

System for Acquisition, Processing and Material Rendering from Images

Andrey Ilyin, Andrey Lebedev, Vitaliy Sinyavsky, Alexey Ignatenko
Computational Mathematics and Cybernetics Faculty
Moscow State University, Moscow, Russia
{ailyin, alebedev, vsinyavsky, ignatenko}@graphics.cs.msu.ru

Abstract

Polygonal models of the objects from the real world often are used in virtual reality scenes. A lot of methods for creation and reconstruction of polygonal models are available nowadays. These models can be obtained as a result of the designers work in the corresponding 3D editing systems or by 3D laser scanning process of the real world objects and even by methods of geometry reconstruction from a photograph or video [1]. The acquisition of complex realistic models of the reflectance properties of materials and their photorealistic interactive rendering are still important problems in the computer graphics. The complexity of these problems is connected with the structural and reflectance properties of materials which determine the human perception of this material. In this work we focus on the process of the acquisition of the complex realistic models of reflectance properties of the material of the real world objects from their photographs.

Keywords: Image-Based Rendering, Material Reconstruction, BRDF, Phong's Model, Clusterization

1. INTRODUCTION

Nowadays there are several methods for the material reconstruction and algorithms for photorealistic interactive rendering [2]. The algorithms proposed and implemented in our system are based on image processing. Our image-based method can robustly determine different material models of the surface of the object. Several intuitively obvious user interaction tools were implemented for the ability of the convenient correction of automatically generated results.

For each material of the object the proposed algorithm can construct the corresponding BRDF (Bidirectional Reflectance Distribution Function) [3] reflectance model with required quality and adapt its properties for further hardware accelerated rendering. For each of the BRDF reflectance model of the material the rough approximations by analytical Phong model [4, 5] of this material can be calculated.

High-quality models of the materials of the real objects can be obtained by using relatively small amount of input data collected by non-professional means.

The presented system physically accurately reproduces reflective properties of objects and afterwards allows interactive rendering of objects with applied constructed materials. Along with high quality photorealistic renderer of the tabular BRDF we propose a fast interactive renderer for the material defined in an adaptive hardware oriented manner.

Note that the reconstruction of the geometry shape of the objects is not considered in this work provided that the geometry shape is known during the material reconstruction process.

The use of realistic models at all stages of the image synthesis is an essential condition for the photorealistic rendering. It is impossible to generate these models manually along with the ever growing need for more visual complexity. Therefore automatic and semi-automatic techniques for the acquisition of real world models are highly demanded.

Methods described in the literature (see Section 2 for details) are either used for spatial-invariant (isotropic) material acquisition, or assume certain type of material (for example, leather), or need precalculated values [6].

In this work we focus on the automatic process of the realistic reconstruction of real world object materials in which we obtain following results:

- the automatized process of the placement and fitting of the polygonal object to the virtual scene;
- the robust and effective process (BRDF model construction) for the reconstruction of the material reflectance properties based on photographs acquired from conventional digital camera;
- the process of the clusterization and decomposition of the object materials over the basis models of materials;
- the photorealistic rendering process of the tabled BRDF was implemented. Along with the interactive rendering of materials represented by Phong models we proposed adaptive method of BRDF usage on GPU

It is necessary to mention that the proposed algorithm needs relatively few photographs (about 10-20 for an object), which accelerates the acquisition process. For example materials obtained from a single image are shown on Figure 14.

As a result of the data gathering process, automatic fitting of the object, material reconstruction and clusterization we calculate the BRDF material models that best fits reflective properties of each material. Rendering of these materials is shown in Figure 11.b. For interactive rendering we propose algorithm which uses adaptive representation of the BRDF model and implemented on GPU (Figure 11.c). In cases where the main requirement is the speed we proposed acquisition and rendering of the materials represented by Phong model (Figure 11.d).

In the next section there is a brief overview of the existing reconstruction methods and results of our previous research. In the Section 3 we give a description of the automatic process of the placement and fitting of the polygonal object in the virtual scene. The process of generation and storage of geometric data is described in Section 4. Then the reconstruction process is

described based on: the process of the initial image segmentation (Section 5), method of the BRDF basis models construction (Section 6) and the stage of the material merging (Section 7). Final material surface clusterization and rendering methods are described in Sections 8, 9. Finally, obtained results and future research plans on this topic are described in Section 10.

