

Lighting Design by Simulated Annealing

F. Anrys and P. Dutré

Report CW 393, September 2004



Katholieke Universiteit Leuven
Department of Computer Science
Celestijnenlaan 200A – B-3001 Heverlee (Belgium)

Lighting Design by Simulated Annealing

F. Anrys and P. Dutré

Report CW 393, September 2004

Department of Computer Science, K.U.Leuven

Abstract

In this paper, we present a lighting design framework for glossy real objects, starting from the reflectance field of an object. Using existing digital imaging software, the lighting designer paints the desired illumination distribution on a photograph of an object. This painted-on illumination is used to determine the light intensities which produce a shading of the real object matching the desired illumination distribution. Certain areas in the painted image can be favored among others by weighting their importance in the optimization algorithm. Our optimization algorithm is driven by a simulated annealing approach to find fixed-sized lighting combinations and accompanying intensities.

Lighting Design By simulated Annealing

F. Anrys[†] and P. Dutré[‡]

Department of Computer Science, Katholieke Universiteit Leuven

Abstract

In this paper, we present a lighting design framework for glossy real objects, starting from the reflectance field of an object. Using existing digital imaging software, the lighting designer paints the desired illumination distribution on a photograph of an object. This painted-on illumination is used to determine the light intensities which produce a shading of the real object matching the desired illumination distribution. Certain areas in the painted image can be favored among others by weighting their importance in the optimization algorithm. Our optimization algorithm is driven by a simulated annealing approach to find fixed-sized lighting combinations and accompanying intensities.

Categories and Subject Descriptors (according to ACM CCS): I.3.4 [Computer Graphics.]: Graphics Utilities. Paint Systems. I.3.6 [Computer Graphics.]: Methodology and Techniques. Interaction techniques. F.2.1 [Analysis Of Algorithms and Problem Complexity.]: Numerical Algorithms and Problems. Computations on matrices

1. Introduction

Given a full scene specification such as the geometry, material properties and camera parameters, the look and feel of an environment highly depends on the lighting setup. Good lighting can significantly improve the character of an environment, reducing under- or over-illuminated regions, enhancing contrast, etc. In certain spaces (workshops, storage, utility rooms) one is usually interested in bright, shadowless light. But in other areas (lobbies, boardrooms, restaurants, museums) one might be looking for a little style, drama, or have the lighting be part of an artistic expression.

Inverse lighting design starts from desired illumination objectives and constraints, such as the position of shadows, the intensity of highlights, or the overall brightness level of the environment. Using a cost function for each of the desired targets, an optimization process determines the most optimal positioning and color of light sources to satisfy the objectives and constraints of the lighting design. This inverse approach requires an *a priori* knowledge of the desired appearance.

With the availability of advanced rendering techniques

and global illumination algorithms, an approach widely used in lighting design, is to make a 3D model of the real environment and find a suitable lighting design for the virtual environment. While this is a very flexible approach in exploring the space of possible lighting designs, applying a design back in reality still remains a problem. Building an exact 3D model with appropriate material parameters for a real environment or object can be a demanding task. One has to address all these problems when performing virtual simulations.

In this paper, we present a novel inverse lighting design technique for real objects, starting from the reflectance field of an object. The designer is able to specify the desired illumination by painting on a photograph of the object using existing digital image software, such as Adobe Photoshop. A simulated annealing algorithm is then used to work backwards and establish a lighting configuration, suiting the designers goals. Additional input can be provided to the optimization algorithm by giving painted areas different importance. Afterwards, we show our lighting design can be reproduced in reality.

Since we only use a reflectance field, real world lighting design for which virtual descriptions of an object are difficult to construct, can benefit from our approach. For instance, designing a lighting configuration for fragile or highly valuable objects for which it is difficult to acquire the geometry

[†] e-mail: Frederik.Anrys@cs.kuleuven.ac.be

[‡] e-mail: Phil.Dutre@cs.kuleuven.ac.be

from can be done with any type of light source. There are no constraints on the type of light sources as well. Our approach might also be applicable in stage design as long as photographs can be made of the possible lighting setup.

