ISBN: 978-972-8924-84-3 © 2009 IADIS

SINGLE-IMAGE APPEARANCE ACQUISITION USING GENETIC ALGORITHMS

Adolfo Munoz, Belen Masia, Alfonso Tolosa, Diego Gutierrez
Universidad de Zaragoza

ABSTRACT

Most of the current appearance acquisition methods require the use of specialized equipment, involved capture sessions and/or large data sets. We propose a single-image approach which greatly simplifies the process, and allows to estimate reflectance properties of both opaque and translucent objects. Given the under-constrained nature of an image-based approach, we leverage two well-known illumination models, Phong and the diffuse approximation, to reduce the high-dimensional parameter space. The formulation is that of an optimization problem and satisfactory results are obtained using an implementation based on genetic algorithms. Additionally, a study is carried out to provide guidance in the election of the configuration parameters of the genetic algorithm.

KEYWORDS

Appearance acquisition, inverse rendering, BSSRDF, reflectance, genetic algorithm.

1. INTRODUCTION

Realistic image synthesis requires precise reflection and scattering models of real-world materials. As rendering algorithms become more sophisticated, efficiently simulating all aspects of light transport, a new area of research has gained importance over the last few years: appearance acquisition. Capturing the appearance of an object implies obtaining its BSSRDF (BRDF in the simpler case of non-translucent materials), in order to be able to model the interaction of light with that object.

The problem of capturing the appearance of a certain material can be faced from different approaches. When the actual object whose material is to be captured is available, specific measurement equipment and methodology allow for the acquisition of the reflectance parameters (see Related Work for details). However, this may not be the case, and then image-based approaches offer an alternative, less costly both in terms of time and equipment. Our work focuses on this second approach, and in particular in single-image acquisition, as opposed to other methods which require large sets of images under different angles or lighting conditions.

Obtaining the reflectance parameters from an image falls within the more general problem of inverse rendering. In a traditional direct rendering approach the lights, material, camera position and geometry of a 3D scene are known parameters used in the generation of the final 2D image. However, in many cases it is useful to obtain unknown information of the 3D scene from a rendered image, a problem known as inverse rendering. This includes inverse lighting (i.e. estimating the position and characteristics of the light sources of the scene), estimation of the camera position and orientation, obtaining the geometry of the scene, and appearance acquisition. As the complete problem of inverse rendering is highly under-constrained, previous knowledge of any of these (lights, geometry, camera position or appearance) is usually leveraged to determine the rest of them. A detailed survey covering this topic can be found in the work by Patow and Pueyo (2003). In our case, the specific problem of appearance acquisition will be explored, assuming that the rest of the information of the 3D scene is known, and starting with a single image as input.

Our approach poses appearance acquisition as an optimization problem. Starting from an initial set of reflectance parameters, successive images are rendered (see Figure 2) and compared with the original input image until the objective function, defined as the error between both images, falls below a certain value (or alternatively until a maximum execution time is exceeded). The method we propose to solve this optimization problem is based on genetic algorithms, which have already proved efficient to solve different combinatorial optimization problems (Goldberg, 1989; Reeves, 1995). Genetic algorithms are probabilistic heuristic algorithms for search and optimization which apply the concepts of biological evolution: each string of parameters to optimize is analogous to a chromosome, and the way in which these strings are generated and evaluated when searching for the best solution applies the concepts of natural selection, reproduction and mutation. Whenever solving an optimization problem, falling into local minima of the objective function is always a concern, and in our case the objective function has a large number of them. Statistically, genetic algorithms have been demonstrated to be less prone to this problem than other well-known optimization methods, as mutation favors diversity, increasing the probability of overcoming local minima.

In this paper we show how genetic algorithms can be used to capture the appearance of an object in an image. Besides, we also provide insight into how to configure the parameters of genetic algorithms when applying them to the specific problem of appearance acquisition, and their influence on the behavior of the algorithm and the final result.

