Research Statement

Frédo Durand
http://gfx.lcs.mit.edu/~fredo/

My research is concerned with the development of algorithms and numerical methods for image generation. The field of rendering has seen tremendous achievements. Feature films exhibit photorealistic computer imagery, and commodity hardware permits stunning games. However, two major challenges remain: the complexity and richness of scene data grow much faster than hardware capabilities, and authoring effective images is still a long and tedious task that requires skilled artists. My goal is to create tools for the efficient generation of compelling and effective pictures.

My previous contributions illustrate my methodology: I find it important to balance my work between theoretical basic research, where the quest for understanding dominates, and the development of practical tools. Image generation must be grounded on two important fundamentals: a solid mathematical and algorithmic framework on the one hand, and a perceptual basis coupled with an understanding of the user's needs, on the other. My PhD work on *visibility* focused on the production of more realistic images and on real-time display, via a deeper understanding of the geometric fundamentals of image generation. While traditional visibility dealt with the portion of a scene visible from a given viewpoint, modern algorithms require the treatment of more global queries, such as mutual visibility between pairs of objects or limits of umbra and penumbra. I was the first to develop a comprehensive theoretical framework that encapsulates all global visibility properties of a 3D scene in a data structure called the *3D visibility complex* [6,12]. This allowed me to develop novel practical tools that were able to address unsolved problems, such as global visibility queries between any pairs of objects [5]. During my post-doctoral studies, I added important dimensions to my work, including human perception and user intervention, through contributions in tone mapping, image-based rendering, and non-photorealistic rendering.

My research work is largely motivated by a multidisciplinary study I made on the topic of *depiction*. As the holy grail of photorealism seems at reach, the development of new approaches such as image-based rendering or non-photorealistic techniques raises the issue of the fundamental goals of rendering research. I undertook a detailed study of this subject, building upon the perception and visual art literature, notably through a new multidisciplinary class I taught at MIT. In a recent article [9], I recast rendering into the broader context of picture production, and show the often-overlooked complexity of depiction. In particular, the complex interactions and mapping between 3D object space and 2D pictorial space, and the optimization nature of depiction — obtaining the most effective picture given constraints — are crucial aspects that raise exciting research challenges. In the long term, my goal is to organize a multidisciplinary center about depiction, and to gather people from computer science, cognitive sciences, visual arts, as well as actors from the imaging industry. In what follows, I outline the type of research that I want to pursue, focusing on four sub-areas: real-time rendering, realistic rendering, image editing, and pictorial style. In addition to long-term goals and approaches, I include with each area several concrete milestones that are suitable for PhD topics.

Real-time rendering: The ever-increasing size of 3D models makes display acceleration a crucial issue. It can be addressed through two complementary approaches: fast culling of invisible portions of the scene, and model simplification. In addition, more attention must be paid to the integration of these techniques and to the design of strict real time rendering - i.e. with a time budget – that can adapt to the available resources. I am active in these areas; In particular my visibility culling technique was the first to treat the cumulative occlusion due to multiple blockers with respect to a volumetric portion of space [4]. The simplification of a model to a smaller number of primitives has for the most part relied on the greedy decimation of triangles, which can be seen as gradient descent in mesh space. Current solutions are very effective for outputs that still contain a significant number of primitives (several hundred triangles). In contrast, *extreme simplification* aims at producing models with a number of primitives orders of magnitude smaller. It raises additional problems, including proper error metric and that of local minima in mesh space. I propose to address extreme simplification through two new approaches: the use of different primitives, and novel optimization strategies. We recently introduced a new representation called *billboard clouds* that bridges the gap between purely geometric and image-based representations. I also want to develop simplification techniques leveraged to the higher-order curved primitives that are becoming available on graphics hardware. In order to achieve truly real-time display, I want to investigate the design of systems that can dynamically adapt the algorithms to both the scene and the hardware.

Realistic rendering: Realistic rendering is faced with three major issues: richer representations, more efficient algorithms, and the often-overlooked control of such algorithms. My work focuses on the design of efficient and controllable algorithms. As a crucial instance of the relationship between efficiency and control, a realistic rendering algorithm should rely on only one parameter that controls the quality/time tradeoff. This involves a better understanding of both the perceptual and numerical aspects of lighting simulation, which I propose to address through wavelet and Fourier-space studies of light transport, and the use of tools from robust statistics.

The problem of depiction can be viewed as an optimization problem, be it solved by the computer or by a user [9]. This makes the study of *parameter spaces* crucial. I want to make the complex and non-linear parameter spaces of graphics algorithms and models meaningful, uniform, and predictable. This requires perceptual tools such as visible difference predicators, as well as mathematical techniques such as Jacobian computation and manifold embedding. I am currently working on the parameter space of material models.

Image-based editing and pictorial techniques: Image-based modeling and rendering is at the interface between computer graphics and computer vision. It is concerned with sampled data such as images, as opposed to the purely geometric approach of traditional computer graphics. Important topics include the acquisition of 3D models from photographs, and the design, capture and rendering of appropriate sampled representations. The use of real images permits realistic pictures with little user intervention, but raises the important issue of controlling and altering the representations. Unfortunately, little work has been dedicated to the problem of *editing* in the context of image-based representations. In recent work, we have extended classical photo editing to images with geometric depth per pixel, which we call *image-based editing* [2]. We have shown how a single photograph can be converted to a suitable image-based representation through intuitive interactive tools. This permits the easy modeling and alteration of scenes of unprecedented richness. More generally, many applications, e.g., design, architecture, or special effects can greatly benefit from the ability to manipulate and alter rich discrete representations through intuitive interactive tools. Among the wealth of future research that this work has opened, I wish to pursue three specific goals: extension to multiple input images, to the temporal domain, and to higher dimensions. Challenges include spatial and/or temporal registration of editing operations, the adaptation of motion flow to an interactive context, or the development of high-level yet controllable tools. Image-based editing also require the design of fast display and editing algorithms that can work in real time on such complex representations with no preprocessing, since the data are altered continuously.

Pictorial techniques are related, because they alter the picture to address the limitations of the medium, such as flatness or limited contrast, or they aim at more effective pictures (e.g., by focalizing the gaze of the beholder). The treatment of contrast limitation was introduced only ten years ago into computer graphics through *tone mapping*, a process that takes an input image with a large range of intensities, and reduces its contrast while preserving the perceptual impression [11,23]. I have recently developed a spatially-varying filtering technique that address the contrast of most challenging photographic situations. However, other pictorial techniques remain largely unexplored, although they are fundamental parts of effective depiction, and can prove a key aspect of the digital photography revolution. They require the development of appropriate operators and interactive image editing techniques based on perceptual phenomena. One of my key projects in this direction is the development of digital photography techniques that decouple the measurement process from a pictorial creative process. The main idea is to vary the imaging parameters of the camera while measurements are taken (e.g. exposure, focus, polarization) and then develop appropriate reconstruction techniques to *a posteriori* control pictorial effects, such as depth of field, motion blur, or contrast management. The development of such reconstruction algorithms raises exciting scientific challenges.

Non-photorealistic rendering and style: The recent field of non-photorealistic rendering aims at the production of pictures that are inspired by traditional media such as pencil drawing or painting. It is motivated by the effectiveness of such styles to convey information, and also by the aesthetic and expressive quality of such rendering. As initiated in my study of depiction [9], I strongly believe that a comprehensive framework taking into account the perception and visual art literature is necessary to organize this domain and to develop techniques that are more inter-operable. My work focuses on the development of a principled framework, both at the theoretical and practical level. This is a necessary step towards a longer-term goal: the study, parameterization, and capture of the elusive notion of style.