2. Previous and Related Work

Lighting design techniques address the problem of finding lighting parameters, such as position, orientation or intensity, that yield a desirable lighting configuration. Inverse lighting design techniques in computer graphics so far, have mostly addressed this problem for virtual environments.

Poulin [PRJ97], [PF92] allowed the designer to specify design goals by sketching shadows and highlights, from which the light source position and surface roughness are inferred. Related to this approach, [PTG02] presented a technique in which shadows are considered as modeling primitives. Costa [CSF99] presents a methodology in which fictitious luminaries can be defined and placed in a virtual scene to describe a desired radiance distribution. Kawai [KPC93] controls light emissions and directions, as well as surface reflectance to create a room for which the user will have a feeling of comfort, or to minimize the overall energy. Dorsey [DSG91] designed and simulated opera lighting and projection effects using global illumination techniques. An entirely different approach for exploring the space of possible design solutions is presented by *Design Galleries* [MAB⁺97]. Their interface presents the user with the broadest selection, automatically generated and organized, of perceptually different lighting configurations. Shackel et al. [SL01] determine various lighting parameters by optimizing quality objective functions such as contrast or histogram equalization.

Inverse lighting design using prerecorded photographs, the approach we follow in this paper, relates strongly to the field of image-based relighting [NSD94][DKNY95][GH00][DHT⁺00][DWT⁺02]. By photographing an object under different illumination conditions, the reflectance field of the object for a fixed viewpoint can be captured. Using this reflection field, the object illuminated by any environment map can be rendered as a linear combination of the basis images. The core technique in our paper is to inversely apply the image-based relighting approach: If an image of an object is given, what is the linear combination of basis images that resembles the target image as close as possible?

Our user interface is somewhat related with work presented by Schoeneman [SDS⁺93]. By painting light on the geometry present in the scene, they optimize the light source intensities necessary to obtain the desired shading. The number and positions of the light sources are known in advance. Using a reflectance field of an object, we remove the constraint of positioning the light sources in advance and determine the position as well as the intensities of the light sources for real objects.

3. Simlated Annealing

Simulated annealing in a nutshell: You are given a single function F that depends on one or more independent variables. You want to find the value of the variables where F takes on a maximum or minimum value. The computational effort of doing so is dominated by the cost of evaluating F . An extremum can be either global or local. Finding a global extremum is, in general, a very difficult problem. A method of solving for extrema is to perturb a local extremum by taking a finite amplitude step away from it, and then see if the routine returns you a better point. Relatively recently, so-called simulated annealing methods have demonstrated important successes on a variety of global extremization problems.

The method of simulated annealing is a technique that has attracted significant attention as suitable for optimization problems of large scale, especially one where a desired global extremum is hidden among many, poorer. Usually, the number of elements in the configuration space is factorially large so that they cannot be explored exhaustively.

At the heart of the method of simulated annealing is an analogy with thermodynamics, specifically with the way that liquids freeze and crystallize, or metals cool and anneal. At high temperatures the molecules of a liquid move freely with respect to one another. If the liquid is cooled slowly, thermal mobility is lost. The atoms are often able to line themselves up and form a pure crystal that is completely ordered over a distance up to millions of times the size of an individual atom in all directions. This crystal is the state of minimum energy for this system. The amazing fact is that, for slowly cooled systems, nature is able to find its minimum energy state. In fact, if a liquid metal is cooled quickly, it does not reach this state but rather ends up in an amorphous state having somewhat higher energy.

Many optimization methods go greedily for the quick nearby solution: From the starting point, go immediately downhill as far you can go. This leads often to a local, but not necessarily a global, minimum. Nature's own minimization algorithm is based on quite a different procedure. The so-called Boltzmann probability distribution, $Prob(E) \propto \exp(-E/kT)$ expresses the idea that a system in thermal equilibrium at temperature T has its energy probabilistically distributed among all different energy states E . Even at low temperature, there is a chance, albeit very small, of a system to get out of a local energy state. Therefore, there is a corresponding chance for the system to get out of a local energy minimum in favor of finding a better, more global one. The quantity k (Boltzmann's constant) is a constant of nature that relates temperature to energy. In other words, the system sometimes goes uphill as well as downhill; but the lower the temperature, the less likely is any significant uphill excursion.