1.1 Related Work

An obvious choice to measure general reflection properties is using a gonioreflectometer (Li et al., 2006), but a complete characterization of a spectral, anisotropic BRDF may require up to 10^5 samples, so several optimization strategies have been introduced. By using a camera instead of a single photoreceptor, lots of samples can be obtained simultaneously (Ward, 1992). However, calibration issues need to be considered, which make measurements less precise.

More general solutions that include sub-surface scattering capture typically use complex measuring equipment (Matusik et al., 2002; Debevec et al., 2000; Goesele et al., 2004; Peers et al., 2006). Image-based approaches, while simpler in conception, usually require large sets of data acquired from different angles and/or lighting conditions (Yu et al., 1999; Lensch et al., 2003; Shen and Takemura, 2006; Ghosh et al., 2008). Reduction of these sets can be achieved by adding some knowledge of the geometry of the object whose optical properties are being captured (Boivin and Gagalowicz, 2001).

Wu and Tang (2006) separate the sub-surface scattering component of a BSSRDF, starting from a single image together with a set of diffuse priors. Other methods to capture a generalized BSSRDF from single images impose constraints on the positions of the camera and light sources (Wang et al., 2008). We refer the reader to the excellent work by Weyrich and colleagues (2008) for a more comprehensive overview of appearance acquisition techniques.

In our work, we are interested in exploring the feasibility of appearance acquisition of complex materials by using genetic algorithms. This approach has been successfully used before in the field of computer graphics for texture synthesis, analysis and parameterization (Sims, 1991; Salek et al., 1999), image-based simulation of facial aging (Hubball et al., 2008), image recognition (Katz and Thrift, 1994; Koljonen and Alander, 2006), or extraction of geometric primitives (Roth and Levine, 1994).

2. GENETIC ALGORITHMS

As any other optimization method, a genetic algorithm tries to find a set of variables, $(x_1, x_2, ..., x_n)$, so that the objective function, $F(x_1, x_2, ..., x_n)$, reaches its minimum (or maximum). Given that each possible set of input variables $(x_1, x_2, ..., x_n)$ is equivalent to a chromosome (i.e. an individual) and each parameter x_i is denominated $gene_i$, the analogy with the theory of evolution is immediate: starting from a population of n chromosomes, each of them delivers a solution to the problem, and only the chromosomes yielding the better solutions survive to produce the next generations and perpetuate their *genetic material*. Genetic diversity is completed by sexual reproduction and random mutations. This section gives an overview of how these algorithms work, but we refer the reader to Winter and colleagues' work (Winter et al., 1995) for a more comprehensive explanation on genetic algorithms and their application.

The algorithm consists of four steps: initialization, selection, reproduction and termination. Selection and reproduction are iterated until the condition for termination is reached.

Initialization. The first step implies the creation of an initial population of individuals (or sets of variables corresponding to the parameters we want to estimate). The genes of these individuals are generated randomly within the search space, unless any prior knowledge exists.

Selection. In order to apply the principle of natural selection, it is necessary to evaluate the performance of each generated individual. To do this, each individual is assigned a rating, called *fitness*, representing the proximity of that individual to the solution. Chromosomes are then ordered according to their *fitness* and the ones with the lowest *fitness* values are eliminated and substituted by the descendants of the surviving chromosomes (*based-on-rank selection*). This way only the genetic material delivering the best results is perpetuated.

Reproduction. This step entails the creation of the next generation using two genetic operators: *crossover* and *mutation*. Crossover is a genetic operator used for exchange of genetic material, in which two chromosomes are randomly selected and an exchange of genes between them is performed. Mutation, on the other hand, ensures genetic diversity from one generation of individuals to the next by randomly modifying the value of some genes.

Termination. Typical termination conditions of the iterative process are a solution being found which satisfies a certain minimum criterion, the specified maximum number of generations being reached or the solution found not being able to be improved any further.

In the following section, we present our adaptation of the genetic algorithms approach to the problem of appearance acquisition, and comment on some implementation details.