2. RELATED WORK

The image-based material representation receives nowadays growing attention in the computer graphics [7].

For the description of the reflectance properties of the material several lighting models can be used: Phong [4], Blinn-Phong [4], cosine lobe [3], Torrance-Sparrow (and its modifications [8]), Lafortune [8] and arbitrary BRDF [3]. In our previous work [5] we also limit reflectance properties of the materials by Phong's model. But in practice the use of this model turns out not very convenient. Particularly, Phong's model is not physically accurate (energy is not conserved) and can represent small class of the isotropic material. We also found that as a result of bad accuracy of collection of the object information and the sparse data after optimization process we obtain models with strongly distorted specular component of the material.

Some methods require the use of complex devices during the acquisition step [2, 7], others need only the photographs and the calibration information of the camera (orientation, position) [9]. There are several methods by which one can obtain the material and geometry information at once. These methods use only a set of the photographs but requires that the scene contains a shiny spherical object. In this case the direction to the light source can be easily obtained from the glare on the spherical object. It is also assumed that the analyzed object must perfectly reflect light [10]. In our work we use methods similar to [2].

At the initial reconstruction stage we use the image segmentation algorithm based on region growing [11].

The hardware acceleration is also widely used for rendering and also for the acquisition process. Most of the methods are optimized for the modern GPUs [3, 7].

Methods which work with arbitrary BRDFs often use the parameterization and factorization [12, 13]. We propose high quality photorealistic rendering mode implemented on CPU. And the interactive rendering of the adopted BRDF models was realized on GPU (similar to [14]).

3. OBJECT FITTING

There are no specific restrictions in our work on the set of the input photographs, they can be obtained with a conventional digital camera (Figure 1). But for virtual scene definition we must know the information about light sources and about calibration parameters (intrinsic and extrinsic) of the camera from which the photographs were made. For automatic calibration the special object is placed in the scene (a paper list with a 'check board' image). The determination of the camera parameters (camera calibration) using a photographs sequence is done by using GML Camera Calibration Toolbox [15]. To determine the position of the light sources we make a photograph from the light source position and then calibrate the camera using this image.



Figure 1: Samples of the test photographs.

The virtual scene is defined by determining the correspondence between images, cameras and light sources. Based on the photometric information a polygonal object must be placed directly at the same position as the examined one. Note that this task cannot be properly solved manually. Unfortunately, the methods of feature point detection also cannot be applied directly to this problem. The problem is that the materials of the fitting object are unknown at this stage so it cannot be rendered with real materials.

This task of the augmented reality [16] in our work is based on the iteration process solved automatically. The embedding process is determined by two stages: the retrieval of the initial approximate object position and the further iterative fitting process.




Figure 2: Approximate object position.

The initial (approximative) object position can be found by combining the centers of the object on the images and the center of the polygonal model (Figure 2). We apply binarization process to the initial photographs sequence in which we also search and accumulate information about places with accurate borders of the object. Binarization process generates error growth because of the noncontrast borders and self shading (regions falsely become a background). For every binarized image we construct the set of certain descriptions of the object: S – area description, l – object elongation, α – orientation of the main axis of persistence, b – regions with accurate borders.

For every image we look for the center of the object P and convert its coordinates from the image plane coordinates to the 3D coordinate space of the virtual scene. For each camera and corresponding photograph lines (1) are constructed. The area of the most probable center of the object obtains by finding the intersection region of the lines l_1, \dots, l_n .

$$l_i : \vec{R}_i = O_i + \vec{V}_i t, \quad (1)$$

where O_i - is the center of the i th camera, \vec{V}_i - direction to the center of the object.

