena. These three 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. Figure 1 illustrates some typical representations, including sketches, hand-made prototypes, 3D models, 3D printed prototypes or instructions.

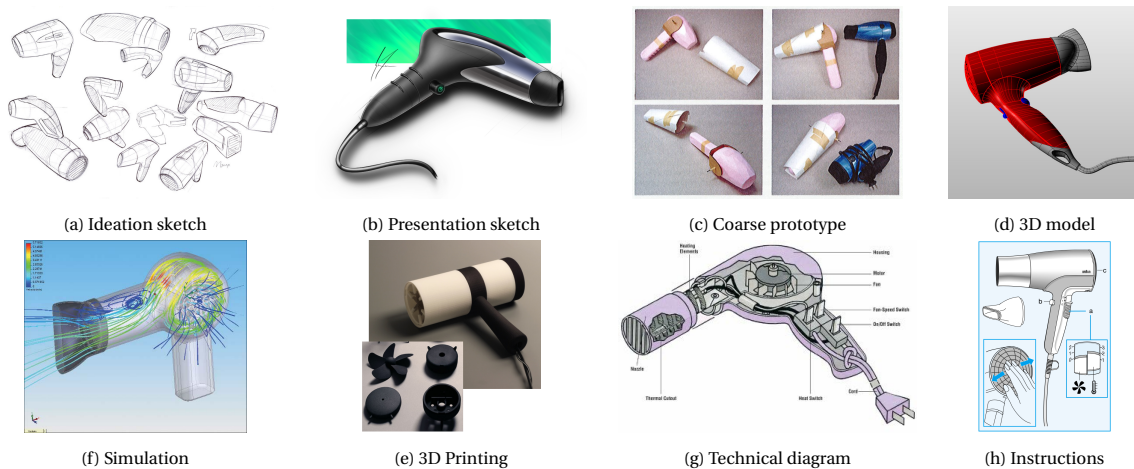


Figure 1: Various representations of a hair dryer at different stages of the design process. Image source, in order: c-maeng on deviantart.com, shauntur on deviantart.com, "Prototyping and Modelmaking for Product Design" Hallgrimsson, B., Laurence King Publishers, 2012, samsher511 on turbosquid.com, 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 analyze 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 modeled. This includes uncertainty in geometry, but also in materials and/or lights – which are often not even available with such content. Since CG algorithms cannot be applied directly, visual effects artists spend hundreds of hours correcting inaccuracies and completing the captured data to make them useable in film and advertising.

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

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) (see Fig. 2).

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



Figure 2: Image-Based Rendering (IBR) techniques use input photographs and approximate 3D to produce new synthetic views.

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.

3.4 Physical Simulation of Natural Phenomena

Our world emerged from the conjunction of natural phenomena at different scales, from the orogenesis of mountains to the evolution of ecosystems or the daily changes in weather conditions.

Understanding and modeling these phenomena is key to visually synthesizing our environments, reducing the uncertainty inherent to the capture of natural sceneries, and anticipating the impacts of natural hazards on our societies. For all these applications, the ability of a user to efficiently direct the simulation is preeminent, which provides us with two key constraints: first, the models should be *fast* to enable interactive interactions between the user and the simulation. Second, the models have to exhibit *efficient control* mechanisms.

The previous work on natural phenomena is as diverse as the number of scientific fields specialized in environmental and Earth sciences but with a main focus on predictability. In contrast, computer graphics has a long history of models focused on efficiency, robustness, and controllability, although originally explored for dynamic visual effects (smoke, explosions) and less so for natural phenomena that are more considered from a procedural or *phenomenon-based* perspective.

We benefit from computer graphics expertise in efficient and controllable physically-based simulations and extend it to natural phenomena. We explore new methods in machine learning and optimization, that enable us to enhance the efficiency of our models and reach a new space of forward and inverse control mechanisms. Coupling these models with physics provides guarantees on the quality of the results and a physical interpretation of the controls.

4 Application domains

Our research on design, simulation 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.

Earth and environmental sciences use simulations to understand and characterize natural processes. One of the common scientific methodologies requires testing several simulations with different sets of parameters to observe emergent behavior or to match observed data. Our fast simulation models accelerate this workflow, while our focus on control gives new tools to efficiently reduce the misfit between simulations and observations.

Natural hazard prevention is becoming ever more critical now that several climatic tipping points are crossed or about to be. Fast and controllable simulations of natural phenomena could allow public authorities to quickly assert different scenarios on the verge of imminent hazards, informing them of the probable impacts of their decisions.

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, including the development of innovative digital tools for circular design. Finally, the work in FUNGRAPH (G. Drettakis) has advanced research in visualization for reconstruction of real scenes. The recent 3D Gaussian Splatting [15] work has resulted in extensive technology transfer with many commercial licenses of the code already completed. These involve diverse industrial domains, including e-commerce, casual capture for 3D reconstruction, special effects for film, real-estate visualization and others.

6 Highlights of the year

6.1 Awards

The paper "3D Gaussian Splatting for Real-Time Radiance Field Rendering" [15] won the best paper award at SIGGRAPH 2023. This award is given to only 5 papers out of hundreds. The software of this paper has seen widespread adoption and has resulted in significant technology transfer, as mentioned above.

6.2 Press release

The ERC D3 project was selected to appear as a *science story* on the [ERC website](#).

7 New software, platforms, open data

7.1 New software

7.1.1 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

7.1.2 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 (<https://repo-sam.inria.fr/fungraph/deep-indoor-relight/>).

Functional Description: A neural relighting algorithm for captured indoors scenes, that allows interactive free-viewpoint navigation. Our method allows illumination to be changed synthetically, while coherently rendering cast shadows and complex glossy materials.

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

Contact: George Drettakis

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

Partner: Adobe

7.1.3 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/>).

Functional Description: This code provides an algorithm for novel-view synthesis based on neural rendering. Given multiple photographs of a scene, a point cloud is reconstructed using Multi-View Stereo. This point cloud is further optimized in the space of input views, including depth and reprojected features, resulting in improved novel-view synthesis.

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

Contact: George Drettakis

7.1.4 CASSIE

Name: CASSIE: Curve and Surface Sketching in Immersive Environments

Keywords: Virtual reality, 3D modeling

Scientific Description: Implementation of the article <https://hal.inria.fr/hal-03149000>

Functional Description: This system is a 3D conceptual modeling user interface in VR that leverages freehand mid-air sketching, and a novel 3D optimization framework to create connected curve network armatures, predictively surfaced using patches.

This project is composed of: - a VR sketching interface in Unity - an optimization 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

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

Contact: Emilie Yu

7.1.5 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, David Jourdan

Partners: Etienne Vouga, Méline Skouras

7.1.6 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: Active Exploration optimizes ground truth data generation during training of a Neural Generator using MCMC. The resulting Neural Generator can render variable scenes with effects at interactive rates that were previously prohibitive.