3. APPEARANCE ACQUISITION

3.1 Variables and Objective Function

To be able to run genetic algorithms for appearance acquisition, we first need to define the variables and the objective function. Our method works both for opaque and translucent materials, and the objective function is an error function, defined in both cases as the difference between the input image and the image rendered in each iteration with the estimated parameters. However, the reflectance parameters being sought, i.e., the variables of our optimization problem, differ in each case. In order to reduce the dimensionality of the problem, we assume that other parameters such as the lighting or the geometry of the scene are known.

It should be noted that whilst the parameters are calculated in RGB, the error function, i.e. the difference between the original and the successively rendered images, is computed in the YCrCb color space, which is a perceptual color space (Poynton, 1996). This yields better results than comparing the images in non-perceptual color spaces such as RGB.

3.1.1 Opaque Materials: Phong Illumination Model

For opaque materials, interaction of light with the surface of the object is rendered using the Phong model (Phong, 1973), but more complete models such as Blinn-Phong could be modeled as well. The illumination on a certain point *p* on the surface is obtained as the sum of the ambient, diffuse and specular components as:

$$I_p = k_a I_a + \sum_{lights} (k_d (\vec{L} \cdot N) I_d + k_s (\vec{R} \cdot \vec{V})^{\alpha} I_s) \qquad (1)$$

$\vec{L}$ indicates the direction of the rays of light from a light source to a point of the surface, N is the normal to the surface, $\vec{R}$ indicates the specular direction and $\vec{V}$ the direction joining the point and the camera. I_a, I_d and I_s are the ambient, diffuse and specular intensities, respectively. The parameters which need to be estimated by the algorithm are the ambient, diffuse and specular reflection constants k_a, k_d and k_s, respectively, plus the Phong exponent α.

Determining these would imply obtaining 12 parameters (four for each of the three channels of the color space). However, in order to reduce the complexity of the problem we have made some assumptions. In most cases k_s and α have no significant spectral dependency, and therefore in our model we consider one same value for the three channels of these parameters. Besides, the term corresponding to ambient illumination ($k_a I_a$) makes a constant contribution throughout the image; we do not calculate it for simplicity and for further reduction of the parameter space, but doing so would be trivial. These reasonable assumptions reduce the number of parameters sought to five ($k_{d,R}$, $k_{d,G}$, $k_{d,B}$, k_s and α). Finally, as the interpretation of α from a perceptual point of view is not straightforward, in order to handle a variable which gives more perceptual information b is calculated instead of α. We define b as the angle covered by the specular highlight, and the relation between them is simply given by:

$$\alpha = \frac{-\ln 2}{\ln(\cos(b))} \quad (2)$$

3.1.2 Translucent Materials: Diffuse Approximation Illumination Model

When working with translucent materials the presence of subsurface scattering requires a more complex illumination model. We have used the diffuse approximation model described by Jensen et al. (2001), which decouples single and multiple scattering. Single scattering is obtained in a precise way, whereas multiple scattering is approximated by means of dipole diffusion. The complete BSSRDF describing the outgoing radiance at point x_0 in direction $\vec{w}_0$ is thus the sum of both components:

$$S(x_0, \vec{w}_0) = S_d(x_0, \vec{w}_0) + S^{(1)}(x_0, \vec{w}_0) \quad (3)$$

where S_d and $S^{(1)}$ represent multiple and single scattering respectively. These terms are given by:

$$S_d(x_0, \vec{w}_0) = \frac{1}{\pi} F_t(\eta, \vec{w}_0) \int_A R_d(\| x_0 - x_i \|) E(x_i) dA(x_i) \quad (4)$$

$$S_0^{(1)}(x_0, \vec{w}_0) = \sigma_s \int_{2\pi} \int_0^{\infty} Fp(\theta) \Psi(s_i' + s) S(x_i, \vec{w}_i) ds d\vec{w}_i \quad (5)$$

where $E(x_i)$ represents incoming irradiance, $F=F_t(\eta, \vec{w}_0)F_t(\eta, \vec{w}_i)$ is the product of two Fresnel transmittances, s'_i and s indicate scattering paths and Ψ is an exponential attenuation function. R_d is the diffuse reflectance function, which is computed as (see Table 1 for a definition of all the symbols):

$$R_d(r) = \frac{\alpha'}{4\pi} \left[z_r (\sigma_{tr} + \frac{1}{d_r}) \frac{e^{-\sigma_{tr} d_r}}{d_r^2} + z_v (\sigma_{tr} + \frac{1}{d_v}) \frac{e^{-\sigma_{tr} d_v}}{d_v^2} \right] \quad (6)$$