Each iteration of the process of automatic fitting of the object is composed of two stages: rotation and translation of the object. At first consecutive rotations of the polygonal object relatively to its own coordinate axis are applied (by using rotation operator, see Figure 3). For each new position we render object with white color and black background. From all of these images the closest to the binarized photography image are chosen. This position becomes the current position. For the translation operation we make similar operations. For the error comparison of the result of every iteration the error functionals are introduced (2, 3). Their minimization defines the required position.

$$\sum_{i=1}^n \rho_i(M_1, M_2), \quad (2)$$

where ρ_i is the value of the functional for i th pair of images, $M(S, l, \alpha, b)$ - is the object description on the binarized image, n - is the number of photographs;

$$\sum_{i=1}^n \mu_i(S_1, S_2, S_{\cap}), \quad (3)$$

where $\mu_i(S_1, S_2, S_{\cap})$ - the value of the functional μ for i image, S_1, S_2 - area characterizations of the pair of the binarized images, S_{\cap} - the area characterization on the intersection of S_1 и S_2 .



Figure 3: The rotation and translation operator.

4. STORAGE OF THE GEOMETRIC INFORMATION

The proposed algorithm for the reconstruction of the reflectance properties of the material requires the information on the local geometric properties of the surface.

This information can be conveniently represented by images. For each camera position special maps (textures) are generated for the position, normal and tangential information of the surface. The color of the pixels in each map are encoded by the set of three integers $r, g, b \in [0, 255]$ corresponding to the red, blue and green component of the color. We describe below the methods for the construction of these maps.

The color of the pixel in normal map (Figure 4.d) is uniquely defined by the coordinates of the unit normal vector:

$$r = 128(n_x + 1), \quad g = 128(n_y + 1), \quad b = 128(n_z + 1), \quad (4)$$

where n_x, n_y, n_z - coordinates of the normal vector.

The tangential maps are constructed in the same manner. The tangential vector can be computed by the following formula: $\vec{t} = [\vec{n} \times \vec{c}]$, \vec{t} - tangential vector, \vec{n} - normal vector, \vec{c} - some constant vector, not collinear to \vec{n} . By substituting t_x, t_y, t_z values into (4) we obtain the similar formula of the color of the pixel in the tangential map (Figure 4.b).

The position map (Figure 4.c) represents the offset of the objects surface points from the center of the object. The color of its pixels is defined as follows:

$$r = 128(\Delta_x Dim_x + 1), \quad g = 128(\Delta_y Dim_y + 1), \quad b = 128(\Delta_z Dim_z + 1),$$

where Δ - the point offset relatively to the center of the object, Dim - the object dimension.




Figure 4: a) the color map; geometric information maps: b) tangential map, c) position map and d) normal map.

5. INITIAL SEGMENTATION OF THE SURFACE OF THE OBJECT

After the position of the object in the scene is determined and the geometry information is properly stored in the corresponding geometrical maps we can start the initial segmentation process. The initial (rough) determination of the areas of the object with same materials is the purpose of this process. For simplification of the searching process of the correspondence of the pixels across photographs we convert photographs of the object to its textures. This step ensures one-to-one correspondence of the pixels on all of the textures and guaranties the fast searching process through them.

The segmentation algorithm is based on the formation of the homogeneous regions and the homogeneity criterion is the color and the spatial proximity of the pixels of the object. The region growing [11] algorithm was implemented. The straightforward use of this algorithm does not give acceptable results: there were significant problems during the process of the analyses of the object surface with different lighting condition.

In the Figure 5 one of these problems is shown schematically: the glares either define individual materials or wrongly link actuality different materials.




Figure 5: The glare. The wrong result of the region growing process.

Since on each texture the glares most likely must be at different places, for the effective suppression of the glares we proposed to average all textures of the object (Figure 6). Good results were obtained by this method. In fact, if the same region of the object in all textures lies in the glare area, then no algorithm can recognize materials in that area based only on the color information.




Figure 6: The textures averaging process.

Then the region growing algorithm is applied to the averaged texture which takes into account the total color and apply special conditions of the boundaries of the regions (Figure 7). As a result of this process each extracted area is marked with a unique marker.



Figure 7: The initial segmentation process.

By discarding the bad regions with weight (the area measured in pixels) less then the threshold of 10-50 pixels we obtain up to 200-350 materials on the difficult objects. After post-filtering and

screening of the bad regions we go over to the process of material merging (Section 7) and generation of basis models of the materials. The process of the creation of the basis BRDF materials is described in the next section.

