Since we not only simulate the shape of plants but their complete biological behaviour, direct applications of our simulation model also concern agronomy. Small scale or local simulation can help understanding the reactivity of plants to various external conditions such as lighting, water supply or even pruning. Large scale simulations allow to virtually explore the different possibilities for optimizing plantations in regions where resources are scarce, for instance cultures in the tropical forests. Finally, virtual plant models can also be used to replace costly or unpracticable experiments on real plant canopies, which is the case for deriving global models of canopy reflectance and in many remote sensing applications.

Among direct commercial applications, we are planning to integrate the environment sensitivity of plants into the commercial version of the CIRAD plant growth simulator, called AMAP. This software is already used by designers of various application fields for integrating realistic plant models into their own scenes. It also provides plug-ins for modelization/animation software such as SoftImage™ and Maya™. Thanks to the results we obtained on simulating these different kind of interactions, and especially concerning lighting simulation, we will soon be able to bridge the gap between these plants and a complete virtual garden.

Links:
SOLEIL projcet:
http://www-imagis.imag.fr/SOLEIL/
AMAP Plant growth simulator:
http://www.cirad.fr/produits/amap/amapen.html

Please contact:
François Sillion – INRIA
Tel: +33 4 7661 5423
E-mail: Francois.Sillion@imag.fr

Cyril Soler – INRIA
Tel: +33 4 7661 5429
E-mail: Cyril.Soler@imag.fr

Image-Based Rendering for Industrial Applications

by Samuel Boivin and André Gagalowicz

Several techniques have been recently developed for the geometric and photometric reconstruction of indoor scenes. The MIRAGES team at INRIA-Rocquencourt designed a new analysis/synthesis method that can recover the reflectances of all surfaces in a scene, including anisotropic ones, using only one single image taken with a standard camera and a geometric model of the scene.

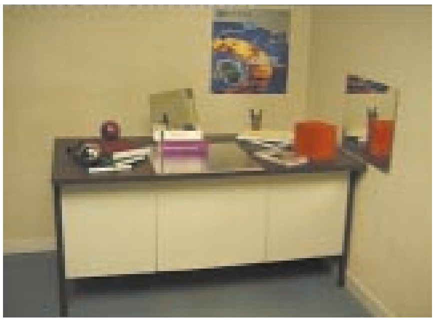
Figure 1: Real image captured with a camera.

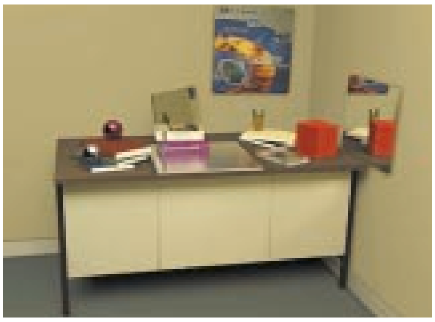
Figure 2: Synthetic image generated using our technique.

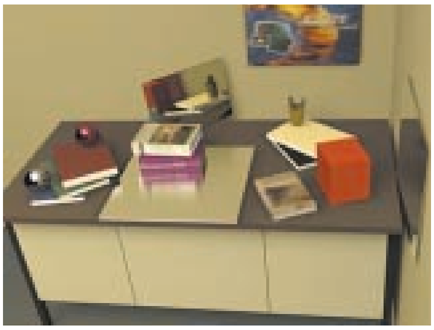
Figure 3: An application example of our technique in augmented reality: the observer's position has been modified.

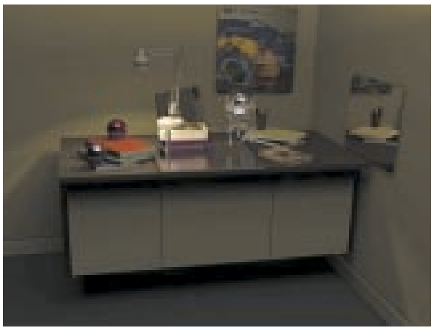
Figure 4: An application example of our technique in augmented reality: some synthetic objects have been added and the light sources positions have been modified. A new light source has been added on the desk to generate different illumination conditions. The legs of the desk have been cut.

The first step of the technique developed by MIRAGES is to capture a single photograph of the scene that we want to reconstruct. The second step is to position the 3D model (including the light sources), that can be built manually using the real image, thanks to a classical modeler such as Maya (Alias|Wavefront). Once the parameters of the camera (position, orientation, ...) have been automatically recovered, our image-based rendering software may recover a full BRDF (Bi-directional Reflectance Distribution Function) of all the objects in the scene. Starting with a first assumption about the reflectances of these surfaces, we generate a computer graphics image using an advanced photorealistic rendering software, called Phoenix, developed by our team. Phoenix is a global illumination renderer that can produce photorealistic images in a very short time (faster than current radiosity software). Phoenix takes advantage from the internal architecture of the computer, to greatly accelerate the process of image generation.