Using this formulation (for the complete details and derivation of these equations, consult Jensen et al. (2001)), it can be shown that the diffuse approximation model depends on only four parameters (Xu et al., 2007): σ_s (scattering coefficient), σ_a (absorption coefficient), η (relative index of refraction) and p (normalized phase function), which may additionally show spectral dependencies.

As previously done with the Phong model, in practice the parameter space is reduced making reasonable assumptions in order to simplify the problem without compromising the final results. The final string of parameters to estimate consists of: η (assuming the same value for the three color space channels), g (mean cosine of the scattering angle, from which p can be calculated using a Henyey-Greenstein formulation), $\sigma_{s,R}$, $\sigma_{s,G}$, $\sigma_{s,B}$, $\sigma_{a,R}$, $\sigma_{a,G}$, $\sigma_{a,B}$ and b (related to the Phong exponent as shown in equation (2) and used to render the highlights of the translucent object), i.e. 9 parameters.

Table 1. Symbols used in the formulation of the diffuse approximation model

σ_s	Scattering coefficient	η	Relative index of refraction
σ_a	Absorption coefficient	p	Normalized phase function
$\sigma_t=\sigma_a+\sigma_s$	Extinction coefficient	$F_{dr}=-1.440/\eta^2+0.710/\eta+0.668+0.0636\eta$	
$\sigma'_s=(1-g)\sigma_s$	Reduced scattering coefficient	$A=(1+F_{dr})/(1-F_{dr})$	
$\sigma'_t=\sigma_a+\sigma'_s$	Reduced extinction coefficient	$z_r=1/\sigma'_t$, $z_v=z_r(1+4A/3)$	
$\alpha'=\sigma'_s/\sigma'_t$	Reduced albedo	$r=\|\|x_i - x_0\|\|$	
$\sigma_{tr} = \sqrt{3\sigma_a\sigma'_t}$	Effective extinction coefficient	$d_r = \sqrt{r^2 + z_r^2}, d_v = \sqrt{r^2 + z_v^2}$	

3.2 Implementation of the Algorithm

We provide here some insight on how the genetic algorithms framework maps to our appearance acquisition problem. A discussion of the influence of the specific configuration parameters is provided in subsection 3.3.

The first step of any genetic algorithm is *initialization*. A set of chromosomes consisting of strings of reflectance parameters (nine in the case of translucent and five in the case of opaque materials, as explained in the previous subsection) are set. In this first generation the parameters could take random values within the search space, but, in order to accelerate convergence, we fixed the initial estimations of the different parameters to common values, shown in Table 2. The number of chromosomes created is a configuration parameter of the algorithm.

Table 2. Initialization values of the sought parameters

Phong model			Diffuse approximation model				
$k_{d,(R,G,B)}$	k_s	b	η	g	b	$\sigma_{s,(R,G,B)}$	$\sigma_{a,(R,G,B)}$
0.5	0.5	10	1.3	0	10	1	1

An image is then rendered for each of the chromosomes created in each generation to calculate the *fitness* value of each chromosome and thus perform the *selection* step. This *fitness* value is calculated with a per-pixel least squares function measuring the difference between the individual channels in the original and the rendered images in the YCrCb space. The set of parameters delivering the most approximate solution are used to create the next generation. The number of chromosomes being replaced is another configuration parameter of the algorithm.

Once the best chromosomes have been selected, *reproduction*, involving crossover and mutation, takes place. In our implementation crossover is performed at only one point of the chromosome (i.e. two *parent* chromosomes are cut at one point and one part of each combined to form the *child* chromosome), which has proven enough for our objectives, but more complex crossover procedures are also possible. During mutation, gene values vary between $\pm$0-30% of their original value. Our research shows that greater variations introduce a too random behaviour and control over the evolution of the algorithm is easily lost, whereas very small variations need many generations for the algorithm to reach a valid solution.

