

RESEARCH CENTRE

Inria Centre at Université Côte
d'Azur

2024

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	10
7 New software, platforms, open data	10
7.1 New software	10
7.1.1 sibr-core	10
7.1.2 3DGaussianSplats	10
7.1.3 NerfShop	11
7.1.4 H3DGS	11
7.1.5 DiffRelightGS	12
7.1.6 3DLayers	12
7.1.7 VideoDoodles	12
7.1.8 pySBM	13
7.1.9 pyLowStroke	13
7.1.10 AnaTerrains	13
7.1.11 Fastflow	14
8 New results	14
8.1 Computer-Assisted Design with Heterogeneous Representations	14
8.1.1 3D-Layers: Bringing Layer-Based Color Editing to VR Painting	14
8.1.2 CADTalk: An Algorithm and Benchmark for Semantic Commenting of CAD Programs	14
8.1.3 DiffCSG: Differentiable CSG via Rasterization	15
8.1.4 Single-Image SVBRDF Estimation with Learned Gradient Descent	16
8.1.5 STIVi: Turning Perspective Sketching Video into Interactive Tutorials	17
8.1.6 Presentation Sketches in Product Design	18
8.1.7 Reconstructing CAD Programs from Hand-drawn Sketches	19
8.1.8 Surface Dissections: Waste Reuse through Shape Approximations	19
8.1.9 Recovering Data from Hand-Drawn Infographics	19
8.1.10 Computational garment reuse	19
8.2 Graphics with Uncertainty and Heterogeneous Content	20
8.2.1 Learning Images Across Scales Using Adversarial Training	20
8.2.2 N-Dimensional Gaussians for Fitting of High Dimensional Functions	20
8.2.3 MatUp: Repurposing Image Upsamplers for SVBRDFs	21
8.2.4 A Diffusion Approach to Radiance Field Relighting using Multi-Illumination Synthesis	22
8.2.5 A Hierarchical 3D Gaussian Representation for Real-Time Rendering of Very Large Datasets	23

8.2.6	Reducing the Memory Footprint of 3D Gaussian Splatting	24
8.2.7	Physically-Based Lighting for 3D Generative Models of Cars	24
8.2.8	Studying the Effect of Volumetric Rendering Approximations for 3D Gaussian Splatting	25
8.2.9	SVBRDF Prediction and Generative Image Modeling for Appearance Modeling of 3D Scenes	26
8.2.10	Physical Reflections in Gaussian Splatting	26
8.2.11	Incremental Scalable Reconstruction	26
8.2.12	Caching Global Illumination using Gaussian Primitives	27
8.2.13	Adding Textures to Gaussian Splatting	27
8.2.14	Dynamic Gaussian Splatting	27
8.2.15	Splat and Replace: 3D Reconstruction with Repetitive Elements	27
8.3	Physical Simulation for Graphics	28
8.3.1	Volcanic Skies: coupling explosive eruptions with atmospheric simulation to create consistent skiescapes	28
8.3.2	Fast simulation of viscous lava flow using Green's functions as a smoothing kernel	28
8.3.3	Unerosion: simulating terrain evolution back in time	29
8.3.4	CHONK 1.0: landscape evolution framework: cellular automata meets graph theory	30
8.3.5	FastFlow: GPU Acceleration of Flow and Depression Routing for Landscape Simulation	30
8.3.6	Efficient Debris-flow Simulation for Steep Terrain Erosion	32
8.3.7	Learning Based Infinite Terrain Generation with Level of Detailing	32
8.3.8	Physically-based analytical erosion for fast terrain generation	33
8.3.9	Windblown sand around obstacles- simulation and validation of deposition patterns	34
8.3.10	Sandstone Erosion	34
8.3.11	Neural surrogate model for wind prediction above a terrain	35
8.3.12	Extracting terrain characteristics using erosion laws	35
9	Partnerships and cooperations	36
9.1	International research visitors	36
9.1.1	Visits of international scientists	36
9.1.2	Visits to international teams	38
9.2	European initiatives	39
9.2.1	Horizon Europe	39
9.2.2	H2020 projects	40
9.3	National initiatives	41
9.3.1	Visits of French scientists	41
9.3.2	ANR projects	42
10	Dissemination	44
10.1	Promoting scientific activities	44
10.1.1	Scientific events: selection	44
10.1.2	Journal	44
10.1.3	Invited talks	44
10.1.4	Leadership within the scientific community	45
10.1.5	Scientific expertise	45
10.1.6	Research administration	45
10.2	Teaching - Supervision - Juries	45
10.2.1	Teaching	45
10.2.2	Supervision	45
10.2.3	Juries	46
10.3	Popularization	46
10.3.1	Participation in Live events	46

11 Scientific production	46
11.1 Major publications	46
11.2 Publications of the year	47
11.3 Cited publications	49

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]
- Adrien Bousseau [INRIA, Senior Researcher]
- Guillaume Cordonnier [INRIA, Researcher]
- Andreas Meuleman [INRIA, Starting Research Position, from Oct 2024]

Post-Doctoral Fellows

- Melike Aydinlilar [INRIA, Post-Doctoral Fellow, from Jun 2024]
- Alban Gauthier [INRIA, Post-Doctoral Fellow]
- Georgios Kopanas [INRIA, Post-Doctoral Fellow, until Jan 2024]
- Andreas Meuleman [INRIA, Post-Doctoral Fellow, until Sep 2024]
- Anran Qi [INRIA, Post-Doctoral Fellow]
- Marzia Riso [INRIA, Post-Doctoral Fellow, from Oct 2024]
- Emilie Yu [INRIA, until Apr 2024]

PhD Students

- Berend Baas [INRIA]
- Aryamaan Jain [INRIA]
- Henro Kriel [INRIA, from Sep 2024]
- Alexandre Lanvin [INRIA, from Dec 2024]
- Panagiotis Papantonakis [INRIA]
- Yohan Poirier-Ginter [UNIV LAVAL QUEBEC]
- Nicolas Rosset [INRIA]
- Petros Tzathas [INRIA]
- Nicolas Violante [INRIA]

Technical Staff

- Alexandre Lanvin [INRIA, Engineer, until Nov 2024]
- Ishaan Shah [INRIA, Engineer, from Dec 2024]
- Jiayi Wei [INRIA, Engineer, from Sep 2024]

Interns and Apprentices

- David Behrens [INRIA, Intern, from Jun 2024 until Aug 2024]
- Juan Sebastian Osorno Bolivar [INRIA, Intern, from May 2024 until Sep 2024]
- Vishal Pani [INRIA, Intern, until Feb 2024]
- Clement Remy [ENS DE LYON, Intern, from Jun 2024 until Jul 2024]

Administrative Assistant

- Sophie Honnorat [INRIA]

Visiting Scientists

- Frederic Durand [MIT, until Aug 2024]
- Gilda Manfredi [University of Basilicata, until Mar 2024]
- Eric Paquette [ETS MONTREAL, from Oct 2024 until Nov 2024]

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.

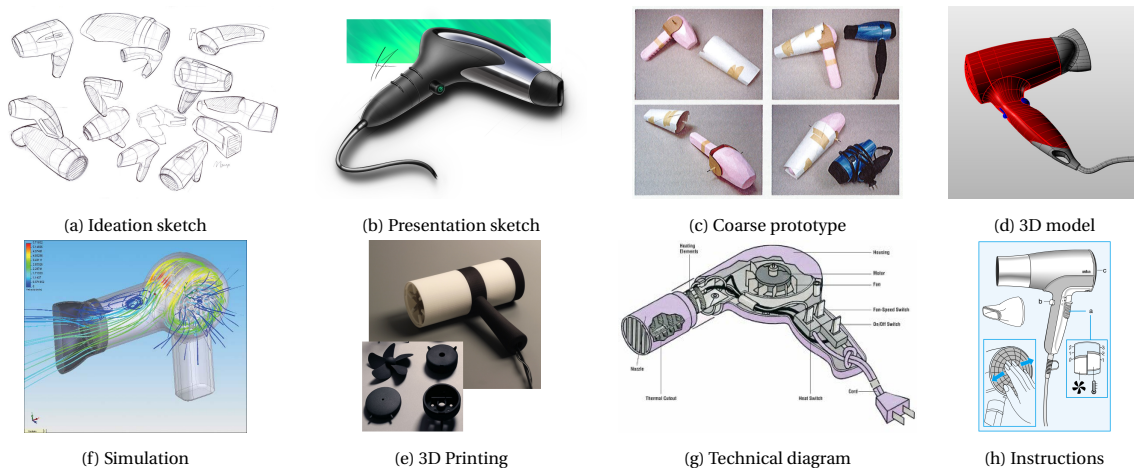


Figure 1: Various representations of a hair dryer at different stages of the design process. Image source, in order: c-maeng on deviantart.com, shauntur on deviantart.com, "Prototyping and Modelmaking for Product Design" Hallgrimsson, B., Laurence King Publishers, 2012, samsher511 on turbosquid.com, 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.



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

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

We strive to develop methods in which heterogeneous – and often uncertain – data can be handled automatically in CG with a principled methodology. Our main focus is on *rendering* in CG, including dynamic scenes (video/animations) (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 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 developement 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 (3DGS) [33] 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, film and television production, virtual and extended reality, real-estate visualization and others. The 3DGS technology significantly reduces the computation time and thus computational resources required compared to the previous state of the art in 3D reconstruction/novel view synthesis that was the Neural Radiance Fields method (for typical scenes, 40 min instead of 48 hours of GPU time).

6 Highlights of the year

The most significant highlight of the year was the success of the 3D Gaussian Splatting method [33] which has had unprecedented impact both scientifically (already more than 2500 citations on google scholar since 2023) and in terms of technology transfer (around 11 licenses commercialized by Inria, see Sec. 5.2).

6.1 Awards

- George Drettakis received the Eurographics Outstanding Career Award.
- George Drettakis received the Grand Prix Inria - Academie de Sciences.
- Emilie Yu received the Ph.D. award from GdR IG-RV and the Gilles Kahn Ph.D. award from Société informatique de France (SIF).
- Yorgos Kopanas received the EDSTIC Doctoral Prize in Computer Science.
- Our paper *FastFlow: GPU Acceleration of Flow and Depression Routing for Landscape Simulation* [14] received the Best Paper Award at Pacific Graphics 2024.

- Our paper *Unerosion: Simulating Terrain Evolution Back in Time* [23] received an Honorable Mention to the Best Paper Award at the ACM SIGGRAPH / Eurographics Symposium on Computer Animation 2024.
- Our paper *CADTalk: An Algorithm and Benchmark for Semantic Commenting of CAD Programs* [31] was selected as a highlight at CVPR (top 10%).

6.2 Press release

- Article on 3D Gaussian Splatting: "[Une technique revolutionne la creation des scenes en 3D](#)", Le Monde, May 2024
- 3D Gaussian Splatting mentioned in the "[Best of Science News 2024](#)" in Le Monde, December 2024,
- [3D Mag website interview](#), October 2024.
- Inria media department articles on [NERPHYS](#): and on the [Grand Prix Inria - Academie de Sciences](#).
- Adrien Bousseau was interviewed for [People of ACM](#).

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>

This year several improvements were added as part of the 3D Gaussian Splatting support.

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.

Contact: George Drettakis

7.1.2 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 that 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.

We have provided several updates to the software this year, most importantly integrating new features for better quality and speed of optimization.

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

Contact: George Drettakis

Participants: Georgios Kopanas, Bernhard Kerbl, Thomas Leimkuhler

7.1.3 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.

The software has seen extensive usage in followup papers during this year.

