Radioptimization — Goal Based Rendering

John K. Kawai
James S. Painter
Michael F. Cohen

UUCS-93-002

Department of Computer Science
University of Utah
Salt Lake City, UT 84112 USA

February 24, 1993

Abstract

This paper presents a method for *designing* the illumination in an environment using optimization techniques applied to a radiosity based image synthesis system. An optimization of lighting parameters is performed based on user specified constraints and objectives for the illumination of the environment. The system solves for the "best" possible settings for: light source emissivities, element reflectivities, and spot light directionality parameters so that the design goals, such as to minimize energy or to give the the room an impression of privacy, are met. The system absorbs much of the burden for searching the design space allowing the user to focus on the goals of the illumination design rather than the intricate details of a complete lighting specification. A software implementation is described and some results of using the system are reported.

The system employs an object space perceptual model based on work by Tumblin and Rushmeier to account for psychophysical effects such as subjective brightness and the visual adaptation level of a viewer. This provides a higher fidelity when comparing the illumination in a computer simulated environment against what would be viewed in the "real" world. Optimization criteria are based on subjective impressions of illumination with qualities such as "pleasantness", and "privateness". The qualities were selected based on Flynn's work in illuminating engineering. These criteria were applied to the radiosity context through an experiment conducted with subjects viewing rendered images, and the respondents evaluated with a Multi-Dimensional Scaling analysis.

Radioptimization — Goal Based Rendering

John K. Kawai, James S. Painter
Department of Computer Science
University of Utah

Michael F. Cohen
Department of Computer Science
Princeton University

February 24, 1993

Abstract

This paper presents a method for *designing* the illumination in an environment using optimization techniques applied to a radiosity based image synthesis system. An optimization of lighting parameters is performed based on user specified constraints and objectives for the illumination of the environment. The system solves for the "best" possible settings for: light source emissivities, element reflectivities, and spot light directionality parameters so that the design goals, such as to minimize energy or to give the the room an impression of privacy, are met. The system absorbs much of the burden for searching the design space allowing the user to focus on the goals of the illumination design rather than the intricate details of a complete lighting specification. A software implementation is described and some results of using the system are reported.

The system employs an object space perceptual model based on work by Tumblin and Rushmeier to account for psychophysical effects such as subjective brightness and the visual adaptation level of a viewer. This provides a higher fidelity when comparing the illumination in a computer simulated environment against what would be viewed in the "real" world. Optimization criteria are based on subjective impressions of illumination with qualities such as "pleasantness", and "privateness". The qualities were selected based on Flynn's work in illuminating engineering. These criteria were applied to the radiosity context through an experiment conducted with subjects viewing rendered images, and the respondents evaluated with a Multi-Dimensional Scaling analysis.

1 Introduction

Historically, lighting design has been a black art. The lighting designer first received a design specification of the customer's expectations and of the room's function. The designer then made a lighting lay out and from experience would sketch what the room would look like from rough lighting calculations. With the advent of computer aided rendering, this process has been simplified allowing the designer to model lighting specifications with a CAD system and have it simulate the lighting calculations giving the designer a quick design check of what the room would look like. This also allows the customer who has no experience with lighting units a realistic preview of the finished room early in the design cycle [26]. Progress in rendering to date has mainly focused on improving the realism of the physical simulation and the development of algorithms with faster performance. Though great advances have been made in these areas, little work has been done on addressing the design problems in creating better quality lighting.

Lighting designers base their art on the belief that spatial lighting patterns are a visual communicative medium, in which some patterns of light suggest or reinforce shared attitudes and impressions to people of the same cultural background [10]. In addition, the designer must be aware of the need to conserve the electrical energy used in implementing their designs. An over-reaction to the wasteful energy consumption of the 1960s and 1970s often led to buildings which were inadequately lit for their designed purposes, hampering the productivity of the residents. A better balance of goals between energy conservation and the quality of the lighting is needed [21]. With office and factory personnel costs ranging from \$150 to \$275 per square foot [20], an extra investment of \$1 or \$2 per square foot per year can potentially result in a large savings through improved productivity.

This paper proposes a *Teleological* [1] or goal based illumination design approach to help a lighting designer search the space of of possible lighting specifications. Though computers will never replace artists, the system may generate configurations not previous considered or optimize on an already considered configuration.

The approach described below allows the designer to concentrate on high level goals such as "visual clarity" and specify constraints such as minimum lighting levels in specific locations. The system then determines optimal settings for the lighting parameters of the modeled environment by searching for the "best" possible settings for

- light source emissivities,
- surface reflectivities, and
- spot light directionality.

Unconstrained optimization techniques are employed in conjunction with classical radiosity [14, 4, 3, 16] to simulate global illumination and our current implementation is thus limited to diffuse environments with fixed geometry. We have however, extended the basic radiosity system to include spot lights as well as diffuse area sources.

Creating an appropriate two-way link between the designer and the rendering system requires two important enhancements to basic rendering methods. First, since the designer is asked to iteratively evaluate the visual impression from a rendered image, the images must provide (as much as possible) a subjective match to a "real" environment. We have applied the work of Tumblin and Rushmeier [28] on the psycho-physical quantities of subjective brightness as related to the adaptation of the viewer in order to map luminance values to brightness values for display on the CRT. This is an important and often overlooked step in providing an image with good subjective fidelity to the real environments.