In 1953, Metropolis et al. [MRR⁺53] first incorporated these kinds of principles into numerical calculations. the

general scheme of always taking a downhill step while sometimes taking an uphill step, has come to be known as the Metropolis algorithm. One must provide the following elements:

- A description of possible system configurations.
- A generator of random changes in the configuration.
- An objective function E whose minimization is the goal of the procedure.
- A control parameter T and an annealing schedule which tells how it is lowered from high to low values. After how many random changes in configuration is each downward step in T taken, and how large is that step.

4. Lighting Design

Given a full scene specification such as the geometry, material properties and camera parameters, the look and feel of an environment highly depends on the lighting setup. Good lighting can significantly improve the character of an environment, reducing under- or over-illuminated regions, enhancing contrast, etc. In certain spaces (workshops, storage, utility rooms) one is usually interested in bright, shadowless light. But in other areas (lobbies, boardrooms, restaurants, museums) one might be looking for a little style, drama, or have the lighting be part of an artistic expression.

Inverse lighting design starts from desired illumination objectives and constraints, such as the position of shadows, the intensity of highlights, or the overall brightness level of the environment. Using a cost function for each of the desired targets, an optimization process determines the most optimal positioning and color of light sources to satisfy the objectives and constraints of the lighting design. This inverse approach requires an *a priori* knowledge of the desired appearance.

We present a system that searches the space of lighting designs spanned by free lighting parameters: position and color of a fixed number of light sources. In the image based lighting design paper [AD04] we were able to optimize light intensities for light sources which were placed in advance around a real object. The major disadvantage of the technique was the fact we had to position our light sources and constructed basis sets of 40 light source positions. We could not constrain the number of light sources used.

By using larger basis sets and capturing the reflection field of an object, we are able to constrain the number of light sources used to obtain a certain desired illumination. This implies that a linear combination of a constrained subset of basis images has to be found in such a way that it minimizes the overall objective function.

5. Lighting Design by simulated annealing

In our image-based lighting design technique, we searched for an optimal linear combination of basis images, in such a

way that the linear combination matches the desired illumination as closely as possible. Recall that every basis image corresponds to a light position in the light-stage. The linear combination refers to the intensities that have to be applied in reality. If this is not clear, we refer the reader to [AD04].

If we want to constrain the number of light sources used, we have to apply our image based lighting design technique to a subset of the basis images and find an optimal subset of images which minimizes the objective function. The linear combination of a subset of the images corresponds to the light intensities of the fixed number of light sources used. The other lights will have intensity zero and so, only a fixed number of lights will be used.

By capturing the reflectance field of an object lit from 1280 directions, we are searching a fixed set of directions for the light sources which match a painted light distribution of an object as closely as possible. Every light direction can be formulated as a 2 dimensional vector consisting of a *phi* and a *theta*. Fixing the lights to e.g. 4, will result in an 8-dimensional dimensional problem of finding 4 *phis* and 4 *thetas* in such a way that

$$O(x(\phi, \theta)) = \sum_i w_i \left[(\Delta t \langle a_i, x(\phi, \theta) \rangle - y_i)^2 \right], \quad (1)$$

is minimal. i iterates over the pixels which have received some color information and $\langle a_i, x \rangle$ is the dot product between the i th row of reflectance field A and its lighting configuration x for the selected *phis* and *thetas*. Additionally, for each pixel, a weight w_i can be provided to modulate the importance of some pixels or areas. Also, since our basis images are taken at maximum intensity of the light sources and emitting negative light is impossible, our weights x are limited to $[0..1]$. I refer to [AD04] for a more elaborate explanation of this objective function.