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

Contact: George Drettakis

7.1.4 H3DGS

Name: Hierarchical 3D Gaussian Splatting

Keywords: 3D modeling, 3D rendering, Differentiable Rendering

Scientific Description: Implementation of the SIGGRAPH 2024 paper A Hierarchical 3D Gaussian Representation for Real-Time Rendering of Very Large Datasets, project page <https://repo-sam.inria.fr/fungraph/hierarchical-3d-gaussians/>

Functional Description: Implementation of the SIGGRAPH 2024 paper A Hierarchical 3D Gaussian Representation for Real-Time Rendering of Very Large Datasets, project page <https://repo-sam.inria.fr/fungraph/hierarchical-3d-gaussians/>

URL: <https://repo-sam.inria.fr/fungraph/hierarchical-3d-gaussians/>

Contact: George Drettakis

Participants: Georgios Kopanas, Alexandre Lanvin, Bernhard Kerbl, Andreas Meuleman, George Drettakis, Michael Wimmer

Partner: Technische Universität Wien

7.1.5 DiffRelightGS

Name: A Diffusion Approach to Radiance Field Relighting using Multi-Illumination Synthesis

Keywords: 3D, 3D reconstruction, 3D rendering, Artificial intelligence, Machine learning

Functional Description: A method to create relightable 3D radiance fields from single-illumination data by exploiting priors extracted from 2D image diffusion models.

URL: <https://repo-sam.inria.fr/fungraph/generative-radiance-field-relighting/>

Publication: hal-04628346v1

Contact: George Drettakis

Participants: Yohan Poirier-Ginter, Alban Gauthier, Julien Philip, Jean-Francois Lalonde, George Drettakis

Partner: Université Laval

7.1.6 3DLayers

Name: VR painting system with layers

Keywords: Virtual reality, 3D, Painting

Functional Description: This is the source code for the prototype implementation of the research paper: "3D-Layers: Bringing Layer-Based Color Editing to VR Painting", Emilie Yu, Fanny Chevalier, Karan Singh and Adrien Bousseau, ACM Transactions on Graphics (SIGGRAPH) - 2024

This is a Unity project that implements a simple VR application compatible with Quest 2/3/Pro headsets. The project features:

- A VR app with basic 3D painting features (painting tube strokes, stroke deletion and transformation, color palette, undo/redo).
- A UI in the VR app to create, paint in, and edit shape and appearance layers, as described in the 3DLayers paper. We have a basic menu UI for users to visualize and navigate in the layer hierarchy.
- A basic in-Unity visualizer for paintings created with our system. It enables users to view and render still frames or simple camera path animations. We used it to create all results in the paper/video.

URL: <https://em-yu.github.io/research/3dlayers/>

Contact: Emilie Yu

7.1.7 VideoDoodles

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

Keywords: 3D web, 3D, 2D animation, 3D animation, Visual tracking

Scientific Description: Implementation for Siggraph 2023 paper VideoDoodles: Hand-Drawn Animations on Videos with Scene-Aware Canvases

Functional Description: 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. 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, perspective deformations and occlusions 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.

Our system is composed of the following elements: * A preprocessing library to convert depth, camera and motion data into data formats compatible with our system * A collection of Python scripts that can be used either offline to execute 3D point tracking (with the possibility to specify 1 to N keyframes) , or as a backend server to our interactive frontend UI * A web UI to author video doodles. It features a renderer capable of displaying the composited video doodles , an editing interface to keyframe canvases , a sketching interface to create frame-by-frame animations on canvases , basic export capabilities.

URL: <https://em-yu.github.io/research/videodoodles/>

Contact: Emilie Yu

7.1.8 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/>
The software is currently part of an ERC PoC development cycle for the creation of a Blender plugin.

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

7.1.9 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.

The software is currently part of an ERC PoC development cycle for the creation of a Blender plugin.

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

Contact: Adrien Bousseau

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

7.1.10 AnaTerrains

Name: Physically-based analytical erosion for fast terrain generation

Keywords: Analytic model, Terrain, Simulation

Functional Description: Source code for the paper "Physically-based analytical erosion for fast terrain generation". This code demonstrates the use of analytical solutions of erosion laws to generate large-scale mountain ranges instantly.

Contact: Guillaume Cordonnier

7.1.11 Fastflow

Name: GPU Acceleration of Flow and Depression Routing for Landscape Simulation

Keywords: Flow routing, Landscape, GPU

Functional Description: Library for the fast computation of flow and depression routing on the GPU for numerical simulations of landscapes in hydrology and geomorphology. This code enables the computation of flow-related properties such as the discharge over large Digital Elevation Models (flow routing). It solves the problem that local minima in topography interrupt the flow path (depression routing). The code is optimized for the GPU, resulting in fast execution time for large domains.

URL: <https://gitlab.inria.fr/landscapes/fastflow>

Contact: Guillaume Cordonnier

8 New results

8.1 Computer-Assisted Design with Heterogeneous Representations

8.1.1 3D-Layers: Bringing Layer-Based Color Editing to VR Painting

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

The ability to represent artworks as stacks of layers is fundamental to modern graphics design, as it allows artists to easily separate visual elements, edit them in isolation, and blend them to achieve rich visual effects. Despite their ubiquity in 2D painting software, layers have not yet made their way to VR painting, where users paint strokes directly in 3D space by gesturing a 6-degrees-of-freedom controller. But while the concept of a stack of 2D layers was inspired by real-world layers in cell animation, what should 3D layers be? We propose to define 3D-Layers as groups of 3D strokes, and we distinguish the ones that represent 3D geometry from the ones that represent color modifications of the geometry. We call the former substrate layers and the latter appearance layers. Strokes in appearance layers modify the color of the substrate strokes they intersect. Thanks to this distinction, artists can define sequences of color modifications as stacks of appearance layers, and edit each layer independently to finely control the final color of the substrate. We have integrated 3D-Layers into a VR painting application and we evaluated its flexibility and expressiveness by conducting a usability study with experienced VR artists. Figure 3 details an artwork created with our system.

This work is a collaboration with Fanny Chevalier and Karan Singh from the University of Toronto and was initiated during the visit of Emilie Yu at University of Toronto in May-August 2023 funded by MITACS. It was published in ACM Transactions on Graphics and presented at SIGGRAPH 2024 [24].

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

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

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 (see Fig. 4). We solve the



Figure 3: Our VR painting system allows artists to achieve rich, editable coloring effects using 3D-Layers. Starting with substrate layers (a) that define the geometry and basic colors of the scene, users can stack multiple appearance layers (b) that are composited onto the substrate to produce the final 3D scene (c). Importantly, strokes painted in appearance layers only recolor the substrate strokes they intersect (b, intersections highlighted with a yellow boundary), which avoids the need to position the appearance strokes precisely on the surface of the substrate. In this example, we used appearance strokes to add texture details (white bands on the lighthouse, dark lines on the house and rocks), to paint shadows (lighthouse, rocks), to depict translucency (semi-transparent water painted on the rocks and seabed, subject to a vertical gradient in opacity over the rocks).

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 report an 83.24% accuracy on the new CADTalk dataset.

This work is a collaboration with Haocheng Yuan, Jing Xu and Changjian Li from University of Edinburgh, Hao Pan from Microsoft Research Asia, and Niloy J. Mitra from University College London and Adobe Research. It was published at the Conference on Computer Vision and Pattern Recognition (CVPR) where it was selected as a highlight (top 10%) [31].

8.1.3 DiffCSG: Differentiable CSG via Rasterization

Participants: Adrien Bousseau, Haocheng Yuan (*University of Edinburgh*), Hao Pan (*Microsoft Research Asia*), Chengquan Zhang (*Nanjing University*), Niloy J. Mitra (*University College London, Adobe Research*), Changjian Li (*University of Edinburgh*).

Differentiable rendering is a key ingredient for inverse rendering and machine learning, as it allows to optimize scene parameters (shape, materials, lighting) to best fit target images. Differentiable rendering requires that each scene parameter relates to pixel values through differentiable operations. While 3D mesh rendering algorithms have been implemented in a differentiable way, these algorithms do not directly extend to Constructive-Solid-Geometry (CSG), a popular parametric representation of shapes, because the underlying boolean operations are typically performed with complex black-box mesh-processing libraries. We present an algorithm, DiffCSG, to render CSG models in a differentiable manner. Our algorithm builds upon CSG rasterization, which displays the result of boolean operations between primitives without explicitly computing the resulting mesh and, as such, bypasses black-box mesh processing. We describe how to implement CSG rasterization within a differentiable rendering pipeline, taking special care to apply antialiasing along primitive intersections to obtain gradients in such critical

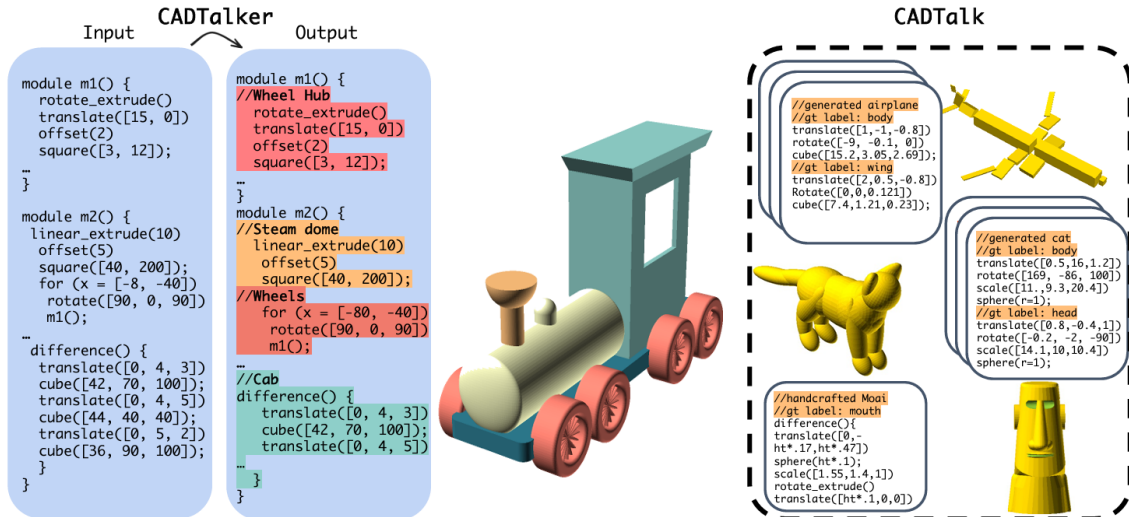


Figure 4: Given a CAD program as input, our algorithm – CADTalker – automatically generates comments before each code block to describe the shape part that is generated by the block (left). We evaluate our algorithm on a new dataset of commented CAD programs – CADTalk – that contains both human-made and machine-made CAD programs (right).

areas. Our algorithm is simple and fast, can be easily incorporated into modern machine learning setups, and enables a range of applications for computer-aided design, including direct and image-based editing of CSG primitives (Fig. 5).

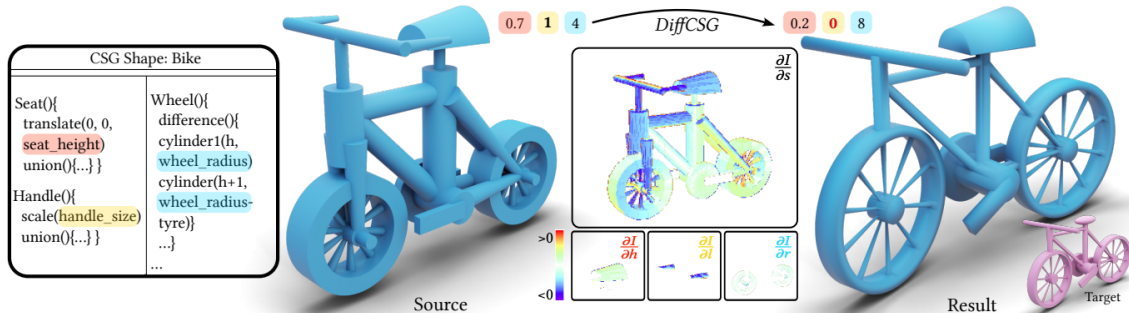


Figure 5: Given the CSG model of a bike (left), our DiffCSG renders the corresponding shape in a differentiable manner, such that its continuous parameters can be optimized to best match multi-view renderings of a target shape (right, pink). Our solution builds upon differentiable rasterization to compute image gradients with respect to CSG parameters (middle). Here we visualize the per-pixel gradient contribution for the global scale parameter s , the seat height h , the handle size l , and the wheel radius r (we cropped the gradient visualizations around the areas of interest for h , l and r). In this example, the optimization decreased the seat height, increased the wheel radius, and made the handles vanish by setting their size to 0 (top middle). The optimization also adjusted the orientation of the pedals.