The processes of selection and reproduction continue iteratively until the termination condition is met. Given that our goal is to study the effectiveness of the algorithm and the influence of its configuration parameters on the final result, we simply define our termination condition as a fixed number of generations. This suffices in our context, although changing the termination condition to an error threshold is straightforward.

3.3 Parameter Space

Genetic algorithms have a series of input parameters (initial number of individuals, crossover and mutation probabilities, etc), whose correct configuration is vital in reaching a consistent solution within a reasonable execution time. In order to select the most adequate values for these parameters, we have performed a series of tests, taking into account both the accuracy of the final solution and the computation time required. The results of these tests for the most relevant configuration parameters are discussed here, and can be seen in Figure 1 for the case of the Phong model. The accuracy was measured as the difference between the real ground truth value and the value obtained by the algorithm, expressed as a percentage of the ground truth value.

Probability of replacement. The probability of replacement accounts for the percentage of individuals which are eliminated in each selection process. Following the evolution simile, the higher this probability is, the faster the population evolves. However, running times also increase significantly, as all the chromosomes and their corresponding image need to be created for each generation. As seen in Figure 1 (*left*) it is one of the most influential parameters of the algorithm, both in time and in accuracy of the result.

Probability of crossover. As explained in Section 2, crossover represents sexual reproduction and takes place after the selection and replacement process. The probability of crossover represents the percentage of individuals which are the result of combining the genes of two survivor chromosomes. Figure 1 (*center left*) shows how an increase in *sexual reproduction* (hence favoring genetic diversity) causes the percentage of error to decrease slightly. As the resulting images of more combinations of genes need to be calculated, execution time increases slowly.

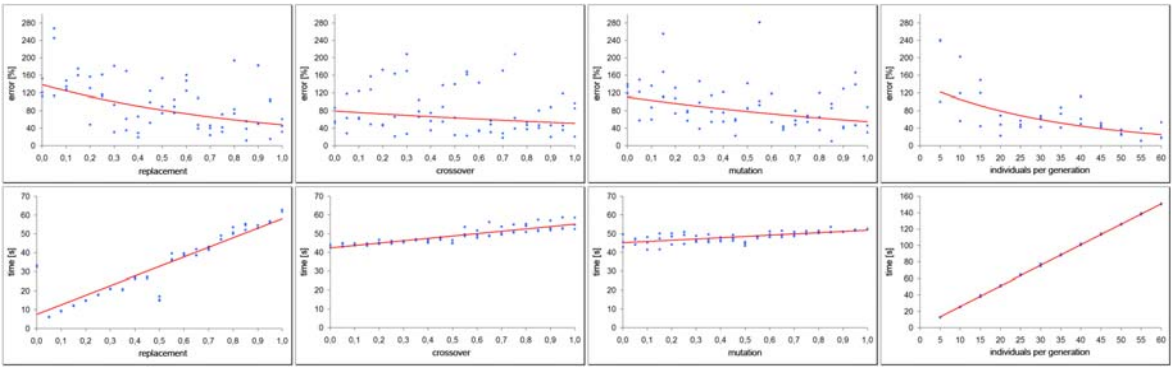

Figure 1. Percentage of error *(top row)* and execution time *(bottom row)* as a function of, *from left to right*, the probability of replacement, the probability of crossover, the probability of mutation and the number of individuals in each generation. Data obtained for the Phong model.

Probability of mutation. Representing the percentage of genes which mutate from one generation to the next, this probability is critical when working with a small number of individuals per generation. Variations in the initial genes are crucial to progressively reach the optimal solution and to avoid falling into local minima; the higher this probability, the lower the percentage of error with minimum time penalty (see Figure 1, *center right*).