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

Authors: Stavros Diolatzis, Julien Philip, George Drettakis

Contact: George Drettakis

Partner: Adobe

7.1.7 3DGaussianSplats

Name: 3D Gaussian Splatting for Real-Time Radiance Field Rendering

Keywords: 3D, View synthesis, Graphics

Scientific Description: Implementation of the method 3D Gaussian Splatting for Real-Time Radiance Field Rendering, see <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>

Functional Description: 3D Gaussian Splatting is a method achieves real-time rendering of captured scenes with quality that equals the previous method with the best quality, while only requiring optimization times competitive with the fastest previous methods

3D Gaussian Splatting represents 3D scenes with 3D Gaussians that preserve desirable properties of continuous volumetric radiance fields for scene optimization while avoiding unnecessary computation in empty space. The method performs interleaved optimization/density control of the 3D Gaussians, notably optimizing anisotropic covariance to achieve an accurate representation of the scene. We provide a fast visibility-aware rendering algorithm that supports anisotropic splatting and both accelerates training and allows realtime rendering.

URL: <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>

Contact: George Drettakis

Participants: Georgios Kopanas, Bernhard Kerbl, Thomas Leimkuhler

7.1.8 NerfShop

Name: Interactive Editing of Neural Radiance Fields

Keywords: 3D, Deep learning

Scientific Description: Software implementation of the paper "Nerfshop: Interactive Editing of Neural Radiance Fields". See <https://repo-sam.inria.fr/fungraph/nerfshop/>

Functional Description: Neural Radiance Fields (NeRFs) have revolutionized novel view synthesis for captured scenes, with recent methods allowing interactive free-viewpoint navigation and fast training for scene reconstruction. NeRFshop is a novel end-to-end method that allows users to interactively edit NeRFs by selecting and deforming objects through cage-based transformations.

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

Contact: George Drettakis

7.1.9 pyLowStroke

Keywords: Vector-based drawing, 3D modeling

Scientific Description: This library contains several functionalities to process line drawings, including loading and saving in svg format, detecting vanishing points, calibrating a camera, detecting intersections. It has been used to develop our reconstruction algorithm Symmetry-driven 3D Reconstruction from Concept Sketches published at SIGGRAPH 2022 <https://ns.inria.fr/d3/SymmetrySketch/>

Functional Description: This is a library for low-level processing of freehand sketches. Example applications include reading and writing vector drawings, removing hooks, classifying lines into straight lines and curves and the calibration of a perspective camera model.

URL: <https://gitlab.inria.fr/D3/pylowstroke>

Contact: Adrien Bousseau

Participants: Felix Hahnlein, Yulia Gryaditskaya, Bastien Wailly, Adrien Bousseau

7.1.10 pySBM

Keywords: 3D modeling, Vector-based drawing

Scientific Description: This project is the official implementation of our paper Symmetry-driven 3D Reconstruction from Concept Sketches, published at SIGGRAPH 2022 <https://ns.inria.fr/d3/SymmetrySketch/>

Functional Description: This is a sketch-based modeling library. It proposes an interface to input a vector sketch, process it and reconstruct it in 3D. It also contains a Blender interface to trace a drawing and to interact with its 3D reconstruction.

URL: <https://gitlab.inria.fr/D3/pysbm>

Contact: Adrien Bousseau

Participants: Felix Hahnlein, Yulia Gryaditskaya, Alla Sheffer, Adrien Bousseau

Partner: University of British Columbia

7.1.11 CAD2Sketch

Keywords: Non-photorealistic rendering, CAD

Scientific Description: This code is the result of the following research project: <https://ns.inria.fr/d3/cad2sketch/>

Functional Description: This is a program to render CAD models as design drawings. The method generates lines for each part of the CAD model, and selects a subset of these lines to produce a drawing that is well constructed with as little clutter as possible. The resulting drawings look like drawings created by product designers.

URL: <https://gitlab.inria.fr/D3/cad2sketch>

Authors: Adrien Bousseau, Niloy Mitra, Changjian Li, Felix Hahnlein

Contact: Adrien Bousseau

7.1.12 AGE

Name: Accelerated Glacial Erosion

Keywords: Simulation, Terrain, Glacier, Erosion

Scientific Description: This software illustrates our method from the article "Guillaume Cordonnier, Guillaume Jovet, Adrien Peytavié, Jean Braun, Marie-Paule Cani, Bedrich Benes, Eric Galin, Eric Guérin, James Gain, Forming Terrains by Glacial Erosion, ACM Transactions on Graphics, Volume 42, Issue 4, 2023".

Functional Description: This software provides: - a model for the fast prediction of ice flow velocity based on learning - a multi-scale advection scheme to move the glacier over long time spans. - a coupled model for glacial, fluvial, and debris flow erosion.

Author: Guillaume Cordonnier

Contact: Guillaume Cordonnier

7.2 Open data

7.2.1 OpenSketch

OpenSketch is a dataset of product design sketches aimed at offering a rich source of information for a variety of computer-aided design tasks. OpenSketch contains more than 400 sketches representing 12 man-made objects drawn by 7 to 15 product designers of varying expertise. Each participant drew the objects from two perspective viewpoints. Together with industrial design teachers, we distilled a taxonomy of line types and used it to label each stroke of the 214 sketches drawn from one of the two viewpoints. While some of these lines have long been known in computer graphics, others remain to be reproduced algorithmically or exploited for shape inference. We also asked participants to produce clean presentation drawings from each of their sketches, resulting in aligned pairs of drawings of different styles. Finally, we registered each sketch to its reference 3D model by annotating sparse correspondences. Our sketches, in combination with provided annotations, form challenging benchmarks for existing algorithms as well as a great source of inspiration for future developments.

7.2.2 Multi-view datasets

We have acquired or generated multiple datasets for novel view synthesis and relighting. We list all of these dataset, along with accompanying source code, on the [Fungraph](#) project webpage.

8 New results

8.1 Computer-Assisted Design with Heterogeneous Representations

8.1.1 VideoDoodles: Hand-Drawn Animations on Videos with Scene-Aware Canvases

Participants: Emilie Yu, Adrien Bousseau, Kevin Blackburn-Matzen (*Adobe Research*), Kevin Nguyen (*Adobe Research*), Oliver Wang (*Adobe Research*), Rubaiat Kazi (*Adobe Research*).