This work is a collaboration with Haocheng Yuan and Changjian Li from University of Edinburgh, Chengquan Zhang from Nanjing University, Hao Pan from Microsoft Research Asia, and Niloy J. Mitra from University College London and Adobe Research. It was published at the conference track of SIGGRAPH Asia [30].

8.1.4 Single-Image SVBRDF Estimation with Learned Gradient Descent

Participants: Adrien Bousseau, Xuejiao Luo (*Technical University of Delft*), Leonardo Scandolo (*Technical University of Delft*), Elmar Eise-
mann (*Technical University of Delft*).

Recovering spatially-varying materials from a single photograph of a surface is inherently ill-posed, making the direct application of a gradient descent on the reflectance parameters prone to poor minima. Recent methods leverage deep learning either by directly regressing reflectance parameters using feed-forward neural networks or by learning a latent space of SVBRDFs (Spatially-Varying Bidirectional Reflection Distribution Functions) using encoder-decoder or generative adversarial networks followed by a gradient-based optimization in latent space. The former is fast but does not account for the likelihood of the prediction, i.e., how well the resulting reflectance explains the input image. The latter provides a strong prior on the space of spatially-varying materials, but this prior can hinder the reconstruction of images that are too different from the training data. Our method combines the strengths of both approaches. We optimize reflectance parameters to best reconstruct the input image using a recurrent neural network, which iteratively predicts how to update the reflectance parameters given the gradient of the reconstruction likelihood. By combining a learned prior with a likelihood measure, our approach provides a maximum a posteriori estimate of the SVBRDF. Our evaluation shows that this learned gradient-descent method achieves state-of-the-art performance for SVBRDF estimation on synthetic and real images (Fig. 6).

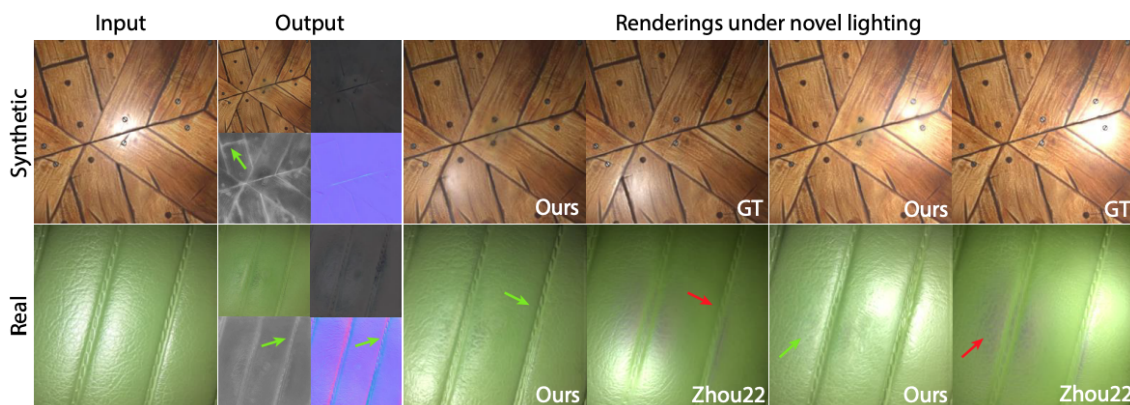


Figure 6: Given a single flash picture as input, our learned gradient descent algorithm solves for SVBRDF maps that closely reproduce the input, while generalizing well to new view and light configurations. We demonstrate the effectiveness of our approach by comparing its output to ground-truth synthetic data (top) and to relit real images produced by state-of-the-art methods (bottom). Note how our method recovers intricate geometric details in the normal map, inpaints small saturated highlights with plausible material properties, and propagates roughness information away from the highlight.

This work is a collaboration with Xuejiao Luo, Leonardo Scandolo and Elmar Eisemann from the Technical University of Delft and was initiated during Adrien Bousseau’s sabbatical visit there in 2022-2023. It was published in Computer Graphics Forum and presented at Eurographics 2024 [16].

8.1.5 STIVi: Turning Perspective Sketching Video into Interactive Tutorials

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

For design and art enthusiasts who seek to enhance their skills through instructional videos, following drawing instructions while practicing can be challenging. STIVi presents perspective drawing

demonstrations and commentary of prerecorded instructional videos as interactive drawing tutorials that students can navigate and explore at their own pace (see Fig. 7). Our approach involves a semi-automatic pipeline to assist instructors in creating STIVi content by 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. We evaluated STIVi's interactive tutorials against a regular video player. We observed that our interface supports non-linear learning styles by providing students alternative paths for following and understanding drawing instructions.




Figure 7: STIVi augments instructional sketching videos to help students follow perspective drawing instructions and practice at their own pace. STIVi's user interface links keywords from the video transcript to elements in the drawing, such as lines and planes, making it easier to navigate and familiarize with key concepts covered in the instructions. It also allows students to interact with perspective constructions (e.g., vanishing lines highlighted on the video) and understand how the geometry of depicted shapes would change if they were drawn from different viewpoints.

This work is a collaboration with Theophanis Tsandilas from Inria, LISN, Université Paris-Saclay, Jan Willem Hoftijzer and Mark Sypesteyn from TU Delft and Maneesh Agrawala from Stanford University. It was presented at the Graphics Interface 2024 conference and published in the conference's proceedings [28].

8.1.6 Presentation Sketches in Product Design

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

Sketching is a core skill for industrial designers. While substantial research has studied sketching for product development, our work focuses on the activity of product presentation. We study product designer's practice of sketching for presentation, and more generally their presentation preparation workflow through three lenses: (1) We analyzed a corpus of presentation sketches to identify their visual characteristics and storytelling techniques. (2) We interviewed nine designers about their workflow and decision-making related to a recent presentation of theirs. (3) We also invited these designers to brainstorm on a sketch-based dynamic presentation authoring tool to explore the potential for novel sketch-based tools supporting their presentation authoring workflow.

This work is a collaboration with Theophanis Tsandilas from Inria, LISN, Université Paris-Saclay, Jan Willem Hoftijzer and Mark Sypesteyn from TU Delft.

8.1.7 Reconstructing CAD Programs from Hand-drawn Sketches

Participants: Henro Kriel, Adrien Bousseau, Gilda Manfredi (*University of Basilicata*), Daniel Ritche (*Brown University*).

We present a system that takes a concept sketch and reconstructs the depicted CAD model. We segment strokes into CAD operations, label these operations, and fit their parameters. We propose a graph neural network architecture to segment and label concept sketches for fitting, and we leverage sketch processing techniques, such as 3D reconstruction and cycle detection, to ease this prediction task.

This project is a collaboration with Gilda Manfredi, who visited our lab for 6 months, and Daniel Ritchie from Brown.

8.1.8 Surface Dissections: Waste Reuse through Shape Approximations

Participants: Berend Baas, Adrien Bousseau, David Bommès (*University of Bern*).

The design and manufacturing industry faces an urgent need to shift away from linear methods of production to more circular production models. In order to aid in this effort, the current generation of computational design tools needs to be rethought in order to accommodate reused materials for reuse scenarios. Motivated by this use case, this research project proposes a new computational inverse problem we label Surface Dissections, generalizing dissection puzzles to curved domains, with the intent to approximate target designs using reclaimed material. We propose an interactive method in an effort to solve this problem.

This is ongoing work with David Bommès from the University of Bern, Switzerland.

8.1.9 Recovering Data from Hand-Drawn Infographics

Participants: Anran Qi, Theophanis Tsandilas (*Inria ExSitu*), Ariel Shamir (*Reichman University*), Adrien Bousseau.

Data collection and visualization have long been perceived as activities reserved for experts. However, by drawing simple geometric figures –or glyphs – anyone can easily record and visualize their own data. Still, the resulting hand-drawn infographics do not provide direct access to the raw data values being visualized, hindering subsequent digital editing of both the data and the glyphs. We present a method to recover data values from glyph-based hand-drawn infographics. Given a visualization in a bitmap format and a user-specified parametric template of the glyphs in that visualization, we leverage deep neural networks to detect and localize all glyphs, and estimate the data values they represent. This reverse-engineering procedure effectively disentangles the depicted data from its visual representation, enabling various editing applications, such as visualizing new data values or testing different visualizations of the same data.

This is ongoing work with Theophanis Tsandilas from Inria ExSitu and Ariel Shamir from Reichman University.

8.1.10 Computational garment reuse

Participants: Anran Qi, Maria Korosteleva (*Meshcapade*), Nico Pietroni (*University of Technology of Sydney*), Olga Sorkine-Hornung (*ETH Zurich*), Adrien Bousseau.

The fashion industry is experiencing mass production, mass consumption, and as a consequence it also produces mass pollution. According to the UN Environment Program, the fashion industry is responsible for about 10% of global carbon emissions. Despite this, only 1% of used clothes are recycled into new clothes due to a lack of technology for recovering virgin fibers. Garment upcycling, a creative practice that transforms old clothes, termed as the source, into a new item, termed as the target, is gaining popularity as a method of recycling that individuals can engage in daily. We propose a computational method to compute the fabrication plan that transforms from the source to the target. We formulate this fabrication plan computation as an optimization that seeks to reuse as much of the existing seams and hems as possible (to minimize fabrication cost), while deviating as little as possible from the target (to best preserve design intent).

This is ongoing work with Maria Korosteleva from Meshcapade, Nico Pietroni from the University of Technology of Sydney and Olga Sorkine-Hornung from ETH Zurich.

8.2 Graphics with Uncertainty and Heterogeneous Content

8.2.1 Learning Images Across Scales Using Adversarial Training

Participants: George Drettakis, Guillaume Cordonnier, Krzysztof Wolski (*MPI*), Adarsh Djeacoumar (*MPI*), Alireza Javanmardi (*MPI*), Hans-Peter Seidel (*MPI*), Christian Theobalt (*MPI*), Karol Myszkowski (*MPI*), Xingang Pan (*Nanyang Technological University, Singapour*), Thomas Leimkühler (*MPI*).

The real world exhibits rich structure and detail across many scales of observation. It is difficult, however, to capture and represent a broad spectrum of scales using ordinary images. We devise a novel paradigm for learning a representation that captures an orders-of-magnitude variety of scales from an unstructured collection of ordinary images. We treat this collection as a distribution of scale-space slices to be learned using adversarial training, and additionally enforce coherency across slices. Our approach relies on a multiscale generator with carefully injected procedural frequency content, which allows to interactively explore the emerging continuous scale space. Training across vastly different scales poses challenges regarding stability, which we tackle using a supervision scheme that involves careful sampling of scales. We show that our generator can be used as a multiscale generative model, and for reconstructions of scale spaces from unstructured patches. Significantly outperforming the state of the art, we demonstrate zoom-in factors of up to 256x at high quality and scale consistency (see Fig. 8).

This work is a collaboration with Thomas Leimkuehler, Krzysztof Wolski, Adarsh Djeacoumar, Alireza Javanmardi, Hans-Peter Seidel, Christian Theobalt, Karol Myszkowski from the Max-Planck-Institut für Informatik in Saarbruecken and Xingang Pan from NTU - Nanyang Technological University in Singapour; the work was started while Thomas was a postdoc at GRAPHDECO. It was published in ACM Transactions on Graphics, and presented at ACM SIGGRAPH [22].