Number of generations. The number of generations is, together with the number of individuals per generation discussed below, the parameter with the greatest influence. It indicates the number of generations which are created before the algorithm terminates and delivers a solution (alternatively, an error threshold can be trivially set as termination parameter). Figure 2 shows how the solution progressively evolves along generations. With an infinite number of generations, the solution would perfectly match the original. In practice, a compromise has to be found between execution time and accuracy of the solution, determined by the number of generations.

Number of individuals per generation. The effect of the number of individuals of each generation in the performance of the algorithm is straightforward: the more individuals, the least the percentage of error, as more possibilities are evaluated. However, there is a substantial increase in the execution time, as shown in Figure 1 (*right*).

4. RESULTS

We have presented a method based on genetic algorithms suitable for capturing the appearance of opaque and translucent materials depicted on a single image. The algorithm converges to an approximate solution in reasonable times with little user interaction. Figure 2 shows how the algorithm works, as the image rendered with the estimated parameters evolves through iterations converging to the input image. As our aim was to study the performance of the algorithm and the influence of the configuration parameters our termination condition was set to a maximum number of generations, i.e. 50, and not to a maximum error between the final and input images.

Figure 2. Evolution of the result of the genetic algorithm as the number of generations increases. *Left:* original image from which the Buddha's reflectance properties are meant to be captured. *Rest (from left to right):* partial results every ten generations, showing convergence to the solution.

Table 3. Evolution of the values of the estimated reflectance parameters as the number of generations increases

	Ground truth	Initial value	Generation #10	Generation #20	Generation #30	Generation #40	Generation #50	Percentage of error
η	1.5	1.3	1.18195	1.29587	1.36760	1.35480	1.35480	9.7
g	0.5	0	0.43301	0.49859	0.50151	0.50151	0.50151	0.3
b	20	10	15.8393	17.6512	18.4832	18.4832	18.4832	7.6
$\sigma_{s,R}$	2.19	1	1.31844	2.26811	2.26811	2.07799	2.07799	5.1
$\sigma_{s,G}$	2.62	1	1.80584	2.76615	2.76615	2.54309	2.54309	2.9
$\sigma_{s,B}$	3	1	2.24059	2.42196	2.42196	2.91114	2.91114	3.0
$\sigma_{a,R}$	0.0021	1	0.04108	0.01706	0.01527	0.01527	0.00685	226.2
$\sigma_{a,G}$	0.0041	1	0.01207	0.03432	0.02742	0.02742	0.00916	123.4
$\sigma_{a,B}$	0.0071	1	0.00650	0.00477	0.00948	0.00948	0.00725	2.1

The actual ground truth values of the reflectance parameters and the evolution of the estimated values can be seen in Table 3. Even though the number of generations is not very large almost all errors are below 10%, and it can be seen that those parameters which have very large errors do not have a significant relevance perceptually in the final result.

Additional results are shown in Figure 3. The probabilities of replacement, crossover and mutation were all fixed to 0.8, the number of generations was 50 and the number of individuals in each generation was set to 40. All images in this paper have been rendered on an AMD Opteron Quad-core machine @3*GHz* and 4*GB* of RAM, and took between 15 and 20 minutes in the case of translucent materials and around one minute in the case of opaque objects. For the diffusion approximation, we have used the fast hierarchical rendering technique of Jensen and Buhler (2002).

5. CONCLUSIONS AND FUTURE WORK

Obtaining the reflectance parameters of a certain material from a single image can be posed as an optimization problem in which the error between the original image and an image generated with the estimated parameters is the objective function. To solve it, an implementation based on genetic algorithms offers a solution within the specified error range in a reasonable execution time as long as the configuration parameters of the algorithm are chosen wisely.

Genetic algorithms have already proved useful in solving non-structured or inverse problems in many other fields, but configuring their input parameters remains a fundamental task which varies across applications, and has to be individually studied for each specific optimization problem. The correct election of these parameters is a key aspect for achieving a valid solution in an acceptable execution time, as expectedly more accuracy in determining the solution implies higher execution times. We have studied the influence of the different configuration parameters of the genetic algorithm in both the accuracy of the result and the execution time of the algorithm, providing guidance for future implementations. Special attention must be paid to the number of individuals per generation and to the percentage of individuals being replaced in each iteration, given their stronger influence in both computational cost and accuracy of the result.