We present an interactive system to ease the creation of so-called video doodles – videos on which artists insert hand-drawn animations for entertainment or educational purposes (see [Figure 3](#)). Video doodles are challenging to create because to be convincing, the inserted drawings must appear as if they were part of the captured scene. In particular, the drawings should undergo tracking (the face and arms follow the tramway in [Figure 3](#)), perspective deformations and occlusions (the bridge is occluded in [Figure 3](#)) as they move with respect to the camera and to other objects in the scene – visual effects that are difficult to reproduce with existing 2D video editing software. Our system supports these effects by relying on planar canvases that users position in a 3D scene reconstructed from the video. Furthermore, we present a custom tracking algorithm that allows users to anchor canvases to static or dynamic objects in the scene, such that the canvases move and rotate to follow the position and direction of these objects. When testing our system, novices could create a variety of short animated clips in a dozen of minutes, while professionals praised its speed and ease of use compared to existing tools.

This work has been published in *ACM Transactions on Graphics*, and presented in SIGGRAPH 2023 [19].

8.1.2 3D Layer Compositing for VR Painting

Participants: Emilie Yu, Adrien Bousseau, Fanny Chevalier (*University of Toronto*), Karan Singh (*University of Toronto*).

VR painting is a novel 3D authoring workflow in which artists create 3D scenes with a painterly non-photorealistic look by arranging many colored 3D strokes. To overcome challenges faced by users in editing strokes color, we take inspiration from the well-established interaction paradigm of layer



Figure 3: Video doodles combine hand-drawn animations with video footage. Our interactive system eases the creation of this mixed media art by letting users place planar canvases in the scene which are then tracked in 3D. In this example, the inserted rainbow bridge exhibits correct perspective and occlusions, and the character's face and arms follow the tram as it runs towards the camera.

compositing from 2D digital painting and propose a data structure and real-time rendering algorithm to support *3D Layers* in VR painting.

8.1.3 STIVi: Turning Perspective Sketching Videos into Interactive Tutorials

Participants: Adrien Bousseau, Capucine Nghiem (*Inria ExSitu*, *LISN*), Theophanis Tsandilas (*Inria ExSitu*, *LISN*), Jan Willem Hoftijzer (*TU Delft*), Mark Sypesteyn (*TU Delft*), Maneesh Agrawala (*Stanford*).

For design and art enthusiasts who seek to enhance their skills through instructional videos, following drawing instructions while practicing can be challenging. STIVi converts the perspective drawing demonstrations and commentary of prerecorded instructional videos into interactive drawing tutorials that students can navigate and explore at their own pace.

Our approach involves a semi-automatic pipeline for creating STIVi content, extracting pen strokes from video frames and aligning them with the accompanying audio commentary. Thanks to this structured data, students can navigate through transcript and in-video drawing, refer to provided highlights in both modalities to guide their navigation, and explore variations of the drawing demonstration to understand fundamental principles, in STIVi's interface. We evaluated STIVi against a regular video player. We observed that our interface supports nonlinear learning styles by providing alternative paths for following and understanding drawing instructions.

8.1.4 CADTalk: An Algorithm and Benchmark for Semantic Commenting of CAD Programs

Participants: Adrien Bousseau, Haocheng Yuan (*University of Edinburgh*), Jing Xu (*University of Edinburgh*), Changjian Li (*University of Edinburgh*), Hao Pan (*Microsoft Research Asia*), Niloy Mitra (*UCL*).

CAD programs are a popular way to compactly encode shapes as a sequence of operations that are easy to parametrically modify. However, without sufficient semantic comments and structure, such programs can be challenging to understand, let alone modify. We introduce the problem of semantic commenting CAD programs, wherein the goal is to segment the input program into code blocks corresponding to semantically meaningful shape parts and assign a semantic label to each block. We solve the problem by combining program parsing with visual-semantic analysis afforded by recent advances in foundational language and vision models. Specifically, by executing the input programs, we create shapes, which we use to generate conditional photorealistic images to make use of semantic annotators for such images. We then distill the information across the images and link back to the original programs to semantically comment on them. Additionally, we collected and annotated a benchmark dataset, CADTalk, consisting of 5,280 machine-made programs and 45 human-made programs with ground truth semantic comments to foster future research. We extensively evaluated our approach, compared to a GPT-based baseline

approach, and an open-set shape segmentation baseline, i.e., PartSLIP, and reported an 83.24% accuracy on the new CADTalk dataset.

8.2 Graphics with Uncertainty and Heterogeneous Content

8.2.1 MesoGAN: Generative Neural Reflectance Shells

Participants: Stavros Diolatzis, George Drettakis, Jan Nova'k (*Nvidia*), Fabrice Roussele (*Nvidia*), Miika Aittala (*Nvidia*), Jonathan Granskog (*Runway*), Ravi Ramamoorthi (*UC San Diego*).

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 (Fig 4). Our primitive can be applied to a surface and used in a path tracer, thanks to our Neural Reflectance Shells. To generate 3D textures without repetition artifacts and with arbitrary extent we introduce randomized Fourier features in the layers of StyleGAN. We augment the 2D base feature texture with learned height, used by the neural field renderer, to generate a 3D texture. To enable training of our end-to-end pipeline within existing memory constraints and to allow high quality rendering, we introduce a filtering approach and train our method on multi-scale synthetic datasets of 3D mesoscale structures. Importantly, our method allows MesoGAN to be conditioned on artistic parameters, e.g., albedo, fiber length, density of strands as well as lighting direction, while maintaining high quality.



Figure 4: Our method combines the strengths of StyleGAN and volumetric neural field rendering to generate a 3D mesoscale texture that can be mapped to objects and used in a path tracer (c). We train on datasets of synthetic patches (a); our method can generate textures that have artistic parameters (such as fur saturation and length) which can be used to create shell maps of arbitrary extent (b).

This work was published in Computer Graphics Forum and presented in Eurographics Symposium on Rendering (EGSR) 2023 [12].

8.2.2 Improving NeRF Quality by Progressive Camera Placement for Free-Viewpoint Navigation

Participants: Georgios Kopanas, George Drettakis.

Neural Radiance Fields, or NeRFs, have drastically improved novel view synthesis and 3D reconstruction for rendering. NeRFs achieve impressive results on object-centric reconstructions, but the quality of

novel view synthesis with free-viewpoint navigation in complex environments (rooms, houses, etc) is often problematic. While algorithmic improvements play an important role in the resulting quality of novel view synthesis, in this work, we show that because optimizing a NeRF is inherently a data-driven process, good quality data play a fundamental role in the final quality of the reconstruction. As a consequence, it is critical to choose the data samples - in this case the cameras - in a way that will eventually allow the optimization to converge to a solution that allows free-viewpoint navigation with good quality. Our main contribution is an algorithm that efficiently proposes new camera placements that improve visual quality with minimal assumptions (Fig 5). Our solution can be used with any NeRF model and outperforms baselines and similar work.