6. CONSTRUCTION OF BRDF MODELS

The definition of the reflectance properties of the materials by BRDF is the physically correct method. The choice of this model does not assume any special limitations on the class of the reconstruction materials (there is a possibility of reconstruction of the anisotropic materials exists). One of the advantages of the BRDF usage is that its generation process depends not only on color information from images but also on the geometric information of the object.




Figure 8: The element of the object surface.

We consider a small part of the surface of the examined object (Figure 8). The normal vector to the surface, position of the light source, camera parameters are known. Let \vec{l} be the vector to the light source, \vec{v} - vector to the camera, \vec{n} - the normal vector to the element. Then BRDF is defined by the following formula:

$$BRDF = \frac{L_0}{L_1 \cdot \cos \varphi_{in} \cdot dw}, \quad (5)$$

where L_0 - is the energy reflected from the object in the \vec{v} direction, L_1 - is the incoming energy in the \vec{l} direction, $\cos \varphi_{in} = -(\vec{l} \cdot \vec{n})$ - cosine of the angle between normal vector \vec{n} and vector to the light source, dw - differential of the solid angle in \vec{l} direction.

Actually BRDF is determined by two vectors \vec{l} and \vec{v} . There is also implicit dependence on the wave length of the light falling at the element from the light source. This dependence is taken into account by considering three values of the BRDF for the red, green and blue wave length. The value for the arbitrary wave length is the linear combination of the BRDF with the same coefficients for the R , G and B values.

Note that because of light reciprocity property BRDF is symmetric with respect to the vectors \vec{l} and \vec{v} .

To determine the parameters \vec{l} and \vec{v} we use maps of the geometric information (Figure 4). In this case, the BRDF is a function of four angles: φ_{in} - angle between \vec{n} and \vec{l} , φ_{out} - angle between \vec{n} and \vec{v} , θ_{in} - angle between \vec{v} and \vec{l} , θ_{out} - angle between \vec{v} and \vec{l} (Figure 8).

Running through the values of all color maps (photograph converted to the texture of the object Figure 4.a) the cell of the BRDF cell is defined by corresponding values of geometric information maps. The current value of the pixel of the color map is added to the given cell of the BRDF. Because of the reversibility of the light we at once adding two values into the table: one - the calculated one, and other - the symmetrical value.

The obtained table is irregular grid function. The work with this representation of the BRDF is not convenient so we convert the irregular grid to the uniform grid. At first we chose the step of the grid (in our work we were using following values: 8, 16, 32 and 64). Running through all values of the BRDF function we obtain the corresponding four-dimensional parallelepipeds of the grid. After that the process of averaging of the values to each parallelepiped is applied. Note that by this process we also verify that the function in the given area is single valued. In case of finding several clouds of values in a single cell the algorithm notifies about a probable merging of different materials and then the given region of the material map can be decomposed to the smaller ones or even deleted from the material markup map. The proposed averaging method also ignores fluctuations by discarding them. Conversely the empty cells tell of the lack of the information in the BRDF. Therefore the number of photographs (more precisely the number of different positions and orientation of the camera and light sources) affects the completeness of the information on the BRDF.

7. MATEREAL MERGING

The initial segmentation process forms a rather big number of the materials (Figure 7). By changing of the segmentation parameters we can reduce the number of materials but with simultaneously growing of the probability of different materials merging. The main task of the initial segmentation is to markup regions of the different materials. At this stage we merge together regions with same materials. The main task is to join all regions with the same materials to the single cluster. The clusterization algorithm is as follows. For each markup region of the initial segmentation the BRDF model is obtained.

This gives a set of grid functions that we need to somehow compare. The situation becomes even more complicated in the sense that in the majority of the grid points there is no information at all (the interpolation process is not applied yet and the grid can be quite coarse). This is due to the boundedness of the markup region and a small number of the photographs of the object. Therefore the comparison of the grid function is quite difficult. We propose the following method that is not based on the direct comparison of the four-dimensional grid functions. Consider an

object whose normals are non parallel, for example, a sphere. Then this sphere is visualized in the same scene as the original object with the material defined by the corresponding markup region (Figure 9).

The next problem is in the comparison of the visualized spheres. Due to the lack of the information about the BRDF the rendered spheres may contain “black holes” (indicating the absence of the information).