8.2.2 N-Dimensional Gaussians for Fitting of High Dimensional Functions

Participants: Georgios Kopanas, Stavros Diolatzis (*Intel Labs*), Tobias Zirr (*Intel Labs*), Alexander Kuznetsov (*Intel Labs*), Anton Kaplanyan (*Intel Labs*).

In the wake of many new ML-inspired approaches for reconstructing and representing high-quality 3D content, recent hybrid and explicitly learned representations exhibit promising performance and quality characteristics. However, their scaling to higher dimensions is challenging, e.g. when accounting for




Figure 8: Given an unregistered collection of image patches depicting an environment at vastly different scales, our approach uses adversarial training to obtain continuous and coherent scale spaces. Here, we showcase the reconstructed scale space of images, from the overall structure (1x) to the finest detail (64x or 256x). Users can freely explore the scale space at interactive rates.

dynamic content with respect to additional parameters such as material properties, illumination, or time. We tackle these challenges with an explicit representation based on Gaussian mixture models. With our solutions, we arrive at efficient fitting of compact N-dimensional Gaussian mixtures and enable efficient evaluation at render time: For fast fitting and evaluation, we introduce a high-dimensional culling scheme that efficiently bounds N-D Gaussians, inspired by Locality Sensitive Hashing. For adaptive refinement yet compact representation, we introduce a loss-adaptive density control scheme that incrementally guides the use of additional capacity toward missing details. With these tools we can for the first time represent complex appearance that depends on many input dimensions beyond position or viewing angle within a compact, explicit representation optimized in minutes and rendered in milliseconds (see Fig. 9).

This work is a collaboration with S. Diolatzis, T. Zirr, Z. Kuznetsov and A. Kaplanyan from Intel. It was published in Transactions on Graphics, and presented in SIGGRAPH 2024 [25].

8.2.3 MatUp: Repurposing Image Upsamplers for SVBRDFs

Participants: Alban Gauthier, Bernhard Kerbl (*TU Wien*), Jérémy Levallois (*Adobe Research*), Robin Faury (*Adobe Research*), Jean-Marc Thiery (*Adobe Research*), Tamy Boubekeur (*Adobe Research*).

We propose MatUp, an upsampling filter for material super-resolution. Our method takes as input a low-resolution SVBRDF and upscales its maps so that their rendering under various lighting conditions fits upsampled renderings inferred in the radiance domain with pre-trained RGB upsamplers. We formulate our local filter as a compact Multilayer Perceptron (MLP), which acts on a small window of the input SVBRDF and is optimized using a sparsity-inducing loss defined over upsampled radiance at various locations. This optimization is entirely performed at the scale of a single, independent material. In doing




Figure 9: Our method optimizes N-Dimensional Gaussians to approximate high dimensional anisotropic functions in a few minutes, enabling the representation of surfaces (top) and volume radiance fields (bottom).

so, MatUp leverages the reconstruction capabilities acquired over large collections of natural images by pre-trained RGB models and provides regularization over self-similar structures. In particular, our lightweight neural filter avoids retraining complex architectures from scratch or accessing any large collection of low/high resolution material pairs - which do not actually exist at the scale RGB upsamplers are trained with. As a result, MatUp provides fine and coherent details in the upscaled material maps (see Fig. 10), as shown in an extensive evaluation we provide.



Figure 10: Given a low-resolution 128^2 input materials, our method produces 512^2 outputs preserving tileability and displaying sharper material structures compared to widely-used per-channel upsampling filters (e.g. bilinear or Lanczos) in less than a minute.

This work is a collaboration with Bernhard Kerbl from Technische Universität Wien and Jérémy Levallois, Robin Faury, Jean-Marc Thiery, and Tamy Boubekeur from Adobe Research. It was published in Computer Graphics Forum, and presented at Eurographics Symposium on Rendering (EGSR) 2024 [12].

8.2.4 A Diffusion Approach to Radiance Field Relighting using Multi-Illumination Synthesis

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

Relighting radiance fields is severely underconstrained for multi-view data, which is most often captured under a single illumination condition; it is especially hard for full scenes containing multiple objects. We introduce a method to create relightable radiance fields using such single-illumination data by exploiting priors extracted from 2D image diffusion models. We first fine-tune a 2D diffusion model on a multi-illumination dataset conditioned by light direction, allowing us to augment a single-illumination capture into a realistic – but possibly inconsistent – multi-illumination dataset from directly defined light directions. We use this augmented data to create a relightable radiance field represented by 3D Gaussian splats. To allow direct control of light direction for low-frequency lighting, we represent appearance with a multi-layer perceptron parameterized by light direction. To enforce multi-view consistency and overcome inaccuracies, we optimize a per-image auxiliary feature vector. We show results on synthetic and real multi-view data under single illumination, demonstrating that our method successfully exploits 2D diffusion model priors to allow realistic 3D relighting for complete scenes (see Fig. 11).




Figure 11: Our method produces relightable radiance fields directly from single-illumination multi-view dataset, by using priors from generative data in the place of an actual multi-illumination capture.

This work is a collaboration with Julien Philip (Adobe Research) and Jean-François Lalonde (Université Laval). It was published in the journal *Computer Graphics Forum* and presented at the Eurographics Symposium on Rendering (EGSR) 2024 [17].

8.2.5 A Hierarchical 3D Gaussian Representation for Real-Time Rendering of Very Large Datasets

Participants: Andreas Meuleman, Alexandre Lanvin, George Drettakis, Bernhard Kerbl (*TU Wien*), Georgios Kopanas (*Google*), Michael Wimmer (*TU Wien*).

Novel view synthesis has seen major advances in recent years, with 3D Gaussian splatting offering an excellent level of visual quality, fast training and real-time rendering. However, the resources needed for training and rendering inevitably limit the size of the captured scenes that can be represented with good visual quality. We introduce a hierarchy of 3D Gaussians that preserves visual quality for very large scenes, while offering an efficient Level-of-Detail (LOD) solution for efficient rendering of distant content with effective level selection and smooth transitions between levels. We introduce a divide-and-conquer approach that allows us to train very large scenes in independent chunks. We consolidate the chunks into

a hierarchy that can be optimized to further improve visual quality of Gaussians merged into intermediate nodes. Very large captures typically have sparse coverage of the scene, presenting many challenges to the original 3D Gaussian splatting training method; we adapt and regularize training to account for these issues. We present a complete solution, that enables real-time rendering of very large scenes and can adapt to available resources thanks to our LOD method. We show results for captured scenes with up to tens of thousands of images with a simple and affordable rig, covering trajectories of up to several kilometers and lasting up to one hour (see Fig. 12).

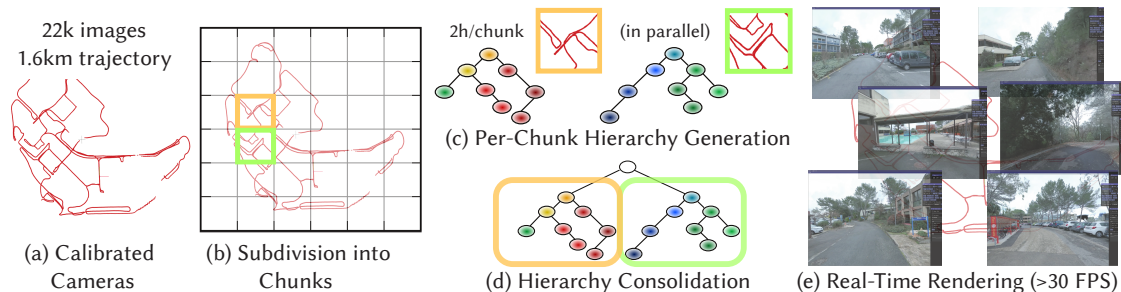


Figure 12: (a) Starting from thousands of calibrated cameras, covering a large area, we subdivide the scene into chunks (b). We introduce a 3D Gaussian Splatting hierarchy to allow efficient rendering of massive data, that we further optimize to enhance visual quality (c). We consolidate the hierarchies (d) enabling us to perform real-time rendering of very large datasets.

This work is a collaboration with Bernhard Kerbl and Michael Wimmer (TU Wien) and Georgios Kopanas (Google). It was published in Transaction on Graphics (TOG) and presented at Siggraph 2024 [15].

8.2.6 Reducing the Memory Footprint of 3D Gaussian Splatting

Participants: Panagiotis Papantonakis, Georgios Kopanas, Bernhard Kerbl (*TU Wien*), 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 for storing and transmission 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 the 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 times reduction in overall size on disk on the standard datasets we tested, along with a 1.7 speedup in rendering speed (see Fig. 13). We demonstrate our method on standard datasets and show how our solution results in significantly reduced download times when using the method on a mobile device.

This work is a collaboration with Bernhard Kerbl from TU Wien and was presented at the 2024 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D) [29].

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

Participants: Nicolas Violante, Alban Gauthier, Stavros Diolatzis (*Intel*), Thomas Leimkühler (*MPI*), George Drettakis.




Figure 13: Our approach to reducing the memory requirements of the 3D Gaussian Splatting (3DGS) representation consists of three methods. A primitive pruning strategy that works during optimization, an SH band assignment that happens once, halfway through the optimization, and a codebook quantization that is applied as post-processing. These result in a reduced number of primitives, smaller footprint per primitive and on disk compression, respectively.

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 work and design a careful training strategy that allows effective disentanglement. Our model is the first that generates a variety of 3D cars that are multi-view consistent and that can be relit interactively with any environment map (see Fig. 14).

This work is a collaboration with Thomas Leimkühler from MPI Informatik and Stavros Diolatzis from Intel. It was published in Computer Graphics Forum and presented in Eurographics 2024 [21].

8.2.8 Studying the Effect of Volumetric Rendering Approximations for 3D Gaussian Splatting

Participants: George Drettakis, Georgios Kopanas, Bernhard Kerbl (*TU Wien*), Adam Celarek (*TU Wien*), Michael Wimmer (*TU Wien*).

We present an in-depth study of the effect of various approximations to volumetric rendering that are performed in the 3D Gaussian Splatting algorithm. These include the fact that constant opacity is used instead of density for attenuation, that self-attenuation is ignored and that primitives are assumed not to intersect. We show that for moderately large numbers of primitives, none of these approximations affects quality.

This work is in collaboration with B. Kerbl, A. Celarek and M. Wimmer from TU Wien, Austria.



Figure 14: Our method enables interactive physically-based rendering (PBR) with 3D GANs under arbitrary illumination conditions in the form of environment maps and unconstrained camera navigation, including 360° rotations. We achieve this with three main contributions: a method to estimate poses of a dataset of car images, a generative pipeline for PBR, and an improved generative network architecture and training solution. We only require a dataset of images of the desired object class (cars in our case) and a dataset of environment maps for training. Our method learns a disentangled representation of shading, enabling relighting with high-frequency reflections on shiny car bodies.

8.2.9 SVBRDF Prediction and Generative Image Modeling for Appearance Modeling of 3D Scenes

Participants: Alban Gauthier, Alexandre Lanvin, Adrien Bousseau, George Drettakis, Valentin Deschaintre (*Adobe Research*), Fredo Durand (*MIT*).

We introduce a method for texturing 3D scenes with SVBRDFs. This method exploits the complementary strengths of two active streams of research – generative image modeling and SVBRDF prediction. We build on the former to generate multiple views of a given scene conditioned on its geometry, while we leverage the latter to recover material parameters for each of these views, which we then merge into a common texture atlas.

This work is in collaboration with V. Deschaintre from Adobe Research and F. Durand from MIT.