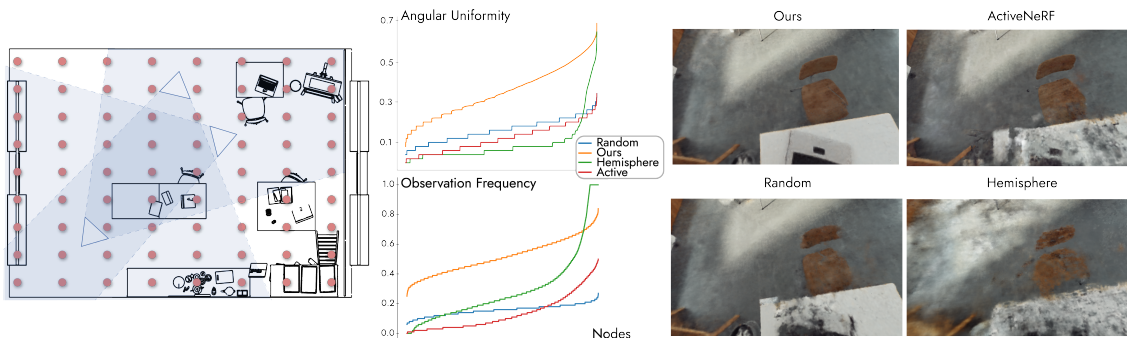


Figure 5: We present a new method that proposes the next best camera placement for NeRF capture (left). We introduce two metrics that can be easily computed, observation frequency and angular uniformity (middle). On the right, we show that our approach outperforms two baseline camera placement strategies, HEMISPHERE which is the typical approach used in most NeRF methods and RANDOM, as well as recent related work ActiveNeRF

It was presented in the International Symposium on Vision, Modeling, and Visualization (VMV) 2023 [20].

8.2.3 NeRFshop: Interactive Editing of Neural Radiance Fields

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

Neural Radiance Fields (NeRFs) have revolutionized novel view synthesis for captured scenes, with recent methods allowing interactive free-viewpoint navigation and fast training for scene reconstruction. However, the implicit representations used by these methods—often including neural networks and complex encodings—make them difficult to edit. Some initial methods have been proposed, but they suffer from limited editing capabilities and/or from a lack of interactivity, and are thus unsuitable for interactive editing of captured scenes. We tackle both limitations and introduce NeRFshop, a novel end-to-end method that allows users to interactively select and deform objects through cage-based transformations (Fig 6). NeRFshop provides fine scribblebased user control for the selection of regions or objects to edit, semi-automatic cage creation, and interactive volumetric manipulation of scene content thanks to our GPU-friendly two-level interpolation scheme. Further, we introduce a preliminary approach that reduces potential resulting artifacts of these transformations with a volumetric membrane interpolation technique inspired by Poisson image editing and provide a process that “distills” the edits into a standalone NeRF representation

This work was published in the journal ACM on Computer Graphics and Interactive Techniques and was presented at the Interactive 3D Graphics and Games (I3D) conference [13].

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



Figure 6: NeRFshop enables intuitive selection via scribbles and interactive editing of arbitrary NeRF scenes. We show duplicative, affine, and non-affine edits (left-to-right) with different viewpoints in the Kitchen scene

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

Recent advances in Neural Radiance Fields enable the capture of scenes with motion. However, editing the motion is hard; no existing method allows editing beyond the space of motion existing in the original video, nor editing based on physics.

We present the first approach that allows physically-based editing of motion in a scene captured with a single hand-held video camera, containing vibrating or periodic motion. We first introduce a Lagrangian representation, representing motion as the displacement of particles, which is learned while training a radiance field. We use these particles to create a continuous representation of motion over the sequence, which is then used to perform a modal analysis of the motion thanks to a Fourier transform on the particle displacement over time.

The resulting extracted modes allow motion synthesis, and easy editing of the motion, while inheriting the ability for free-viewpoint synthesis in the captured 3D scene from the radiance field (Fig. 7). We demonstrate our new method on synthetic and real captured scenes.




Figure 7: ModalNeRF allows physically-plausible synthesis of new motion in NeRF reconstructed dynamic scenes.

This work is a collaboration with Ayush Tewari from MIT and is published at the Eurographics Symposium on Rendering [16].

8.2.5 3D Gaussian Splatting for Real-Time Radiance Field Rendering

Participants: Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler (*Max-Planck Institute for Informatics*), George Drettakis.

This work is a collaboration with Thomas Leimkühler from the Max-Planck Institute for Informatics, Saarbrücken, Germany.

Radiance Field methods have recently revolutionized novel-view synthesis of scenes captured with multiple photos or videos. However, achieving high visual quality still requires neural networks that are costly to train and render, while recent faster methods inevitably trade off speed for quality. For unbounded and complete scenes (rather than isolated objects) and 1080p resolution rendering, no current method can achieve real-time display rates. We introduce three key elements that allow us to achieve state-of-the-art visual quality while maintaining competitive training times and importantly allow high-quality real-time (≥ 30 fps) novel-view synthesis at 1080p resolution. First, starting from sparse points produced during camera calibration, we represent the scene with 3D Gaussians that preserve desirable properties of continuous volumetric radiance fields for scene optimization while avoiding unnecessary computation in empty space; Second, we perform interleaved optimization/density control of the 3D Gaussians, notably optimizing anisotropic covariance to achieve an accurate representation of the scene; Third, we develop a fast visibility-aware rendering algorithm that supports anisotropic splatting and both accelerates training and allows realtime rendering. We demonstrate state-of-the-art visual quality and real-time rendering on several established datasets. See Fig 8.




Figure 8: Demonstration of the quality and speed achievable with 3D Gaussian Splatting. We use the PNSR (Peak signal-to-noise ratio) as an objective metric of image quality.

This work was published in ACM TOG [15], and presented at ACM SIGGRAPH where it received one of five best paper awards.

8.2.6 Trim Regions for Online Computation of From-Region Potentially Visible Sets

Participants: Philip Voglreiter (*Graz University of Technology*), Bernhard Kerbl, Alexander Weinrauch (*Graz University of Technology*), Joerg Hermann Mueller (*Graz University of Technology*), Thomas Neff (*Graz University of Technology*), Markus Steinberger (*Graz University of Technology*), Dieter Schmalstieg (*Graz University of Technology*).