Figure 9: The examples of the markup regions and spheres rendered with the materials defined by the markup region.

We consider the intersection of the non black regions of the sphere. If the intersection is empty, then the materials are recognized as non-comparable. Otherwise we compare materials as follows:

Let J_1 be the image of sphere with the first material, J_2 image of the second sphere. The error function $f(J_1, J_2)$ is defined as follows:

$$f(J_1, J_2) = \sum_{y=1}^h \sum_{x=1}^w (J_1(x, y) - J_2(x, y)) \delta_1[x, y] \delta_2[x, y], \quad (6)$$

$$\text{where } \delta_i[x, y] = \begin{cases} 0, & \text{if pixel } J_i(x, y) \text{ is black} \\ 1, & \text{otherwise} \end{cases}$$

By means of this function the conclusion on the identity or difference of the materials is drawn. Rendering of grid BRDF with various resolution and the corresponding comparison with the initial material are shown in Figure 10.

Other metrics can be used for the comparison of the materials based on the color or the relationship between the R , G and B channels of the BRDF. A certain advantage of this method is that the user of the system can himself distinguish the materials according to the rendered spheres. If the algorithm tells that the materials are non-comparable, then one can compare spheres with interpolated materials or the spheres can be compared by the user himself. In this sense the algorithm is adaptive.




Figure 10: a) Photography, b) $BRDF_8$ $f = 5.2$, c) $BRDF_{16}$ $f = 2.3$, d) $BRDF_{32}$ $f = 0.8$.

8. FINAL CLUSTERIZATION OF MATERIALS AND RENDERING

After the basis materials have been obtained (clearly different materials are obtained automatically) a final clusterization of the material of the object takes place with the respect to them.

Solving an optimization process (7) the unmarked regions of the object are distributed among materials

$$E(m, p) = \sum_{i=1}^n (I_t(m, p) - I_r(p))^2 \xrightarrow{\forall m \in M} \min, \quad (7)$$

where $E(m, p)$ - is the error of the basis material m and point p . $I_t(m, p)$ - the BRDF value for material m and parameters of the point p , $I_r(p)$ - is the color of the pixel p on the photograph, M - is the collection of the basis materials, n - is the number of photographs.

Two ways of the final clusterization of the object materials are realized in the system. The first assumes the one-to-one correspondence of the points of the object surface to one of the basis materials. For its implementation there exist a lot of rendering methods. In the second method the material is represented as a linear combination of several basis materials (we limit ourselves to two, three) in which the coefficients are also defined by the problem (7) but this time for several basis BRDFs. For this most photorealistic material presentation method we currently propose only the software rendering algorithm implemented on CPU (Figure 10.b). For the rendering speed optimization many algorithms are optimized for hardware

acceleration on GPU. But unfortunately GPU have limited resources for material information storage and rendering. The volume of the grid functions grows as the fourth power of the dimensions of the four-dimensional BRDF array. Because of that we cannot efficiently use initial high resolution BRDF on GPU without its adaptation (dimension reduction). Our algorithm of the interactive rendering is also hardware accelerated. The material markup map and definitions of the basis BRDF are passed to GPU and then the material rendering is taking place (Figure 10.c). The next section shows some of the main results of our system.

9. RESULTS

In Figures 10 and 14 the rendering results of the test object are shown (with the same light conditions as in photography 10.a). The geometric shape of the object was constructed via laser scanning process. The polygonal model consists of 50 thousands of triangles. Image 10.b is generated by using of middle quality BRDF 32 (32 steps for each angle). The resulting 800x600 image was rendered on CPU with 1-3 fps speed. The image 10.c was hardware rendered on GPU with 100 fps speed and scaled to the BRDF 8 material model. The rendered image 10.d generated from materials represented by Phong model (speed more than 100 fps).



Figure 11: a) Photography; Rendering: b) $BRDF_{32}$,
c) $BRDF_8$, d) Phong model.

For the proper comparison the corresponding regions of the object are enlarged in Figure 11. The rendering quality is increased simultaneously with chose of the material model from Phong to BRDF 32. In addition to the human perception comparison of the results we proposed the numeric error measure process based on the image “subtraction”. For each of the material models results of error measure are shown on Figure 12.