8.2.10 Physical Reflections in Gaussian Splatting

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

Radiance fields are highly effective at capturing the appearance of complex lighting effects. In theory, they implement reflections with viewpoint-dependent coloration e.g. spherical harmonics in Gaussian splatting. In practice, they struggle to do so robustly and rely on duplicate copies of objects standing in for reflected images. Besides damaging geometric reconstruction, such entanglement precludes reasoning about the content of a scene and its physical properties. We seek to model all reflections in Gaussian splatting properly and explicitly, with hardware-accelerated raytracing and a principled BRDF model.

8.2.11 Incremental Scalable Reconstruction

Participants: Andreas Meuleman, Ishaan Shah, Alexandre Lanvin, George Drettakis, Bernhard Kerbl (*TU Wien*).

Existing methods for large-scale 3D reconstruction are slow and require preprocessed poses. We present an algorithm for novel view synthesis from casually captured videos of large-scale scenes, providing real-time feedback.

This work is in collaboration with B. Kerbl from TU Wien, Austria.

8.2.12 Caching Global Illumination using Gaussian Primitives

Participants: Ishaan Shah, Alban Gauthier, George Drettakis.

Achieving real-time rendering global illumination has been a long-standing goal in Computer Graphics. Modeling global illumination is computationally expensive making it impractical for real-time applications. Caching global illumination is a popular approach alleviating the computational cost. In this work, we propose a novel caching technique that uses Gaussian primitives to store global illumination. Our method is quick to optimize (5-10s) and is able to represent high-frequency illumination such as sharp shadows and glossy reflections.

8.2.13 Adding Textures to Gaussian Splatting

Participants: Panagiotis Papantonakis, Georgios Kopanas, Frédo Durand (*MIT*), George Drettakis.

Gaussian Splatting is a method for Novel View Synthesis, that represents static scenes as a set of Gaussian primitives. Each primitive is characterized by a set of parameters that controls its appearance and geometry. In the original method, a primitive could only have a single color throughout its extent. We observed that this limitation can lead to an increased number of primitives in areas where the geometry is simple (e.g. a planar surface) but the appearance is complex, as is the case of a textured surface. In these cases, the overhead of geometric parameters becomes a burden and contributes to an increased memory usage. We try to remove this limitation by supporting multiple color samples per primitive. Out of the possible options, we have decided to proceed with a texture grid attached on top of our primitives. In that way, apart from increasing the expressivity of appearance by querying the texture grid, we hope to benefit from other advantages of textures, like prefiltering and efficient storage.

This work is in collaboration with F. Durand from MIT.

8.2.14 Dynamic Gaussian Splatting

Participants: Petros Tzathas, Andreas Meuleman, Guillaume Cordonnier, George Drettakis.

Gaussian Splatting has quickly established itself as the prevailing method for novel view synthesis. Nevertheless, as originally formulated, the method restricts itself to static scenes. We seek to extend its modeling capabilities to scenes that include motion. In contrast to other works sharing the same goal, we aim to reformulate the specific manner in which the representation is adapted to better model the data by examining the resulting error.

8.2.15 Splat and Replace: 3D Reconstruction with Repetitive Elements

Participants: Nicolas Violante, Alban Gauthier, Andreas Meuleman, Fredo Durand (*MIT*), Thibault Groueix (*Adobe Research*), George Drettakis.

We propose to leverage the information of repetitive elements to improve the reconstruction of low-quality parts of scenes due to poor coverage and occlusions in the capture. After a base 3D reconstruction, we detect multiple occurrences of the same object, or instances, in the scene, and merge them together into a common shared representation that effectively aggregates multi-view and multi-illumination information from all instances. This representation is then finetuned and used to replace poorly-reconstructed instances, improving the overall geometry while also accounting for the specific appearance variations of each instance.

This work is in collaboration with F. Durand from MIT and T. Groueix from Adobe Research.

8.3 Physical Simulation for Graphics

8.3.1 Volcanic Skies: coupling explosive eruptions with atmospheric simulation to create consistent skiescapes

Participants: Guillaume Cordonnier, Cilliers Pretorius (*University of Cape Town*), James Gain (*University of Cape Town*), Maud Lastic (*LIX*), Damien Rohmer (*LIX*), Marie-Paule Cani (*LIX*), Jiong Chen (*LIX, Inria GeomeriX*).

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 paper, we introduce the first general-purpose model for bi-directional interaction between the atmosphere and a volcano plume. In line with recent interactive volcano models, we approximate the plume dynamics with Lagrangian disks and spheres and the atmosphere with sparse layers of 2D Eulerian grids, enabling us to focus on the transfer of physical quantities such as temperature, ash, moisture, and wind velocity between these sub-models. We subsequently generate volumetric animations by noise-based procedural upsampling keyed to aspects of advection, convection, moisture, and ash content to generate a fully-realized volcanic skyscape. Our model captures most of the visually salient features emerging from volcano-sky interaction, such as windswept plumes, enmeshed cap, bell and skirt clouds, shockwave effects, ash rain, and sheathes of lightning visible in the dark (see Fig. 15).



Figure 15: Impressive phenomena arise from the interaction between the plume of explosive volcanoes and the atmosphere. Our framework is able to capture (from left to right): embedded skirt clouds, cap clouds, ash rain, and volcanic lightning.

This work is a collaboration with Cilliers Pretorius and James Gain from University of Cape Town, Maud Lastic, Damien Rohmer and Marie-Paule Cani from LIX, and Jiong Chen from LIX and Inria GeomeriX. It was published in Computer Graphics Forum and presented at Eurographics 2024 [18].

8.3.2 Fast simulation of viscous lava flow using Green's functions as a smoothing kernel

Participants: Guillaume Cordonnier, Yannis Kedadry (*Ecole Normale Supérieure*).

We present a novel approach to simulate large-scale lava flow in real-time. We use a depth-averaged model from numerical volcanology to simplify the problem to 2.5D using a single layer of particle

with thickness. Yet, lava flow simulation is challenging due to its strong viscosity which introduces computational instabilities. We solve the associated partial differential equations with approximated Green's functions and observe that this solution acts as a smoothing kernel. We use this kernel to diffuse the velocity based on Smoothed Particle Hydrodynamics. This yields a representation of the velocity that implicitly accounts for horizontal viscosity which is otherwise neglected in standard depth-average models. We demonstrate that our method efficiently simulates large-scale lava flows in real-time (see Fig. 16).




Figure 16: Simulation of a viscous lava flow over Mount St Helens.

This work is a collaboration with Yannis Kedadry from Ecole Normale Supérieure, initiated during Yannis' internship in GraphDeco in 2023. It was presented as a poster at the ACM SIGGRAPH / Eurographics Symposium on Computer Animation 2024 [27].

8.3.3 Unerosion: simulating terrain evolution back in time

Participants: Guillaume Cordonnier, Zhanyu Yang (*Purdue University*), Bedrich Benes (*Purdue University*), Marie-Paule Cani (*LIX*), Christian Perrenoud (*Institut de Paléontologie Humaine*).

While the past of terrain cannot be known precisely because an effect can result from many different causes, exploring these possible pasts opens the way to numerous applications ranging from movies and games to paleogeography. We introduce unerosion, an attempt to recover plausible past topographies from an input terrain represented as a height field (see Fig. 17). Our solution relies on novel algorithms for the backward simulation of different processes: fluvial erosion, sedimentation, and thermal erosion. This is achieved by re-formulating the equations of erosion and sedimentation so that they can be simulated back in time. These algorithms can be combined to account for a succession of climate changes backward in time, while the possible ambiguities provide editing options to the user. Results show that our solution can approximately reverse different types of erosion while enabling users to explore a variety of alternative pasts. Using a chronology of climatic periods to inform us about the main erosion phenomena, we also went back in time using real measured terrain data. We checked the consistency with geological findings, namely the height of river beds hundreds of thousands of years ago.




Figure 17: A terrain that has been eroded (left) by fluvial and thermal erosion is uneroded by automatically tracing the direction of the main flow (indicated by the arrows) to its previous state.

This work is a collaboration with Zhanyu Yang and Bedrich Benes from Purdue University, Marie-Paule Cani from LIX, and Christian Perrenoud from the Institut de Paléontologie Humaine. It was published in Computer Graphics Forum and presented at the ACM SIGGRAPH / Eurographics Symposium on Computer Animation 2024 [23], where it received an Honorable Mention for the Best Paper Award.

8.3.4 CHONK 1.0: landscape evolution framework: cellular automata meets graph theory

Participants: Guillaume Cordonnier, Boris Gailleton (*Géosciences Rennes*), Luca C. Malatesta (*GFZ Helmholtz Research Centre for Geosciences*), Jean Braun (*GFZ Helmholtz Research Centre for Geosciences*).

Landscape evolution models (LEMs) are prime tools for simulating the evolution of source-to-sink systems through ranges of spatial and temporal scales. A plethora of various empirical laws have been successfully applied to describe the different parts of these systems: fluvial erosion, sediment transport and deposition, hillslope diffusion, or hydrology. Numerical frameworks exist to facilitate the combination of different subsets of laws, mostly by superposing grids of fluxes calculated independently. However, the exercise becomes increasingly challenging when the different laws are inter-connected: for example when a lake breaks the upstream–downstream continuum in the amount of sediment and water it receives and transmits; or when erosional efficiency depends on the lithological composition of the sediment flux. In this contribution, we present a method mixing the advantages of cellular automata and graph theory to address such cases (see Fig. 18). We demonstrate how the former ensures interoperability of the different fluxes (e.g. water, fluvial sediments, hillslope sediments) independently of the process law implemented in the model, while the latter offers a wide range of tools to process numerical landscapes, including landscapes with closed basins. We provide three scenarios largely benefiting from our method: (i) one where lake systems are primary controls on landscape evolution, (ii) one where sediment provenance is closely monitored through the stratigraphy and (iii) one where heterogeneous provenance influences fluvial incision dynamically. We finally outline the way forward to make this method more generic and flexible.

This work is a collaboration with Boris Gailleton from Géosciences Rennes, and Luca C. Malatesta and Jean Braun from the GFZ Helmholtz Research Centre for Geosciences. It was published in Geoscientific Model Development [11].

8.3.5 FastFlow: GPU Acceleration of Flow and Depression Routing for Landscape Simulation

Participants: Aryamaan Jain, Guillaume Cordonnier, Bernhard Kerbl (*TU Wien*), James Gain (*University of Cape Town*), Brandon Finley (*University of Lausanne*).




Figure 18: Illustrative landscape highlighting several key attributes of the sedimentary system that CHONK is designed to solve with a novel approach blending cellular automata and graph-based methods. The different domains connected by the river network and hillslope transfers of material highlight the interconnected nature of the different processes.

Terrain analysis plays an important role in computer graphics, hydrology and geomorphology. In particular, analyzing the path of material flow over a terrain with consideration of local depressions is a precursor to many further tasks in erosion, river formation, and plant ecosystem simulation. For example, fluvial erosion simulation used in terrain modeling computes water discharge to repeatedly locate erosion channels for soil removal and transport. Despite its significance, traditional methods face performance constraints, limiting their broader applicability.

We propose a novel GPU *flow routing* algorithm that computes the water discharge in $\mathcal{O}(\log n)$ iterations for a terrain with n vertices (assuming n processors). We also provide a *depression routing* algorithm to route the water out of local minima formed by depressions in the terrain, which converges in $\mathcal{O}(\log^2 n)$ iterations. Our implementation of these algorithms leads to a $5\times$ speedup for flow routing and $34\times$ to $52\times$ speedup for depression routing compared to previous work on a 1024^2 terrain, enabling interactive control of terrain simulation (see Fig. 19).