Visibility computation is a key element in computer graphics applications. More specifically, a from-region potentially visible set (PVS) is an established tool in rendering acceleration, but its high computational cost means a from-region PVS is almost always precomputed. Precomputation restricts the use of PVS to static scenes and leads to high storage cost, in particular, if we need fine-grained regions. For dynamic applications, such as streaming content over a variable-bandwidth network, online PVS computation with configurable region size is required. We address this need with trim regions, a new method for generating from-region PVS for arbitrary scenes in real time. Trim regions perform controlled erosion of object silhouettes in image space, implicitly applying the shrinking theorem known from previous work. Our algorithm is the first that applies automatic shrinking to unconstrained 3D scenes, including non-manifold meshes, and does so in real time using an efficient GPU execution model. We

demonstrate that our algorithm generates a tight PVS for complex scenes and outperforms previous online methods for from-viewpoint and from-region PVS. It runs at 60 Hz for realistic game scenes consisting of millions of triangles and computes PVS with a tightness matching or surpassing existing approaches. See Fig 9.




Figure 9: Illustration of the Trim Regions Method. (a) The view in the inset is seen by an observer standing at the location indicated by the inset’s leader line. Our method computes a potentially visible set (PVS) corresponding to a viewcell (region) of a given radius around the viewpoint in real time. The six images each show a PVS for viewcell sizes from 5-30 cm. The truly visible part of the scene is shown shaded, while false positives are shown in blue and the remaining scene is shown in grey. No false negatives are visible. (b) Note how the width of the visible “corridor” on the floor progressively expands with the viewcell size. (c) The base of the crane, which was a false positive from 5 to 15 cm, becomes a true part of the PVS at 20 cm and above. It is typical that false positives become true positives as the viewcell size expands, since they are “almost visible” when first observed.

It was published in ACM Transactions on Graphics, and presented at ACM SIGGRAPH [18].

8.2.7 Level-of-Detail 3D Gaussian Splatting

Participants: Bernhard Kerbl, Andreas Meulemann, Georgios Kopanas, George Drettakis.

The recently presented 3D Gaussian Splatting methodology provides high quality and performance, but is not inherently scalable. In this project, we develop the necessary data structures and methods to enable 3D Gaussian Splat Rendering with adaptive level-of-detail for massively large scenes.

8.2.8 Scaling Point Cloud Diffusion Models

Participants: Georgios Kopanas, George Drettakis, Bernhard Kerbl.

Diffusion models have been quite successful in the last years. They have shown incredible state-of-the-art results in the context image and video generation. More recently, many works try to use diffusion models for 3D model generation. This is quite a challenging task compared to image or video generation for a few reasons. First, the access to available 3D data is limited compared to the billions of samples one can gather for images or videos. Second, the curse of dimensionality makes the diffusion process exponentially slower or memory intensive or both as the number of dimensions of the underlying distribution increases. Third, often 3D models are represented as unordered sets of primitives i.e., points

or triangles etc. while images and videos are signals sampled on a regular grid. This makes it easier to create efficient models that parse these signals. Image and video models are based on convolution operators which is not trivial to apply to geometry signals efficiently.

In this project we focus on generative models that describe a distribution of 3D geometries in the form of point clouds that can scale.

This work is an ongoing collaboration with Adobe.

8.2.9 Relighting Radiance Fields with Generative Multi-Illumination Data

Participants: Yohan Poirier-Ginter, George Drettakis, Julien Philip (*Adobe Research*), Jean-François Lalonde (*Université Laval*).

Producing relightable radiance fields is challenging without multi-illumination capture data. In this work, we explore the possibilities offered by mutually inconsistent, inaccurate, yet hyperrealistic multi-illumination data generated by a large 2D diffusion model fined-tuned for single-view relighting. We show that inconsistencies can be resolved, and accuracy recovered by leveraging physically based renders for 3D alignment.

8.2.10 Reducing the Memory Footprint of 3D Gaussian Splatting

Participants: Panagiotis Papantonakis, Bernhard Kerbl, Georgios Kopanas, Alexandre Lanvin, George Drettakis.

3D Gaussian splatting provides excellent visual quality for novel view synthesis, with fast training and real-time rendering. Unfortunately, the memory requirements of this method are unreasonably high. We first analyze the reasons for this, identifying three main areas where storage can be reduced: the number of 3D Gaussian primitives used to represent a scene, the number of coefficients for the spherical harmonics used to represent directional radiance and the precision required to store Gaussian primitive attributes. We present a solution to each of these issues. First, we propose an efficient, resolution-aware primitive pruning approach, reducing primitive count by half. Second, we introduce an adaptive adjustment method to choose the number of coefficients used to represent directional radiance for each Gaussian primitive, and finally a codebook-based quantization method, together with a half-float representation for further memory reduction. Taken together, these three components result in a 27 reduction factor in overall memory usage on the standard datasets we tested, resulting in a 1.5 times speedup in rendering speed. We demonstrate our method on standard datasets, and show how our solution greatly reduces download time when using the method on a mobile device.

8.2.11 Physically-Based Lighting for 3D Generative Models of Cars

Participants: Nicola Violante, Alban Gauthier, George Drettakis.

Recent work has demonstrated that Generative Adversarial Networks (GANs) can be trained to generate 3D content from 2D image collections, by synthesizing features for neural radiance field rendering. However, most such solutions generate radiance, with lighting entangled with materials. This results in unrealistic appearance, since lighting cannot be changed and view-dependent effects such as reflections do not move correctly with the viewpoint. In addition, many methods have difficulty for full, 360° rotations, since they are often designed for mainly front-facing scenes such as faces. We introduce a new 3D GAN framework that addresses these shortcomings, allowing multi-view coherent 360° viewing and at the same time relighting for objects with shiny reflections, which we exemplify using a car dataset. The success of our solution stems from three main contributions. First, we estimate initial camera poses for a dataset of car images, and then learn to refine the distribution of camera parameters while training the

GAN. Second, we propose an efficient Image-Based Lighting model, that we use in a 3D GAN to generate disentangled reflectance, as opposed to the radiance synthesized in most previous work. The material is used for physically-based rendering with a dataset of environment maps. Third, we improve the 3D GAN architecture compared to previous works and design a careful training strategy that allows effective disentanglement. Our model is the first that generate a variety of 3D cars that are multi-view consistent and that can be relit interactively with any environment map.

8.3 Physical Simulation for Graphics

8.3.1 Forming Terrains by Glacial Erosion

Participants: Guillaume Cordonnier, Guillaume Jovet (*University of Lausanne*), Jean Braun (*GFZ - Potsdam University*), Bedrich Benes (*Purdue University*), James Gain (*Cape Town University*), Marie-Paule Cani (*Ecole Polytechnique*), Eric Galin (*LIRIS*), Eric Guérin (*LIRIS*), Adrien Peytavie (*LIRIS*).