We solve this non-linear problem with simulated annealing since it performs well in finding solutions in high dimensional environments. As mentioned above, the following items have to be taken into account:

System Configurations

The simulated annealing problem can be formulated as finding the minimum of $O(x)$ where x is an N -dimensional vector. Since we are optimizing our light directions, parameterized in ϕ and θ , the possible configurations for the system involve the ϕ constrained between 0 and 360 degrees, while θ may vary between 0 and 90 degrees. Only the upper hemisphere is taken into account.

A generator of random changes in the configuration

To determine random changes in our system we use a modification of a downhill simplex method. This amounts to replacing the single point x as a description of the system

state by a simplex of $N + 1$ points. The moves are controlled by reflexions, contractions or expansions of the simplex [PTVF99].

The annealing schedule.

As in all applications of simulated annealing, there can be quite a lot problem-dependent subtlety in the phrase: sufficiently slowly. Success or failure is quite often determined by the choice of annealing schedule. A strategy which works well is to reduce the temperature T fractional to the smallest function value currently represented in the simplex, and the best function ever encountered. If they are close, one should slowly reduce temperature to explore this region more. However, never reduce T by more than some fraction γ at the time.

6. results

For capturing the reflectance field, we use the setup as presented in [MPW04]. Our setup is shown in image 1. The camera and the object are mounted on a turntable. A semi-circular brace with 40 light sources is mounted over the turntable. By rotating the turntable 180 degrees in 32 steps and switching on one light source at a time, the object can be illuminated from 64×20 regularly sampled directions. Recording a reflectance field lasts for about 2 hours and must only be done once for an object.

A combination of the C# programming language and MATLAB numerical analysis software was used to create a user friendly interface to our technique. The optimization procedure takes only into account the painted pixels. Other pixels can be added if necessary. The optimization process takes about 2 minutes to complete on a standard computer (1.2 Ghz). If the result is not satisfying, user defined areas or pixels can be given more weight.

Figure 2 shows lighting configurations achieved for a metal sportscar based by designing the upper part of the car to be illuminated by green light, while the bottom part resides in red. (a) displays the target image. (b) shows a result for which 2 light sources are used. Notice the system used the shadow of the car to fill up with red light. (c) and (d) show the results, achieved by using 4, respectively 10 lights.

7. discussion and future work

Due to the image based approach, all basis images must be photographed in advance. For each possible position and direction a light source might be placed, a photograph must be made. This manual placement might be a difficult job for architectural complex environments or rooms where mounting real lights is not possible. We used a light stage to capture the reflectance field of real objects. Once a reflectance field is captured, the space of possible lightpositions and intensities can be searched to find an optimal configuration.



Figure 1: The setup for our data acquisition. The scene, a set of stones, and the camera are placed on a turntable. A semi-circular brace with 40 light sources is mounted over the turntable. By rotating the turntable 180 degrees in 32 steps and switching on one light source at a time, the object can be illuminated from 64×20 regularly sampled directions.

Another subject of future investigations might be designing a target illumination for multiple camera viewpoints. Once a viewer will walk away from the chosen camera position, the illumination will be perceived differently. Highlights might shift position on glossy surfaces, or obscured parts of the objects might come into view. Also, more complex environments, such as rooms, could be taken into account.

Our system only allows working on one photograph of the object, from a single point of view. Currently, this prohibits specifying any illumination constraints on parts of the object not visible in the photograph. Also, fine details such as sharp highlights are not within our reach because the optimization function has too narrow valleys around the highlights. So we are limited to objects with glossy materials since the valleys are somewhat broader.

Since we are using existing imaging software, we cannot give feedback about whether the painted illumination can be achieved in reality. Some feedback could be given based on the consequences of the designers actions: the designer can




Figure 2: Lighting configurations are optimized for a metal sportscar based by designing the upper part of the car to be illuminated by green light, while the bottom part resides in red. (a) displays the target image. (b) shows a result for which 2 light sources are used. Notice the system used the shadow of the car to fill up with red light. (c) and (d) show to results, achieved by using 4, respectively 10 lights.

be restricted to paint certain colors, he can be provided with a set of possible illumination solutions for pixels he did not touched yet, etc.