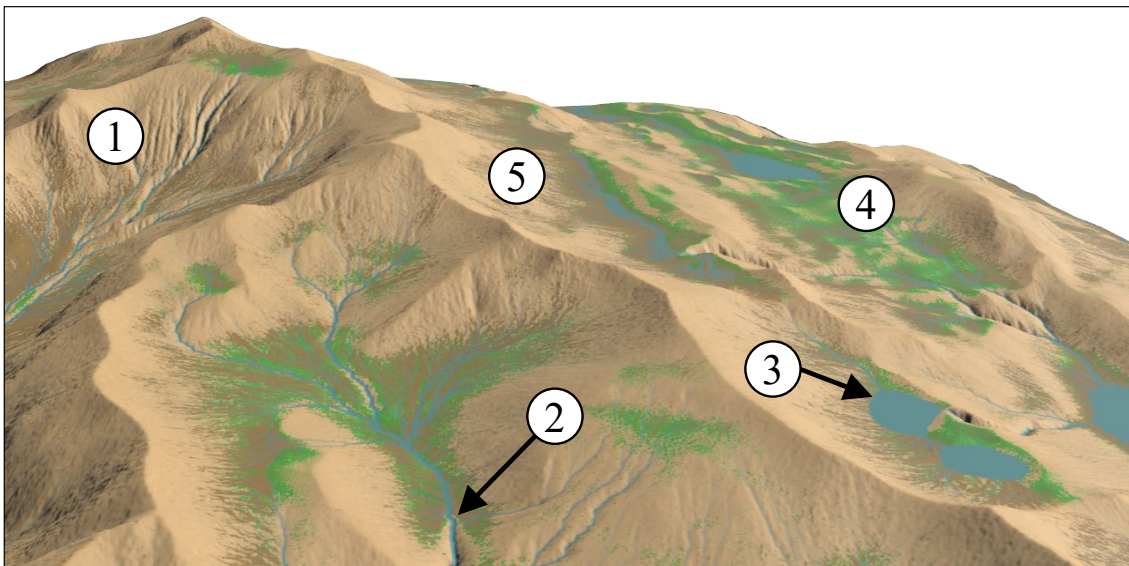


Figure 19: Our GPU flow and depression routing algorithms can be applied to accelerate multiple aspects of landscape simulation, including 1) fluvial erosion, 2) rivers, 3) lakes, 4) ecosystems, and 5) sediment deposition.

This work is a collaboration with Bernhard Kerbl from TU Wien, Austria, James Gain from University of Cape Town, South Africa, and Brandon Finley from University of Lausanne, Switzerland. It was published in Computer Graphics Forum, and presented in Pacific Graphics 2024 where it received the Best Paper Award [14].

8.3.6 Efficient Debris-flow Simulation for Steep Terrain Erosion

Participants: Aryamaan Jain, Guillaume Cordonnier, Bedrich Benes (*Purdue University*).

Erosion simulation is a common approach used for generating and authoring mountainous terrains. While water is considered the primary erosion factor, its simulation fails to capture steep slopes near the ridges. In these low-drainage areas, erosion is often approximated with slope-reducing erosion, which yields unrealistically uniform slopes. However, geomorphology observed that another process dominates the low-drainage areas: erosion by debris flow, which is a mixture of mud and rocks triggered by strong climatic events. We propose a new method to capture the interactions between debris flow and fluvial erosion thanks to a new mathematical formulation for debris flow erosion derived from geomorphology and a unified GPU algorithm for erosion and deposition. In particular, we observe that sediment and debris deposition tend to intersect river paths, which motivates the design of a new, approximate flow routing algorithm on the GPU to estimate the water path out of these newly formed depressions. We demonstrate that debris flow carves distinct patterns in the form of erosive scars on steep slopes and cones of deposited debris competing with fluvial erosion downstream (see Fig. 20).




Figure 20: A landscape carved by simulated erosion and deposition. (1) Deep scars of debris flow erosion, (2) deposit fan, (3) low-slope fluvially shaped hill, (4) sharp ridges carved by debris flow erosion, (5) lake found by our depression routing algorithm.

This work is a collaboration with Bedrich Benes from Purdue University, USA. It was published in ACM Transactions on Graphics, and presented in SIGGRAPH 2024 [13].

8.3.7 Learning Based Infinite Terrain Generation with Level of Detailing

Participants: Aryamaan Jain, Avinash Sharma (*IIT Jodhpur*), KS Rajan (*IIIT Hyderabad*).

Infinite terrain generation is an important use case for computer graphics, games and simulations. However, current techniques are often procedural which reduces their realism. We introduce a learning-based generative framework for infinite terrain generation along with a novel learning-based approach for level-of-detailing of terrains. Our framework seamlessly integrates with quad-tree-based terrain rendering algorithms. Our approach leverages image completion techniques for infinite generation and progressive super-resolution for terrain enhancement. Notably, we propose a novel quad-tree-based training method for terrain enhancement which enables seamless integration with quad-tree-based rendering algorithms while minimizing the errors along the edges of the enhanced terrain. Comparative evaluations against existing techniques demonstrate our framework's ability to generate highly realistic terrain with effective level-of-detailing (see Fig. 21).




Figure 21: (a) We outline an approach for generating infinite terrain (learned from real-world DEM), while incorporating level of detailing within a learning-based framework. (b) To efficiently manage the generated terrain data, a quad-tree structure is employed, which facilitates operations such as view frustum culling. (c) The entire process is executed in real-time, with the terrain being rendered simultaneously.

This work is a collaboration with Avinash Sharma from IIT Jodhpur, India, and KS Rajan from IIIT Hyderabad, India. It was published in 3DV conference proceedings, and presented in 3DV 2024 [26].

8.3.8 Physically-based analytical erosion for fast terrain generation

Participants: Petros Tzathas, Guillaume Cordonnier, Boris Gailleton (*Géosciences Rennes*), Philippe Steer (*Géosciences Rennes & Institut universitaire de France*).

Terrain generation methods have long been divided between procedural and physically-based. Procedural methods build upon the fast evaluation of a mathematical function but suffer from a lack of geological consistency, while physically-based simulation enforces this consistency at the cost of thousands of iterations unraveling the history of the landscape. In particular, the simulation of the competition between tectonic uplift and fluvial erosion expressed by the stream power law raised recent interest in computer graphics as this allows the generation and control of consistent large-scale mountain ranges, albeit at the cost of a lengthy simulation. This work explores the analytical solutions of the stream power law and proposes a method that is both physically-based and procedural, allowing fast and consistent large-scale terrain generation. In our approach, time is no longer the stopping criterion of an iterative process but acts as the parameter of a mathematical function, a slider that controls the aging of the input terrain from a subtle erosion to the complete replacement by a fully formed mountain range. While analytical solutions have been proposed by the geomorphology community for the 1D case, extending them to a 2D heightmap proves challenging. We propose an efficient implementation of the analytical solutions with a multigrid accelerated iterative process and solutions to incorporate landslides and hillslope processes – two erosion factors that complement the stream power law.



Figure 22: Our method allows the fast generation of a variety of physically-based terrains. Here a canyon-like landscape is demonstrated.

This work is a collaboration with Boris Gailleton and Philippe Steer from the Géosciences department of the University of Rennes. It was published in the Computer Graphics Forum and presented at Eurographics 2024 [20].

8.3.9 Windblown sand around obstacles- simulation and validation of deposition patterns

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

Sand dunes are iconic landmarks of deserts, but can also put human infrastructures at risk, for instance by forming near buildings or roads. We present 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 via a qualitative comparison of the erosion and deposition patterns produced by our simulator against real-world patterns measured by prior work under controlled conditions. We can see an example of the desert scenes our model produces in Fig. 23.

It was published in the Proceedings of the ACM on Computer Graphics and Interactive Techniques 2024 [19].

8.3.10 Sandstone Erosion

Participants: Aryamaan Jain, Guillaume Cordonnier, Zhanyu Yang (*Purdue University*), Zhaopeng Wang (*Purdue University*), Marie-Paule Cani (*Ecole polytechnique*), Bedrich Benes (*Purdue University*).

We introduce a novel computational framework for simulating the long-term erosion and deposition of sandstone landscapes, focusing on the interplay between fabric interlocking, wind-driven erosion, and




Figure 23: Desertscape simulated with our method. Note the erosion and deposition patterns around the barn, trees, and cactus. 3D models of farm buildings and trees from quaternius.com under a CC0 license.

fluvial processes. We employ the Material Point Method (MPM) with the MLS-MPM variant to model dynamic stress distribution within sandstone structures, capturing the strengthening effect of fabric interlocking, where increased vertical stress leads to enhanced cohesion and erosion resistance. Wind erosion is addressed through a hybrid approach combining uniform background deflation, modeled using a modified Mohr-Coulomb criterion and Rayleigh distribution of wind forces, with localized intensification simulated via Smoothed Particle Hydrodynamics (SPH) to capture directional wind abrasion and deflation. Fluvial erosion is incorporated using a stream power law model, with flow discharge and depression routing derived from a grid-based representation of the sandstone surface, adapted from existing 2D algorithms to handle 3D scenarios. Furthermore, we introduce a deposition model to simulate the accumulation of sediment, complementing the erosional processes and contributing to the dynamic evolution of the landscape. This allows for a more comprehensive simulation of sandstone evolution under diverse environmental conditions, providing insights into the formation of iconic landforms and the evolution of sandstone landscapes over geological timescales.

This work is a collaboration with Zhanyu Yang, Zhaopeng Wang and Bedrich Benes from Purdue University, USA, and Marie-Paule Cani from Ecole Polytechnique, France.

8.3.11 Neural surrogate model for wind prediction above a terrain

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

Driven by the long term goal of getting a fast and invertible sand simulation pipeline, we investigate neural-based solutions for predicting the wind above a terrain. In particular, we predict time-averaged wind fields that are easier to predict though informative enough to model sand transport.

8.3.12 Extracting terrain characteristics using erosion laws

Participants: Melike Aydinlilar, Guillaume Cordonnier.

Stream Power Law (SPL) describes morphological characteristics of the fluvially eroded terrains. Assuming a steady state solution for our augmented SPL formulation, we extract the parameters of the solution, therefore characteristics of the given terrain. Using these parameters, we use the same laws for synthesizing novel terrains and interpolating between the generated terrains.

9 Partnerships and cooperations

9.1 International research visitors

9.1.1 Visits of international scientists

Other international visits to the team

Fredo Durand

Status: Professor

Institution of origin: MIT

Country: USA

Dates: August 2023 to July 2024

Context of the visit: Sabbatical year

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

Eric Paquette

Status: Professor

Institution of origin: École de Technologie Supérieure, Montréal

Country: Canada

Dates: October 14 to November 15, 2024

Context of the visit: Sabbatical year

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

Gilda Manfredi

Status: PhD Student

Institution of origin: University of Basilicata and University Salento

Country: Italy

Dates: Until March 31st, 2024

Context of the visit: Research collaboration

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

Bernhard Kerbl

Status: Postdoctoral researcher

Institution of origin: TU Wien

Country: Austria

Dates: January 8-26nd, 2024

Context of the visit: Research collaboration

Mobility program/type of mobility: Research stay

Jun-Yan Zhou**Status:** Professor**Institution of origin:** CMU**Country:** USA**Dates:** May, 2024**Context of the visit:** Invited talk in the Graphdeco Deep Learning Seminar**Mobility program/type of mobility:** Self-funded research stay**Alyosha Efros****Status:** Professor**Institution of origin:** University of California, Berkeley**Country:** USA**Dates:** May, 2024**Context of the visit:** Invited talk in the Graphdeco Deep Learning Seminar**Mobility program/type of mobility:** Self-funded research stay**Steve Marschner****Status:** Professor**Institution of origin:** Cornell**Country:** USA**Dates:** June 13, 2024**Context of the visit:** Invited talk in the Graphdeco Seminar "Simulations of optical scattering for material appearance modeling"**Mobility program/type of mobility:** Invited talk**Daniel Ritchie****Status:** Associate Professor**Institution of origin:** Brown University**Country:** USA**Dates:** October 27nd to November 1st, 2024**Context of the visit:** Group retreat**Mobility program/type of mobility:** Invited talk

Peter Hedman

Status: Staff Scientist

Institution of origin: Google DeepMind

Country: United Kingdom

Dates: October 27nd to November 1st, 2024

Context of the visit: Group retreat

Mobility program/type of mobility: Invited talk

Chris Wotjan

Status: Research Scientist

Institution of origin: IST Austria

Country: Austria

Dates: December 9-11th, 2024

Context of the visit: Ph.D. defense of Nicolas Rosset

Mobility program/type of mobility: Invited talk in the workshop on machine learning, geometry and fluid simulation

9.1.2 Visits to international teams**Research stays abroad****Nicolàs Violante**

Visited institution: Adobe Research San Francisco

Country: USA

Dates: September to November 2024

Context of the visit: Industrial internship

Mobility program/type of mobility: Internship

Nicolas Rosset

Visited institution: Laboratoire d'Informatique Graphique de l'Université de Montréal (LIGUM)

Country: Canada

Dates: May 2024

Context of the visit: Presentation of recent work

Mobility program/type of mobility: Research visit

Henro Kriel**Visited institution:** Brown Visual Computing group**Country:** USA**Dates:** December 2024 (three weeks)**Context of the visit:** Research collaboration**Mobility program/type of mobility:** Research stay

9.2 European initiatives

9.2.1 Horizon Europe

NERPHYS

Participants: George Drettakis, Andreas Meuleman, Alban Gauthier, Alexandre Lanvin, Panagiotis Papantonakis, Yohan Poirier-Ginter, Petros Tzathas, Nicolás Violante, Ishaan Shah.