We introduce the first solution for simulating the formation and evolution of glaciers, together with their attendant erosive effects, for periods covering the combination of glacial and inter-glacial cycles. Our efficient solution includes both a fast yet accurate deep learning-based estimation of high-order ice flows and a new, multi-scale advection scheme enabling us to account for the distinct time scales at which glaciers reach equilibrium compared to eroding the terrain. We combine the resulting glacial erosion model with finer-scale erosive phenomena to account for the transport of debris flowing from cliffs. This enables us to model the formation of terrain shapes not previously adequately modeled in Computer Graphics, ranging from U-shaped and hanging valleys to fjords and glacial lakes. (Fig. 10).




Figure 10: Landscape carved by our simulation of glacier erosion. We observe specific glacial features: U-shape valleys, hanging valleys, horns, ridges, cirques and perched lakes.

This work was published in ACM Transactions on Graphics (TOG) [11] and presented at Siggraph 2023.

8.3.2 Ice-flow model emulator based on physics-informed deep learning

Participants: Guillaume Cordonnier, Guillaume Jovet (*University of Lausanne*).

Convolutional neural networks (CNN) trained from high-order ice-flow model realizations have proven to be outstanding emulators in terms of fidelity and computational performance. However, the dependence on an ensemble of realizations of an instructor model renders this strategy difficult to generalize to a variety of ice-flow regimes found in the nature. To overcome this issue, we adopt the approach of physics-informed deep learning, which fuses traditional numerical solutions by finite differences/elements and deep-learning approaches. Here, we train a CNN to minimize the energy associated with high-order ice-flow equations within the time iterations of a glacier evolution model. As a result, our emulator is a promising alternative to traditional solvers thanks to its high computational efficiency (especially on GPU), its high fidelity to the original model, its simplified training (without requiring any data), its capability to handle a variety of ice-flow regimes and memorize previous solutions, and its relatively simple implementation. Embedded into the 'Instructed Glacier Model' (IGM) framework, the potential of the emulator is illustrated with three applications including a large-scale high-resolution (2400x4000) forward glacier evolution model, an inverse modeling case for data assimilation, and an ice shelf (Fig. 11).




Figure 11: Ice thickness of the alpine icefield obtained at 21'000 years BP modeled at 200 meters of resolution. Our physically-based learning model enables the modeling of glacier evolution at unprecedented spatial and temporal scales.

This work was published in the *Journal of Glaciology* [14].

8.3.3 Interactive design of 2D car profiles with aerodynamic feedback

Participants: Nicolas Rosset, Guillaume Cordonnier, Adrien Bousseau, Régis Duvigneau (*Inria Acumes*).

The design of car shapes requires a delicate balance between aesthetic and performance. While fluid simulation provides the means to evaluate the aerodynamic performance of a given shape, its computational cost hinders its usage during the early explorative phases of design, when aesthetic is decided upon. We present an interactive system to assist designers in creating aerodynamic car profiles. Our system 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. Compared to prior work that focused 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. Furthermore, we architected our model to support gradient-based shape optimization within a learned latent space of car profiles. In addition to regularizing the optimization process, this latent space and an associated encoder-decoder allows us to input and output car profiles in a bitmap form, without any explicit parameterization of the car boundary. Finally, we designed our model to support pointwise queries of fluid properties around car shapes, allowing us to adapt computational cost to application needs. As an illustration, we only query our model along streamlines for flow visualization, we query it in the vicinity of the car for drag optimization, and we query it behind the car for vortex attenuation. See Fig 12.




Figure 12: Our system takes as input the profile of a car (a) and predicts the flow field around the car (b). We perform shape optimization in a latent space of cars to suggest how to improve the aerodynamic properties of the profile (c, here by reducing drag by 11%). Both the fluid flow visualization and the shape optimization are computed within milliseconds, enabling an interactive workflow where designers can iterate between sketching a car profile and evaluating its performance.

This work is was published in the Computer Graphics Forum and presented at Eurographics [17].

8.3.4 Volcanic Skies: coupling explosive eruptions with atmospheric simulation to create consistent skylscapes

Participants: Guillaume Cordonnier, Pieter Pretorius (*University of Cape Town*), James Gain (*University of Cape Town*), Maud Lastic (*Ecole Polytechnique*), Chen Jiong (*Ecole Polytechnique*), Damien Rohmer (*Ecole Polytechnique*), Marie-Paule Cani (*Ecole Polytechnique*).

Explosive volcanic eruptions rank among the most terrifying natural phenomena, and are thus frequently depicted in films, games, and other media, usually with a bespoke once-off solution. In this project, we study a general-purpose model for bi-directional interaction between the atmosphere and a volcano plume. We approximate the plume dynamics with Lagrangian disks and spheres and the atmosphere with sparse layers of 2D Eulerian grids, then generate volumetric animations by noise-based procedural upsampling.

8.3.5 Modeling Debris Flow for Large Scale Terrain Generation

Participants: Aryamaan Jain, Guillaume Cordonnier.

Terrain generation plays a crucial role in various graphics applications, with an emphasis on realistic representation. However, the phenomenon of debris flow, a significant natural process shaping terrains,

has often been overlooked or inaccurately modeled in previous works. Existing models either demonstrated simplicity or lacked geological accuracy. We introduce a geologically validated debris flow model that addresses these limitations. Our model incorporates erosion and deposition processes, enhancing the fidelity of terrain generation in graphics applications. To facilitate user control and interaction, we have developed a user-friendly tool within Houdini. This research contributes to the advancement of realistic terrain generation by presenting a more accurate and versatile debris flow model.

8.3.6 Fast terrain erosion through analytical solutions of the stream power law

Participants: Petros Tzathas, Guillaume Cordonnier, Boris Gailleton (*University of Rennes*), Philippe Steer (*University of Rennes*).

Terrain generation methods have long been divided between procedural and physically-based. While procedural methods are faster, they lack the geological consistency which naturally results from physically-based simulation. In particular, the simulation of the competition between tectonic uplift and fluvial erosion expressed by the stream power law, has been favored since it allows the generation and control of large-scale mountain ranges. We aim at combining advantages of both approaches by utilizing analytical solutions to the stream power law. In our approach, time does not dictate the duration of the algorithm but acts as a parameter to a function, a slider that controls the aging of the input terrain.

8.3.7 Windblown sand around obstacles – simulation and validation of deposition patterns

Participants: Nicolas Rosset, Guillaume Cordonnier, Adrien Bousseau, Régis O.

Sand dunes are iconic landmarks of deserts, but can also put human infrastructures at risk, for instance by forming near buildings or roads. We developed a simulator of sand erosion and deposition to predict how dunes form around and behind obstacles under wind. Inspired by both computer graphics and geo-sciences, our algorithm couples a fast wind flow simulation with physical laws of sand saltation and avalanching, which suffices to reproduce characteristic patterns of sand deposition. In particular, we validate our approach by reproducing real-world erosion and deposition measurements collected by prior work under controlled conditions.