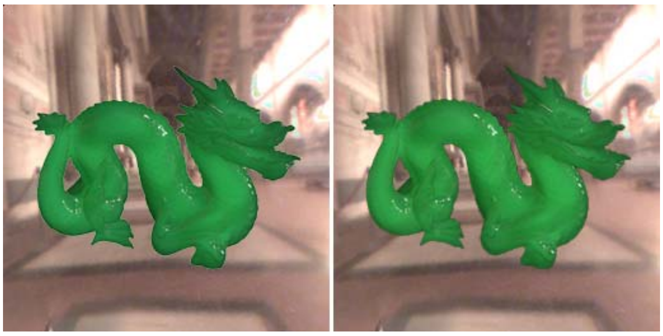

	Ground truth	Generation #50	Percentage of error
η	1.5	1.48675	0.9
g	0.5	0.34655	30.7
b	20	18.1909	9.1
$\sigma_{s,R}$	0.5	0.49969	0.1
$\sigma_{s,G}$	4.5	3.64219	19.1
$\sigma_{s,B}$	2.6	1.77309	31.8
$\sigma_{a,R}$	1.5	1.40514	6.3
$\sigma_{a,G}$	0.11	0.08945	18.7
$\sigma_{a,B}$	1	0.97393	2.6

	Ground truth	Generation #50	Percentage of error
$k_{d,R}$	0.3	0.3019	0.6
$k_{d,G}$	0.7	0.7043	0.6
$k_{d,B}$	0.3	0.3015	0.5
k_s	1	0.9686	3.1
b	40	37.439	6.4

Figure 3. *Left column:* Original image. *Right column:* Image rendered with our algorithm. In the top row the material is translucent, and thus modeled with the diffuse approximation model. The bottom row shows an opaque material, modeled with the Phong model. Rendering times are around 20 minutes for the translucent one and 1 minute for the opaque one. Tables next to each pair of images show the values of the sought parameters in the input image, the estimation obtained after 50 iterations of the algorithm and the relative error between them.

One of the main lines for future research is exploring the possibilities that mutation techniques can offer with the objective of accelerating convergence to the solution and of overcoming local minima. Besides, the operation with the highest cost is rendering the scene with each set of parameters for evaluation by comparison with the original image, so creating the chromosomes of possible solutions intelligently instead of relying on brute force is vital, and more sophisticated mutation functions could also help in this direction.

To reduce the parameter space, we have assumed that information of light sources, geometry and camera position was known, and only the reflectance characteristics of an object in the image were unknown. It would be interesting to stress our approach further and see how genetic algorithms perform as the problem becomes even more ill-posed.

Further strategies which can improve the implementation include parallelization. As mentioned, the bottleneck of the implementation lies in generating an image for evaluation for each string of parameters; given that these strings are completely independent between them, several evaluations could be performed in parallel to reduce the execution time.

ACKNOWLEDGEMENTS

We would like to thank the anonymous reviewers for their valuable comments. This research has been funded by the Spanish Ministry of Science and Technology (TIN2007-63025), the Diputacion General de Aragon (CONAID) and CAI (project IT21/08).

REFERENCES

Boivin, S. and Gagalowicz, A., 2001. Image-based rendering of diffuse, specular and glossy surfaces from a single image. *SIGGRAPH*, ACM. Los Angeles, California, pp. 107-116.

Debevec, P. et al, 2000. Acquiring the reflectance field of a human face. *SIGGRAPH*, ACM. New Orleans, LA, 145-156.

Ghosh, A. et al, 2008. Practical modeling and acquisition of layered facial reflectance. In *ACM Trans. Graphic.* (Proceedings of SIGGRAPH Asia 2008), Vol. 27, No. 5, pp. 1-10.

Goesele, M. et al, 2004. Disco: acquisition of translucent objects. *SIGGRAPH*, ACM. Los Angeles, CA, pp. 835-844.

Goldberg, D. E., 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Publishing Co., Inc., Redwood City, California.