Figure 12: Enlarged region of the images: a) Photography,
b) $BRDF_{32}$, c) $BRDF_8$, d) Phong model.




Figure 13: Error measure of the result comparison (the lighter the worse). a) $BRDF_{32}$, b) $BRDF_8$, c) Phong model.




Figure 14: a) Photography; b) Rendered with $BRDF_{32}$ material reconstructed from single image.

10. CONCLUSION AND FUTURE PLANS

A study of many existing material reconstruction systems was made and new methods for reconstruction and rendering were developed and implemented. We developed an application software for reconstruction and rendering of materials which demonstrates that proposed algorithms work. Note that the total process of the material recognition usually took about 15-20 minutes per object. The directions of the future work are as follows.

In the short-term we are planning to propose more effective representations of the BRDF models for hardware accelerated rendering. Then we are planning to address the problem of the reconstruction of the full information of the object (along with materials we are planning to reconstruct the geometric shape of the object).

This work was granted by RFBR research grant №08-01-00886-a.

11. REFERENCES

- [1] Richard Hartley, Andrew Zisserman, *Multiple view geometry in computer vision*. Cambridge University Press, New York, NY, 2001.
- [2] Hendrik P. A., Lensch Jan Kautz, Michael Goesele, Wolfgang Heidrich, Hans-Peter Seidel. *Image-Based Reconstruction of Spatially Varying Materials*. Proceedings of the 12th Eurographics Workshop on Rendering Techniques, 2001.
- [3] Chris Wynn. *An Introduction to BRDF-Based Lighting*, NVIDIA Corporation. <http://developer.nvidia.com/attach/6568>.
- [4] Yoichi Sato, Mark D. Wheeler, and Katsushi Ikeuchi. *Object Shape and Reflectance Modeling from Observation*. Proceedings of the 24th annual conference on Computer graphics and interactive techniques, 1997.
- [5] P.Sikachev, A.Ilyin, A.Ignatenko *User-Assisted Acquisition, Processing and Rendering of Materials from Images*. Proc. of Graphicon'2007, pp. 131-134, Moscow, Russia, June 2007.
- [6] Hendrik P. A. Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. *Image-Based Reconstruction of Spatially Varying Materials*. In *Rendering Techniques '01*, London, Great Britain, 2001, pages 104-115.
- [7] G. Müller, J. Meseth, M. Sattler, R. Sarlette, R. Klein *Acquisition, Synthesis and Rendering of Bidirectional Texture Functions* in *Computer Graphics forum*, Vol. 24, No. 1, March 2005, pp. 83-109.
- [8] Eric P. F. Lafortune, Sing-Choong Foo, Kenneth E. Torrance, Donald P. Greenberg. *Non-Linear Approximation of Reflectance Functions*. Proceedings of the 24th annual conference on Computer graphics and interactive techniques, 1997
- [9] Hartmut Schirmacher, Wolfgang Heidrich, Martin Rubick, Detlef Schiron, and Hans-Peter Seidel. *Image-Based BRDF Reconstruction*. In B. Girod, H. Niemann, H.-P. Seidel, Proc. *Vision, Modeling and Visualization '99*.
- [10] Aaron Hertzmann and Steven M. Seitz, *Example-Based Photometric Stereo: Shape Reconstruction with General, Varying BRDFs*. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 27, August 2005
- [11] Besl and Jain. *Segmentation Through Variable-Order Surface Fitting*, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, pp. 167-192, 1988.
- [12] Rafal Jaroszkiwicz. *Fast Extraction of BRDFs and Material Maps from Images*. Master's Thesis, 2003.
- [13] Jason Lawrence, Szymon Rusinkiewicz, Ravi Ramamoorthi. *Efficient BRDF Importance Sampling Using A Factored Representation*. *ACM SIGGRAPH 2004 Papers*. 496 – 505, 2004.
- [14] Chris Wynn. *Real-Time BRDF-based Lighting using Cube-Maps* NVIDIA Corporation
- [15] GML C++ Camera Calibration Toolbox <http://research.graphicon.ru/calibration/gml-c-camera-calibration-toolbox-5.html>
- [16] R.Raskar, G.Welch, H.Fuchs *Spatially Augmented Reality* Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, U.S.A.