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

Title: Empowering Neural Rendering Methods with Physically-Based Capabilities**Duration:** From December 1, 2024 to November 30, 2029**Partners:**

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

Inria contact: George Drettakis**Coordinator:** George Drettakis

Summary: While long restricted to an elite of expert digital artists, 3D content creation has recently been greatly simplified by deep learning. Neural representations of 3D objects have revolutionized real-world capture from photos, while generative models are starting to enable 3D object synthesis from text prompts. These methods use differentiable neural rendering that allows efficient optimization of the powerful and expressive “soft” neural representations, but ignores physically-based principles, and thus has no guarantees on accuracy, severely limiting the utility of the resulting content.

Differentiable physically-based rendering on the other hand can produce 3D assets with physics-based parameters, but depends on rigid traditional “hard” graphics representations required for light-transport computation, that make optimization much harder and is also costly, limiting applicability.

In NERPHYS, we will combine the strengths of both neural and physically-based rendering, lifting their respective limitations by introducing polymorphic 3D representations, i.e., capable of morphing between different states to accommodate both efficient gradient-based optimization and physically-based light transport. By augmenting these representations with corresponding polymorphic differentiable renderers, our methodology will unleash the potential of neural rendering to produce physically-based 3D assets with guarantees on accuracy.

NERPHYS will have ground-breaking impact on 3D content creation, moving beyond today’s simplistic plausible imagery, to full physically-based rendering with guarantees on error, enabling the use of powerful neural rendering methods in any application requiring accuracy. Our polymorphic approach will fundamentally change how we reason about scene representations for geometry and appearance, while our rendering algorithms will provide a new methodology for image synthesis, e.g., for training data generation or visual effects.

DLift

Participants: Adrien Bousseau, Jiayi Wei.

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

Title: Lifting Design Drawings to 3D

Duration: From April 1, 2024 to September 30, 2025

Partners:

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

Inria contact: Adrien BOUSSEAU

Coordinator: Adrien BOUSSEAU

Summary: Building on the outcome of the ERC Starting Grant D3 (ERC-2016-STG 714221), the objective of DLift is to demonstrate how our technology for reconstructing drawings in 3D can streamline the Computer-Aided-Design (CAD) workflow, and to take the first steps in industrializing this technology.

Drawing is a fundamental tool of product design. However, while 2D drawings are easily understood by humans, they are currently not interpretable by computers. To confront their ideas with physical reality, designers have to separately create 3D models that form the necessary input for digital engineering tools. Skilled 3D modelers often need several hours to convert a drawing into a 3D model, at an hourly rate of 50 euros on average, making 3D modeling a major bottleneck that hinders rapid iterations between ideation and prototyping.

The ERC Starting Grant D3 aimed at bridging design exploration and design engineering by offering the first algorithm capable of automatically lifting 2D design drawings to 3D. The objective of DLift is to optimize our algorithm to integrate it within leading CAD software. This objective will be achieved through three iterative steps:

1. Feature development. Based on preliminary discussions with CAD users and software editors, we have identified a set of key features to unleash the potential of our technology.
2. User testing. We will hire professional designers and 3D modelers to stress test our tool, first to make it robust to the diversity and complexity of real-world design drawings, but also to assemble a portfolio of artworks that will illustrate diverse ways in which our technology can be used in practice.
3. Transfer. Within each iteration of feature development and testing, we will work hand-in-hand with software companies interested in our technology to assess how it addresses their specific needs. We aim at licensing our technology to one or several of these companies such that it can be integrated in their own solutions.

9.2.2 H2020 projects**FUNGRAPH**

Participants: George Drettakis, Georgios Kopanas, Andreas Meuleman, Alban Gauthier, Alexandre Lanvin, Panagiotis Papantonakis, Yohan Poirier-Ginter, Petros Tzathas, Nicolás Violante, Ishaan Shah, Vishal Pani, Clement Remy.

[FUNGRAPH project on cordis.europa.eu](https://cordis.europa.eu); [Project site](#)

Title: A New Foundation for Computer Graphics with Inherent Uncertainty

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

Partners:

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

Inria contact: George Drettakis

Coordinator: George Drettakis

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

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

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

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

9.3 National initiatives

9.3.1 Visits of French scientists

Julie Digne

Status: Research Scientist

Institution of origin: CNRS, LIRIS

Country: France

Dates: December 9-10th, 2024

Context of the visit: Ph.D. defense of Nicolas Rosset

Mobility program/type of mobility: Invited talk in the workshop on machine learning, geometry and fluid simulation

Kiwon Um**Status:** Research Scientist**Institution of origin:** Telecom Paris**Country:** France**Dates:** December 10th, 2024**Context of the visit:** Ph.D. defense of Nicolas Rosset**Mobility program/type of mobility:** Invited talk in the workshop on machine learning, geometry and fluid simulation**9.3.2 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 INVTERRA**Participants:** Guillaume Cordonnier, Aryamaan Jain, Melike Aydinlilar, Juan-Sebastian Osorno-Bolivar.[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.

ANR-NSF NaturalCAD

Participants: Adrien Bousseau, Henro Kriel, Marzia Riso.

[Project description on anr.fr](#)

Title: Learning to Translate Freehand Design Drawings into Parametric CAD Programs

Duration: From March 1, 2024 to October 31, 2027

Partners:

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

Inria contact: Adrien BOUSSEAU

Coordinator: Adrien BOUSSEAU

Summary: Computer Aided Design (CAD) is a multi-billion dollar industry responsible for the digital design of almost all manufactured goods. It leverages parametric modeling, which allows dimensions of a design to be changed, facilitating physically-based optimization and design re-mixing by non-experts. But CAD's potential is diminished by the difficulty of creating parametric models: in addition to mastering design principles, professionals must learn complex CAD software interfaces. To promote effective modeling strategies and creative flow, design educators advocate freehand drawing as a preliminary step to parametric modeling. Unfortunately, CAD systems do not understand these drawings, so designers must re-create their entire design using complex CAD software. Can we automatically convert freehand drawings to parametric CAD models? Sketch-based modeling techniques do not produce parametric CAD programs; classic CAD reverse-engineering techniques cannot handle drawings as input; the newer field of visual program induction is promising but has been demonstrated only on simple shapes and programs. By leveraging the visual vocabulary shared by drawing and CAD modeling, we will develop a system to translate from the natural language of drawing to the formal language of CAD.

10 Dissemination

Participants: Adrien Bousseau, Guillaume Cordonnier, George Drettakis.

10.1 Promoting scientific activities

10.1.1 Scientific events: selection

Member of the conference program committees

- Adrien Bousseau was on the program committee of SIGGRAPH 2024.
- Guillaume Cordonnier was on the program committee of Pacific Graphics 2024.
- George Drettakis was on the program committee of ACM SIGGRAPH 2024 and EGSR 2024.

Reviewer

- Adrien Bousseau was a reviewer for SIGGRAPH Asia and Eurographics.
- Guillaume Cordonnier was a reviewer for Eurographics, SIGGRAPH and SIGGRAPH Asia.
- Alban Gauthier was a reviewer for SIGGRAPH Asia 2024 and Pacific Graphics 2024.
- Andreas Meuleman was a reviewer for SIGGRAPH and SIGGRAPH Asia.
- Nicolas Violante was a reviewer Eurographics and SIGGRAPH Asia.

10.1.2 Journal

Member of the editorial boards

- George Drettakis and Adrien Bousseau are Associate Editors for Computer Graphics Forum.

Reviewer - reviewing activities

- Adrien Bousseau was a reviewer for ACM Transactions on Graphics.
- Andreas Meuleman was a reviewer for IEEE Transactions on Visualization and Computer Graphics.
- Guillaume Cordonnier was a reviewer for Stochastic Environmental Research and Risk Assessment.

10.1.3 Invited talks

- George Drettakis presented invited talks at Kuytai.ai and Valeo in Paris in April, at the XR2C2 event in Sophia-Antipolis in May, at a Taiwan/French Embassy event in June, at the Journées Scientifiques Inria in Grenoble in August, at the University of Crete/FORTH in September, at Aalto University, and at the SophiaAI summit in Sophia-Antipolis in November.
- Nicolas Rosset gave a talk at the École de technologie supérieure in Montréal.
- Alban Gauthier gave a talk at Inria Grenoble.
- Andreas Meuleman gave an invited talk at Centre Inria de l'Université de Lorraine.
- Adrien Bousseau gave an invited keynote talk to the “*Journée IHM GdR IG-RV*” in Lille, May 30-31 and a talk at the XR2C2 event in Sophia-Antipolis, also in May.

10.1.4 Leadership within the scientific community

- George Drettakis is a member of the Aalto University School of Science Scientific Advisory Board and performed an evaluation onsite in November this year.
- George Drettakis chairs the ACM SIGGRAPH Papers Advisory Group and the Eurographics Working Group on Rendering (which is also the Steering Committee for EGSR).

10.1.5 Scientific expertise

- George Drettakis reviewed ERC grant proposals.
- George Drettakis wrote several tenure and other promotion letters for candidates in North America and Europe.
- Adrien Bousseau reviewed a grant proposal for the Swiss National Science Foundation.

10.1.6 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.
- Guillaume Cordonnier is a member of the *Comité du centre* of the Centre Inria d'Université Côte d'Azur.
- George Drettakis is an elected member of the Inria Scientific Board and in charge of the Morgenstern Colloquium in Sophia-Antipolis.
- George Drettakis co-leads the Working Group on Rendering of the GT IGRV with R. Vergne and organized the GT workshop in Paris in June 2024. He is member of the Administrative Council of the Eurographics French Chapter.
- Adrien Bousseau co-chairs the committee in charge of electing a young research fellow each year for the French Chapter of Eurographics.
- Adrien Bousseau is a member of the committee for doctoral studies (CSD) of the Centre Inria d'Université Côte d'Azur.