Computing the design with photographs provides us with a solution which is perfectly reproducible in reality without further knowledge of geometry or material attributes. This is certainly not the case with virtual objects, since rebuilding a virtual design in reality will have problems with light attachment, simulating the virtual lights with real ones, etc. In virtual environments there is a lot of freedom with does not exist in reality.

References

- [AD04] F. Anrys and P. Dutré. Image based lighting design. In *The 4th IASTED International Conference on Visualization, Imaging, and Image Processing*, 2004.
- [CSF99] A.C. Costa, A.A. Sousa, and F.N. Ferreira.

Lighting design: A goal based approach using optimisation. In *Eurographics Rendering Workshop 1999*, June 1999.

- [DHT⁺00] P. Debevec, T. Hawkins, C. Tchou, H. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, pages 145–156, July 2000.
- [DKNY95] Y. Dobashi, K. Kaneda, H. Nakatani, and H. Yamashita. A quick rendering method using basis functions for interactive lighting design. *Computer Graphics Forum*, 14(3):229–240, August 1995.
- [DSG91] J. O’B. Dorsey, F.X. Sillion, and D.P. Greenberg. Design and simulation of opera lighting and projection effects. In *Computer Graphics (Proceedings of SIGGRAPH 91)*, volume 25,

- pages 41–50, July 1991.
- [DWT⁺02] P. Debevec, A. Wenger, C. Tchou, A. Gardner, J. Waese, and T. Hawkins. A lighting reproduction approach to live-action compositing. *ACM Transactions on Graphics*, 21(3):547–556, July 2002.
- [GH00] R. Gershbein and P.M. Hanrahan. A fast relighting engine for interactive cinematic lighting design. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, pages 353–358, July 2000.
- [KPC93] J.K. Kawai, J.S. Painter, and M.F. Cohen. Radiotimization - goal based rendering. In *Proceedings of SIGGRAPH 93*, Computer Graphics Proceedings, Annual Conference Series, pages 147–154, August 1993.
- [MAB⁺97] J. Marks, B. Andalman, P. A. Beardsley, W. Freeman, S. Gibson, Jessica K. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. Design galleries: A general approach to setting parameters for computer graphics and animation. In *Proceedings of SIGGRAPH 97*, Computer Graphics Proceedings, Annual Conference Series, pages 389–400, August 1997.
- [MPW04] V. Masselus, P. Peers, and Y. Willems. Smooth reconstruction and compact representation of reflectance functions for image-based relighting. In *Eurographics Symposium on Rendering 2004*, pages 21–23, June 2004.
- [MRR⁺53] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Metropolis. *Journal of Chemical Physics*, 21:1087–1092, 1953.
- [NSD94] J.S. Nimeroff, E. Simoncelli, and J. Dorsey. Efficient re-rendering of naturally illuminated environments. In *Fifth Eurographics Workshop on Rendering*, pages 359–373, June 1994.
- [PF92] P. Poulin and A. Fournier. Lights from highlights and shadows. In *1992 Symposium on Interactive 3D Graphics*, volume 25, pages 31–38, March 1992.
- [PRJ97] P. Poulin, K. Ratib, and M. Jacques. Sketching shadows and highlights to position lights. In *Computer Graphics International 1997*, June 1997.
- [PTG02] F. Pellacini, P. Tole, and D.P. Greenberg. A user interface for interactive cinematic shadow design. *ACM Transactions on Graphics*, 21(3):563–566, July 2002.
- [PTVF99] R. Press, A. Teukolsky, T. Vetterling, and P. Flannery. *Numerical Recipes in C*. Cambridge University Press, 1999.
- [SDS⁺93] C. Schoeneman, J. Dorsey, B. Smits, J. Arvo, and D. Greenberg. Painting with light. In *Proceedings of SIGGRAPH 93*, Computer Graphics Proceedings, Annual Conference Series, pages 143–146, August 1993.
- [SL01] R. Shacked and D. Lischinski. Automatic lighting design using a perceptual quality metric. *Computer Graphics Forum*, 20(3):215–226, 2001.