Hubball, D. et al, 2008. Image-based aging using evolutionary computing. In *Comp. Graph. Forum*, 27, 2, pp. 607-616.

Jensen, H. W. and Buhler, J., 2002. A rapid hierarchical rendering technique for translucent materials. *SIGGRAPH*, ACM. San Antonio, Texas, pp. 576-581.

Jensen, H. W. et al, 2001. A practical model for subsurface light transport. *SIGGRAPH*, ACM. Los Angeles, pp. 511-518.

Katz, A. and Thrift, P., 1994. Generating image filters for target recognition by genetic learning. In *PAMI*, Vol. 16, No. 9, pp. 906-910.

Koljonen, J. and Alander, J., 2006. Genetic algorithms for optimizing fuzzy image pattern. *Proceedings of the Ninth Scandinavian Conference on Artificial Intelligence (SCAI 2006)*. Helsinki, Finland, pp. 46-53.

Lensch, H. P. A. et al, 2003. Image-based reconstruction of spatial appearance and geometric detail. In *ACM Transactions on Graphics*, Vol. 22, No. 2, pp. 234-257.

Li, H. et al, 2006. Automated three-axis gonioreflectometer for computer graphics applications. In *Optical Engineering*, Vol. 45, No. 4, pp. 1-11.

Matusik, W. et al, 2002. Acquisition and rendering of transparent and refractive objects. *Thirteenth Eurographics Workshop on Rendering*, EG. Pisa, Italy, pp. 267-277.

Patow, G. and Pueyo, X., 2003. A Survey of Inverse Rendering Problems. In *Comp. Graph. Forum*, 22, 4, pp. 663-687.

Peers, P. et al, 2006. A compact factored representation of heterogeneous subsurface scattering. *SIGGRAPH*, ACM. Boston, Massachusetts, pp. 746-753.

Phong, B., 1973. Illumination of computer-generated images. PhD thesis, University of Utah.

Poynton, C., 1996. *A Technical Introduction to Digital Video*. John Wiley & Sons, Inc.

Reeves, C. R. (ed.), 1995. "Modern Heuristic Techniques for Combinatorial Problems". In *Advanced Topics in Computer Science*. McGraw-Hill, Chap. 4.

Roth, G. and Levine, M., 1994. Geometric primitive extraction using a genetic algorithm. In *PAMI*, 16, 9, pp. 901-905.

Salek, P. et al, 1999. Application of genetic algorithms to texture analysis. In *Crystal Research and Tech.*, 34, 1073-1079.

Shen, L. and Takemura, H., 2006. Spatial reflectance recovery under complex illumination from sparse images. *Computer Vision and Pattern Recognition*, IEEE. New York, NY, pp. 1833–1838.

Sims, K., 1991. Artificial evolution for computer graphics. *SIGGRAPH*, ACM. Las Vegas, Nevada, pp. 319–328.

Wang, J. et al, 2008. Modeling and rendering of heterogeneous translucent materials using the diffusion equation. In *ACM Transactions on Graphics (TOG)*, Vol. 27, No. 1, pp. 1-18.

Ward, G. J., 1992. Measuring and modeling anisotropic reflection. *SIGGRAPH*, ACM. Chicago, Illinois, pp. 265-272.

Weyrich, T. et al, 2008. Principles of appearance acquisition and representation. *SIGGRAPH*, ACM. Los Angeles, California, pp. 1-119.

Winter, G. et al, 1995. *Genetic algorithms in engineering and computer science*. John Wiley & Sons, Inc.

Wu, T. P. and Tang, C. K., 2006. Separating subsurface scattering from photometric image. *International Conference on Pattern Recognition (ICPR)*, IEEE Computer Society. Hong Kong, China, pp. 207-210.

Xu, K. et al, 2007. Real-time homogenous translucent material editing. In *Computer Graph. Forum*, 26, 3, pp. 545-552.

Yu, Y. et al, 1999. Inverse global illumination: Recovering reflectance models of real scenes from photographs. *SIGGRAPH*, ACM. Los Angeles, California, pp. 215-224.