The first image created by Phoenix using the initial hypothesis about the reflectances of the objects, is directly compared to the real one applying a difference or a more complex operation. This error image is then used to correct the estimated reflectances until the result synthetic image becomes as close as possible from the real one. Our image-based rendering technique is iterative because it iteratively corrects the BRDF of the 3D surfaces, and it is hierarchical because it proceeds with more and more complex assumptions about the photometry of the objects inside the scene. If a particular reflectance property has been tried for a surface, and if it produces an image that still has important errors on all the surface area, this reflectance is automatically forced to be a more complex one (perfect specular or isotropic for example). Using these new assumptions, it is now possible to create a new synthetic image that will be again compared to the real image, to assess if the error between these two images decreased. If the error for an object is smaller than a user-defined threshold, then it is confirmed that its photometry has been correctly simulated and it does not need to have a more complex BRDF. Moreover, when there are high optical and thermal interactions between two objects (suppose for example that a book is lying on an aluminium sheet and that this book is reflecting on this sheet), our technique uses several complex algorithms to estimate the reflectances of these two objects, because they simultaneously interact on one another. This is a very complex procedure that can require several minutes of computation time, depending on the complexity of the BRDF of the surface: the most complicated case that we treat is anisotropic surfaces on which highly textured objects are reflected.

We show some results in this presentation, including a real image (figure 1) and the full (photometrically and geometrically) reconstructed synthetic image (figure 2). The mirrors and the aluminium sheet are simulated respectively as real mirrors and real complex BRDF surfaces: they are not simulated as textured objects which is the case of most of the techniques developed today. In other words, with one single image we know how to reconstruct mirrors, isotropic and anisotropic surfaces. This procedure possesses the ability, on one hand, to analyse through the mirror some objects that were not directly seen in the real image, and on the other hand, allows several industrial applications. The first application is called Augmented Reality and it is very used for special in movies in cinema and postproduction applications. Figure 3 presents an example of this application, with a modification of the viewpoint position, while figure 4 shows the same scene containing added synthetic objects under novel illumination conditions. A lot of other direct applications are possible, especially in high rate compression of video sequences (with no visible loss). Our technique can also be applied to automatic positioning of mirrors and light sources, advanced reflectance analysis of surfaces, etc.

As a conclusion, we have presented a new image-based rendering technique that is very fast and that recovers the BRDF (for diffuse, specular, isotropic, anisotropic and textured objects) of all surfaces inside a scene, including very complex ones. Our method is very efficient and needs very few data: one single image taken with a camera and a 3D geometric model (including light sources) of the scene are sufficient for the full reflectances recovery. Some future directions are currently investigated in order to develop new applications and to take into account scenes containing more complex BRDF and objects that combine both anisotropic reflections and texture properties.

Please contact:
Samuel Boivin – INRIA
Tel: +33 1 3963 5186
E-mail: Samuel.Boivin@inria.fr

Colour Texture Modelling

by Michal Haindl and Vojtěch Havlíček

The Textures Modelling Project carried out at CRCIM - UTIA and supported by the Czech Grant Agency, aims at developing mathematical models capable to simulate natural colour textures.

Virtual models require object surfaces covered with realistic nature-like colour textures to enhance realism in virtual scenes. Although such models can work also with digitised natural textures, synthetic textures are far more convenient, not only because of their huge data compression ratio (dozens model parameters irrespective of the required texture size) but can be designed to have certain desirable properties (eg, they can be made smoothly periodic).

Among several possible modelling approaches which are capable to learn a given natural texture sample (eg, fractals, random mosaic models, syntactic models, SARMA, etc.) the Markov random fields (MRF) are the most powerful and flexible ones, because choosing different neighbour sets and different types of probability distribution for the random variables a variety of distinct image types (sharp edges, patchy images, etc.) can be generated. Unfortunately MRF models require in general time-consuming Markov chain Monte Carlo synthesis methods and MRF analysis is complicated task even for simple MRFs.

Modelling multi-spectral images requires three-dimensional models. However if we are willing to sacrifice some spectral information a 3D model can be approximated with a set of simpler 2D models without compromising its visual realism. Random field based models quite successfully represent high frequencies present in natural textures though low frequencies are much more difficult for them. One possibility how to overcome this drawback is to use a multi-scale random field model. The resolution hierarchy provides then a transition between pixel-level features and global