10.2 Teaching - Supervision - Juries

10.2.1 Teaching

- Alban Gauthier was a teaching assistant for the Python course at Ecole Centrale Méditerranée's bachelor.
- Capucine Nghiem was a teaching assistant for the following courses: Informatique Graphique pour la Science des Données, Université Paris-Saclay, L2 Informatique; Introduction à l'IHM, Université Paris-Saclay, L2 Informatique; Fundamentals of HCI, Université Paris-Saclay, HCI Master

10.2.2 Supervision

- Nicolas Rosset, co-supervised by Guillaume Cordonnier and Adrien Bousseau, defended in December 2024 [32].
- 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 as part of a "co-tutelle" Ph.D., since September 2023.
- Petros Tzathas, co-supervised by George Drettakis and Guillaume Cordonnier since November 2023.
- Panagiotis Papantonakis, supervised by George Drettakis since September 2023.
- Berend Baas, supervised by Adrien Bousseau since December 2023.
- Aryamaan Jain, supervised by Guillaume Cordonnier since October 2023.
- Henro Kriel, supervised by Adrien Bousseau since September 2024.
- Alexandre Lanvin, co-supervised by George Drettakis and Adrien Bousseau since December 2024.

10.2.3 Juries

- Adrien Bousseau was a reviewer for the Ph.D. thesis of Emmanuel Rodriguez (Université Grenoble Alpes) and an examiner for the Ph.D. thesis of Florent Robert (Université Côtes d'Azur).

10.3 Popularization

10.3.1 Participation in Live events

- Adrien Bousseau presented his work to high school students twice in the context of "1 scientifique, 1 classe: chiche!".

11 Scientific production

11.1 Major publications

- [1] 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>.
- [2] 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>.
- [3] 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>.
- [4] 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/revue/Basilic/2018/HPPFDB18>.
- [5] A. Jain, B. Benes and G. Cordonnier. 'Efficient Debris-flow Simulation for Steep Terrain Erosion'. In: *ACM Transactions on Graphics* 43.4 (2024), p. 58. DOI: [10.1145/3658213](https://doi.org/10.1145/3658213). URL: <https://hal.science/hal-04574826>.
- [6] G. Jouvét, G. Cordonnier, B. Kim, M. Lüthi, A. Vieli and A. Aschwanden. 'Deep learning speeds up ice flow modelling by several orders of magnitude'. In: *Journal of Glaciology* (22nd Dec. 2021), pp. 1–14. DOI: [10.1017/jog.2021.120](https://doi.org/10.1017/jog.2021.120). URL: <https://inria.hal.science/hal-03525458>.
- [7] 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>.

- [8] 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>.
- [9] 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>.
- [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>.

11.2 Publications of the year

International journals

- [11] B. Gailleton, L. C. Malatesta, G. Cordonnier and J. Braun. ‘CHONK 1.0: landscape evolution framework: cellular automata meets graph theory’. In: *Geoscientific Model Development* 17.1 (5th Jan. 2024), pp. 71–90. DOI: [10.5194/gmd-17-71-2024](https://doi.org/10.5194/gmd-17-71-2024). URL: <https://hal.science/hal-04525782> (cit. on p. 30).
- [12] A. Gauthier, B. Kerbl, J. Levallois, R. Faury, J. Thiery and T. Boubekeur. ‘MatUp: Repurposing Image Upsamplers for SVBRDFs’. In: *Computer Graphics Forum* 43.4 (24th July 2024). DOI: [10.1111/cgf.15151](https://doi.org/10.1111/cgf.15151). URL: <https://hal.science/hal-04669003> (cit. on p. 22).
- [13] A. Jain, B. Benes and G. Cordonnier. ‘Efficient Debris-flow Simulation for Steep Terrain Erosion’. In: *ACM Transactions on Graphics* 43.4 (2024), p. 58. DOI: [10.1145/3658213](https://doi.org/10.1145/3658213). URL: <https://hal.science/hal-04574826> (cit. on p. 32).
- [14] A. Jain, B. Kerbl, J. Gain, B. Finley and G. Cordonnier. ‘FastFlow: GPU Acceleration of Flow and Depression Routing for Landscape Simulation’. In: *Computer Graphics Forum* 43.7 (2024). URL: <https://hal.science/hal-04684270> (cit. on pp. 9, 32).
- [15] B. Kerbl, A. Meuleman, G. Kopanas, M. Wimmer, A. Lanvin and G. Drettakis. ‘A Hierarchical 3D Gaussian Representation for Real-Time Rendering of Very Large Datasets’. In: *ACM Transactions on Graphics* 44.3 (2024). URL: <https://inria.hal.science/hal-04568786>. In press (cit. on p. 24).
- [16] X. Luo, L. Scandolo, A. Bousseau and E. Eisemann. ‘Single-Image SVBRDF Estimation with Learned Gradient Descent’. In: *Computer Graphics Forum* 43.2 (Apr. 2024). URL: <https://inria.hal.science/hal-04484036> (cit. on p. 17).
- [17] Y. Poirier-Ginter, A. Gauthier, J. Phillip, J.-F. Lalonde and G. Drettakis. ‘A Diffusion Approach to Radiance Field Relighting using Multi-Illumination Synthesis’. In: *Computer Graphics Forum* 43.4 (July 2024). URL: <https://inria.hal.science/hal-04628346> (cit. on p. 23).
- [18] C. Pretorius, J. Gain, M. Lastic, G. Cordonnier, J. Chen, D. Rohmer and M.-P. Cani. ‘Volcanic Skies: coupling explosive eruptions with atmospheric simulation to create consistent skies’. In: *Computer Graphics Forum* 43.2 (22nd Apr. 2024). DOI: [10.1111/cgf.15034](https://doi.org/10.1111/cgf.15034). URL: <https://hal.science/hal-04527448> (cit. on p. 28).
- [19] N. Rosset, R. Duvigneau, A. Bousseau and G. Cordonnier. ‘Windblown sand around obstacles - simulation and validation of deposition patterns’. In: *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 7.1 (2024), p. 13. DOI: [10.1145/3651284](https://doi.org/10.1145/3651284). URL: <https://inria.hal.science/hal-04566241>. In press (cit. on p. 34).
- [20] P. Tzathas, B. Gailleton, P. Steer and G. Cordonnier. ‘Physically-based analytical erosion for fast terrain generation’. In: *Computer Graphics Forum* 43.2 (Apr. 2024). DOI: [10.1111/cgf.15033](https://doi.org/10.1111/cgf.15033). URL: <https://hal.science/hal-04525371> (cit. on p. 34).
- [21] N. Violante, A. Gauthier, S. Diolatzis, T. Leimkühler and G. Drettakis. ‘Physically-Based Lighting for 3D Generative Models of Cars’. In: *Computer Graphics Forum* 43.2 (Apr. 2024). URL: <https://hal.science/hal-04493971> (cit. on p. 25).

- [22] K. Wolski, A. Djeacoumar, A. Javanmardi, H.-P. Seidel, C. Theobalt, G. Cordonnier, K. Myszkowski, G. Drettakis, X. Pan and T. Leimkühler. ‘Learning Images Across Scales Using Adversarial Training’. In: *ACM Transactions on Graphics* 43.4 (July 2024), p. 131. DOI: [10.1145/3658190](https://doi.org/10.1145/3658190). URL: <https://hal.science/hal-04593775> (cit. on p. 20).
- [23] Z. Yang, G. Cordonnier, M.-P. Cani, C. Perrenoud and B. Benes. ‘Unerosion: Simulating Terrain Evolution Back in Time’. In: *Computer Graphics Forum* 43.8 (2024). DOI: [url{https://doi.org/10.1111/cgf.15182}](https://doi.org/10.1111/cgf.15182). URL: <https://hal.science/hal-04714138> (cit. on pp. 10, 30).
- [24] E. Yu, F. Chevalier, K. Singh and A. Bousseau. ‘3D-Layers: Bringing Layer-Based Color Editing to VR Painting’. In: *ACM Transactions on Graphics* 43.4 (July 2024). DOI: [10.1145/3658183](https://doi.org/10.1145/3658183). URL: <https://inria.hal.science/hal-04595138> (cit. on p. 14).

International peer-reviewed conferences

- [25] S. Diolatzis, T. Zirr, A. Kuznetsov, G. Kopanas and A. Kaplanyan. ‘N-Dimensional Gaussians for Fitting of High Dimensional Functions’. In: *ACM Digital library. SIGGRAPH 2024 Special Interest Group on Computer Graphics and Interactive Techniques Conference*. Vol. 38. SIGGRAPH ’24: ACM SIGGRAPH 2024 Conference Papers. Denver (Colorado), United States: ACM, 13th July 2024, pp. 1–11. DOI: [10.1145/3641519.3657502](https://doi.org/10.1145/3641519.3657502). URL: <https://hal.science/hal-04723115> (cit. on p. 21).
- [26] A. Jain, A. Sharma and K. S. Rajan. ‘Learning Based Infinite Terrain Generation with Level of Detailing’. In: *3DV 2024 - 11th International Conference on 3D Vision*. Davos, Switzerland, 18th Mar. 2024. URL: <https://inria.hal.science/hal-04474972> (cit. on p. 33).
- [27] Y. Kedadry and G. Cordonnier. ‘Fast simulation of viscous lava flow using Green’s functions as a smoothing kernel’. In: *Eurographics Digital Library. SCA 2024 - 23RD ACM SIGGRAPH / Eurographics symposium on computer animation - posters*. Montreal, Canada, 2024. DOI: [url{https://doi.org/10.2312/sca.20241172}](https://doi.org/10.2312/sca.20241172). URL: <https://hal.science/hal-04714171> (cit. on p. 29).
- [28] C. Nghiem, A. Bousseau, M. Sypesteyn, J. W. Hoftijzer, M. Agrawala and T. Tsandilas. ‘STiVi: Turning Perspective Sketching Videos into Interactive Tutorials’. In: *Graphics Interface (GI’24)*. Halifax, Canada, 3rd June 2024. DOI: [10.1145/3670947.3670969](https://doi.org/10.1145/3670947.3670969). URL: <https://hal.science/hal-04594231> (cit. on p. 18).
- [29] P. Papantonakis, G. Kopanas, B. Kerbl, A. Lanvin and G. Drettakis. ‘Reducing the Memory Footprint of 3D Gaussian Splatting’. In: *Proceedings of the ACM on Computer Graphics and Interactive Techniques*. Vol. 7. 1. Philadelphia, United States, May 2024. DOI: [10.1145/3651282](https://doi.org/10.1145/3651282). URL: <https://hal.science/hal-04550892> (cit. on p. 24).
- [30] H. Yuan, A. Bousseau, H. Pan, C. Zhang, N. J. Mitra and C. Li. ‘DiffCSG: Differentiable CSG via Rasterization’. In: *SIGGRAPH Asia 2024 - 17th ACM SIGGRAPH Conference and Exhibition on Computer Graphics and Interactive Techniques in Asia*. Tokyo (Japan), Japan, 3rd Dec. 2024. URL: <https://inria.hal.science/hal-04733011> (cit. on p. 16).
- [31] H. Yuan, J. Xu, H. Pan, A. Bousseau, N. J. Mitra and C. Li. ‘CADTalk: An Algorithm and Benchmark for Semantic Commenting of CAD Programs’. In: *IEEE Xplore. CVPR 2024 - IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). Seattle, United States: IEEE, 16th June 2024, pp. 3753–3762. DOI: [10.1109/CVPR52733.2024.00360](https://doi.org/10.1109/CVPR52733.2024.00360). URL: <https://inria.hal.science/hal-04732991> (cit. on pp. 10, 15).

Doctoral dissertations and habilitation theses

- [32] N. Rosset. ‘Fast simulation of wind-obstacle interactions. Applications to desertscape modeling and car design’. Université Côte d’Azur, 10th Dec. 2024. URL: <https://theses.hal.science/te1-04907415> (cit. on p. 45).

11.3 Cited publications

- [33] B. Kerbl, G. Kopanas, T. Leimkuehler and G. Drettakis. '3D Gaussian Splatting for Real-Time Radiance Field Rendering'. In: *ACM Trans. Graph.* 42.4 (July 2023). DOI: [10.1145/3592433](https://doi.org/10.1145/3592433) (cit. on p. 9).