This work is a collaboration with Régis Duvigneau.

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 Visits of international scientists

Other international visits to the team

Fredo Durand

Status: Professor

Institution of origin: MIT

Country: USA

Dates: August 2022 to July 2023

Context of the visit: Sabbatical year

Mobility program/type of mobility: Self-funded sabbatical visit

Andeas Bærentzen

Status: Professor

Institution of origin: Technical University of Denmark

Country: Denmark

Dates: December 20-22nd

Context of the visit: Ph.D. defense of Emilie Yu

Mobility program/type of mobility: Invited talk in the workshop on User Interaction, Immersive Storytelling, 3D Modeling and Character Animation

Hendrik Lensch

Status: Professor

Institution of origin: Univ.of Tübingen

Country: Germany

Dates: November 23-25th

Context of the visit: Ph.D. defense of Georgios Kopanas

Mobility program/type of mobility: Invited talk in the workshop on Neural Radiance Fields and Generative Models

Thomas Müller

Status: Researcher

Institution of origin: Nvidia

Country: Switzerland

Dates: November 23-25th

Context of the visit: Ph.D. defense of Georgios Kopanas

Mobility program/type of mobility: Invited talk in the workshop on Neural Radiance Fields and Generative Models

Changjian Li

Status: Assistant Professor

Institution of origin: Univ. Edinburgh

Country: United Kingdom

Dates: November 12-14th

Context of the visit: Group retreat

Mobility program/type of mobility: Invited talk

Gilda Manfredi

Status: Ph.D. student

Institution of origin: University of Basilicata and University Salento

Country: Italy

Dates: October 1st 2023 - March 31st 2024

Context of the visit: Research collaboration

Mobility program/type of mobility: Self-funded visiting Ph.D. student

10.1.2 Visits to international teams**Sabbatical programme****Adrien Bousseau**

Visited institution: TU Delft

Country: Netherlands

Dates: August 2022 to July 2023

Context of the visit: Sabbatical stay

Mobility program/type of mobility: Inria sabbatical program

Research stays abroad**Georgios Kopanas**

Visited institution: Adobe

Country: USA

Dates: July-August 2023

Context of the visit: Industrial internship

Mobility program/type of mobility: internship

Emilie Yu**Visited institution:** Dynamic Graphics Project at the University of Toronto**Country:** Canada**Dates:** June-August 2023**Context of the visit:** Research visit**Mobility program/type of mobility:** MITACS**10.2 European initiatives****10.2.1 H2020 projects****FUNGRAPH**

Participants: George Drettakis, Stavros Diolatzis, Clement Jambon, Bernhard Kerbl, Thomas Leimkuhler, Panagiotis Papantonakis, Automne Petitjean, Yohan Poirier-Ginter, Siddhant Prakash, Petros Tzathas, Alban Gauthier, Georgios Kopanas, Andreas Meuleman, Nicolas Violante, Alexandre Lanvin.

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

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 revisits 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.

10.3 National initiatives

10.3.1 Visits of French scientists

Wendy MacKay

Status: Researcher

Institution of origin: Inria

Country: France

Dates: December, 20-22nd

Context of the visit: Ph.D. defense of Emilie Yu

Mobility program/type of mobility: Invited talk in the workshop on User Interaction, Immersive Storytelling, 3D Modeling and Character Animation

Damien Rhomer

Status: Professor

Institution of origin: Ecole Polytechnique

Country: France

Dates: December, 20-22nd

Context of the visit: Ph.D. defense of Emilie Yu

Mobility program/type of mobility: Invited talk in the workshop on User Interaction, Immersive Storytelling, 3D Modeling and Character Animation

Nicolas Bonneel

Status: Researcher

Institution of origin: CNRS

Country: France

Dates: November, 12-14th

Context of the visit: Group retreat

Mobility program/type of mobility: Invited talk

10.3.2 ANR projects

ANR GLACIS

Participants: Adrien Bousseau, Anran Qi.

[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. In particular, GraphDeco will work on analysing sketches of data visualizations to translate them into parametric graphical objects that can be binded to data.

ANR INTERRA

Participants: Guillaume Cordonnier, Aryamaan Jain, Petros Tzathas.

[Project description on anr.fr](#)

Title: Inverse Control of Physically Consistent Terrains

Duration: From February 1, 2023 to January 31, 2027

Partners:

- Inria Centre de Recherche Inria Sophia Antipolis - Méditerranée, France

Inria contact: Guillaume Cordonnier

Coordinator: Guillaume Cordonnier

Summary: In a world where digital exchanges drive a pressing need for virtual environments, a challenge lies in the authoring of the root of these synthetic worlds: the mountains, plains, and other landforms concatenated and represented as terrains. This problem is notoriously difficult because terrains result from the interplay of physical events over geological time scales. This project aims to explore the inversion of simulation parameters as a novel paradigm for terrain generation in virtual worlds, combining geological consistency, natural diversity, and expressive user control for the first time.

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: selection

Chair of conference program committees

- Adrien Bousseau was co-chair of the Eurographics 2023 State-of-the-art-reports program.

Member of the conference program committees

- Guillaume Cordonnier was a member of the program committees of Computer Animation and Social Agents (CASA), Journées Françaises d'Informatique Graphique (JFIG) and Eurographics short papers
- George Drettakis was a member of the program committee of Eurographics Symposium on Rendering (EGSR'23).
- Adrien Bousseau was a member of the program committee of SIGGRAPH 2023 and Shape Modeling International 2023.

Reviewer

- Guillaume Cordonnier was a reviewer for ACM SIGGRAPH, SIGGRAPH Asia and Eurographics.
- George Drettakis was a reviewer for ACM SIGGRAPH and SIGGRAPH Asia.
- Adrien Bousseau was a reviewer for ACM Conference on Human Factors in Computing Systems (CHI), ACM SIGGRAPH Asia.
- Georgios Kopanas was a reviewer for Eurographics and CVPR.
- Emilie Yu was reviewer for ACM Symposium on User Interface Software and Technology (UIST) 2023, Pacific Graphics 2023, SIGGRAPH Posters 2023, ACM Conference on Human Factors in Computing Systems (CHI) 2024, and Eurographics 2024.

11.1.2 Journal

Member of the editorial boards

- Adrien Bousseau and George Drettakis are associate editors of Computer Graphics Forum.

Reviewer - reviewing activities

- Adrien Bousseau was a reviewer for ACM Transactions on Graphics and IEEE Transactions on Visualization and Computer Graphics.

11.1.3 Invited talks

- Guillaume Cordonnier gave an invited talk for the 20 years anniversary of Faculty of Earth and Environmental Sciences at University of Lausanne. He also gave invited talks in the Geosciences department at University of Rennes and at the Inria-DFKI European Summer School on AI (IDESSAI 2023).
- George Drettakis gave a talk at the meeting of the French Working Group on Graphics and VR in Lyon in May, and at the Athens University of Economics and Business (CS department) in December. He also gave a talk at Stanford University in August together with Georgios Kopanas.
- Adrien Bousseau gave an invited keynote at Smart Tools and Applications in Graphics 2023 in Matera, Italy.
- Georgios Kopanas gave multiple talks including to: NVIDIA, Google, MPI, Berkeley, Stanford.

11.1.4 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.

11.1.5 Scientific expertise

- Adrien Bousseau evaluated a grant proposal for U.S.-Israel Binational Science Foundation.
- Adrien Bousseau was a committee member for the Gilles Kahn Ph.D. award of the Société Informatique de France.
- George Drettakis wrote several tenure and other promotion letters for candidates in North America and Europe.

11.1.6 Research administration

- Guillaume Cordonnier is a member of the direction committee of the working group on computer graphics, geometry, virtual reality and visualization (GT CNRS IGRV) and chairs the Ph.D. award given by this working group.
- 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](#).
- George 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.

11.1.7 Supervision

- Georgios Kopanas, supervised by George Drettakis, defended in December 2023
- Emilie Yu, supervised by Adrien Bousseau, defended in December 2023
- Nicolas Rosset, co-supervised by Guillaume Cordonnier and Adrien Bousseau since October 2021.
- Capucine Nghiem, co-supervised by Adrien Bousseau with Theophanis Tsandilas since October 2021.
- Nicolás Violante, supervised by George Drettakis since October 2022.
- Yohan Poirier-Ginter, co-supervised by George Drettakis with Jean Francois Lalonde from Univ. Laval since September 2023.
- Petros Tzathas, co-supervised by George Drettakis and Guillaume Cordonnier since November 2023.
- Panagiotis Papantonakis, supervised by George Drettakis since Septembre 2023.
- Berend Baas, supervised by Adrien Bousseau since December 2023.
- Aryamaan Jain, supervised by Guillaume Cordonnier since October 2023.

11.1.8 Juries

- G. Drettakis was a member of the Ph.D. committee of Lois Paulin (Lyon) and Andrea Maggiordomo (Milano).

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] 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://inria.hal.science/hal-03594357>.
- [3] 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>.
- [4] 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>.
- [5] 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>.
- [6] 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/revues/Basilic/2018/HPPFDB18>.
- [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] G. Cordonnier, G. Jouvét, A. Peytavie, J. Braun, M.-P. Cani, B. Benes, E. Galin, E. Guérin and J. Gain. ‘Forming Terrains by Glacial Erosion’. In: *ACM Transactions on Graphics* 42.4 (26th July 2023), pp. 1–14. DOI: [10.1145/3592422](https://doi.org/10.1145/3592422). URL: <https://inria.hal.science/hal-04090644>.
- [12] S. Diolatzis, J. Novák, F. Rousselle, J. Granskog, M. Aittala, R. Ramamoorthi and G. Drettakis. ‘MesoGAN: Generative Neural Reflectance Shells’. In: *Computer Graphics Forum* (2023). DOI: [10.1111/cgf.14846](https://doi.org/10.1111/cgf.14846). URL: <https://inria.hal.science/hal-04105889>.
- [13] C. Jambon, B. Kerbl, G. Kopanas, S. Diolatzis, T. Leimkühler and G. Drettakis. ‘NeRFshop: Interactive Editing of Neural Radiance Fields’. In: *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 6.1 (May 2023). DOI: [10.1145/nmnnnnn.nmnnnnn](https://doi.org/10.1145/nmnnnnn.nmnnnnn). URL: <https://inria.hal.science/hal-04027538>.

- [14] G. Jouvet and G. Cordonnier. ‘Ice-flow model emulator based on physics-informed deep learning’. In: *Journal of Glaciology* (2023), pp. 1–15. DOI: [10.1017/jog.2023.73](https://doi.org/10.1017/jog.2023.73). URL: <https://hal.science/hal-04232949>.
- [15] B. Kerbl, G. Kopanas, T. Leimkühler and G. Drettakis. ‘3D Gaussian Splatting for Real-Time Radiance Field Rendering’. In: *ACM Transactions on Graphics* 42.4 (26th July 2023), pp. 1–14. DOI: [10.1145/3592433](https://doi.org/10.1145/3592433). URL: <https://inria.hal.science/hal-04088161>.
- [16] A. Petitjean, Y. Poirier-Ginter, A. Tewari, G. Cordonnier and G. Drettakis. ‘ModalNeRF: Neural Modal Analysis and Synthesis for Free-Viewpoint Navigation in Dynamically Vibrating Scenes’. In: *Computer Graphics Forum* 42.4 (2023), e14888. DOI: [10.1111/cgf.14888](https://doi.org/10.1111/cgf.14888). URL: <https://hal.science/hal-04131503>.
- [17] N. Rosset, G. Cordonnier, R. Duvigneau and A. Bousseau. ‘Interactive design of 2D car profiles with aerodynamic feedback’. In: *Computer Graphics Forum* 42.2 (8th Feb. 2023). DOI: [10.1111/cgf.14772](https://doi.org/10.1111/cgf.14772). URL: <https://inria.hal.science/hal-03975369>.
- [18] P. Voglreiter, B. Kerbl, A. Weinrauch, J. H. Mueller, T. Neff, M. Steinberger and D. Schmalstieg. ‘Trim Regions for Online Computation of From-Region Potentially Visible Sets’. In: *ACM Transactions on Graphics* 42.4 (Aug. 2023), pp. 1–15. DOI: [10.1145/3592434](https://doi.org/10.1145/3592434). URL: <https://inria.hal.science/hal-04273528>.
- [19] E. Yu, K. Blackburn-Matzen, C. Nguyen, O. Wang, R. Habib Kazi and A. Bousseau. ‘VideoDoodles: Hand-Drawn Animations on Videos with Scene-Aware Canvases’. In: *ACM Transactions on Graphics* 42.4 (Aug. 2023). DOI: [10.1145/3592413](https://doi.org/10.1145/3592413). URL: <https://inria.hal.science/hal-04093195>.

International peer-reviewed conferences

- [20] G. Kopanas and G. Drettakis. ‘Improving NeRF Quality by Progressive Camera Placement for Unrestricted Navigation in Complex Environments’. In: *Eurographics - The European Association for Computer Graphics*. VMV 2023 - Vision, Modeling, and Visualization. Braunschweig, Germany, 27th Sept. 2023. URL: <https://inria.hal.science/hal-04182002>.