Finally, the optimization objectives presented to the designer are based on John Flynn's work [10] from the architectural lighting community. His experiments allow one to quantify subjective impressions of lighting patterns. In our work, we have conducted experiments with subjects viewing computer generated images to create a mapping from Flynn's criteria to quantifiable qualities in the radiosity simulations.

After a brief outline of the underlying technology supporting our work, we will describe a software system which implements the ideas discussed and report initial results from using the system.

2 Previous Work

There are three bodies of technology and related literature that are central to the work reported here: numerical optimization, radiosity based image synthesis, and knowledge about human perception as it relates to subjective impressions of lighting and to subjective impressions from images presented on a CRT. We will briefly review each of these area concentrating on the pertinent subtopics in each that relate directly to our work.

2.1 Optimization

The basic constrained optimization problem is to minimize (or maximize) the scalar quantity of an objective function of n system parameters $\mathbf{X} = (x_1, x_2, ..., x_n)$ while satisfying a set of constraints. The j^{th} constraint, C_j, may be posed as an equality or inequality:

$$C_j(\mathbf{X}) - K_j = 0 \quad \text{or} \quad C_j(\mathbf{X}) - K_j < 0 \tag{1}$$

Thus the full problem with m constraints can be stated as:

$$\begin{array}{ll} \text{minimize} & f(x_1, x_2, x_3, ..., x_n) \\ & \\ \text{subject to} & C_1(K_1, x_1, x_2, x_3, ..., x_n) = 0 \\ & C_2(K_2, \mathbf{X}) = 0 \\ & C_3(K_3, \mathbf{X}) = 0 \\ & ... \\ & C_7(K_7, \mathbf{X}) < 0 \\ & C_8(K_8, \mathbf{X}) < 0 \\ & ... \\ & C_m(K_m, \mathbf{X}) < 0 \end{array}$$

Constrained optimization problems arise in a wide variety of domains. One might want to find the optimal way to allocate a limited supply of resources to feed the most people or to maximize a return on investment of a limited amount of money. In the problem addressed in this work, a designer may want to minimize a cost associated with a choice of lighting while maintaining particular design criteria. For example, minimizing electricity consumption subject to the lighting level on desk tops remaining above a minimum value.

Unfortunately, there is no computational algorithm for optimization which will always find a global constrained minimum when the objective and the constraints are allowed to be general non-linear functions. Research on optimization techniques, has resulted in a number of very useful texts under a number of headings, such as Mathematical Programming [6, 17], Operations Research [7], Optimal Control [19], and Optimization [22, 8, 13]. The important aspects of a particular optimization problem, leading to a choice of algorithm include:

- the nature of the objective function, e.g., linear vs. non-linear, convex vs. non-convex, differentiability,
- the nature of the constraints, e.g., linear vs. non-linear, equality vs. inequality, differentiability,
- abilities and needs, e.g., the ability to specify good starting guesses, to generate analytical derivatives, and the need for a global vs. local optimum, and
- the nature of the variables, e.g., continuous vs. discrete, and scalar vs. vector valued.

2.1.1 Constrained Optimization

The lighting optimization problem introduces non-linearities through the objectives, constraints, and the implicit constraints of the radiosity relationships themselves. Most methods for dealing with constraints involve a transformation from a constrained problem to an (approximately) equivalent unconstrained optimization, with the solution to the unconstrained problem found with one or more of the methods described in the following section.

Transformations from a constrained to unconstrained problem involve either removing a constraint by explicitly solving for one optimization variable, or by adding a new function into the objective. If the constraints are simple, and variables can be solved for as explicit expressions of other variables, then the first alternative is attractive, as variables can be directly removed from the optimization problem[24, 25]. Other techniques introduce new variables as in the Lagrange Multiplier methods [22] in which the new unconstrained problem is taken to be the sum of the objective and a linear combination of the constraints. Perhaps the simplest method involves converting the constraints to *penalty* functions, i.e., add a function of the constraint violations into the objective, and minimize the new problem.

The methodology employed in this work involves a combination of explicit replacement of variables, in this case the radiosities for their equivalent expressions, and penalty methods based on quadratic and quartic functions of the constraint violations. Note that the penalty method can not guarantee that the constraints will be exactly satisfied since the method leads to a solution in which a balance is reached between satisfying the constraints and minimizing the objective. This can, however, often be viewed as an advantage in a design setting.

2.1.2 Unconstrained Optimization

Once the constraints are removed or transformed, the problem reduces to finding a minimum of the objective. If the objective is continuous and differentiable, and has a bounded minimum, the minimum point will always be characterized by having a zero gradient. Zero gradients may, however, also occur at local minima and at saddle points on the objective hypersurface. Most optimization methods are performed iteratively from a starting point, $\mathbf{X}(0)$ in the multidimensional search space. Local information about the value, gradient, and Hessian (matrix of second order partial derivatives) of the function is gathered and a step direction is selected to move the solution to a new guess. A local minimum is considered to be found if the gradient is zero and the local region is convex, i.e., has a positive definite Hessian. Techniques for selecting a step direction vary, from simple gradient descent (a step in the negative gradient direction), to conjugate gradient methods (steps in a series of conjugate directions), to Newton's Method which solves for a step direction as the inverse of the Hessian times the negative gradient.

$$\nabla^2 f \cdot \Delta\mathbf{X} = -\nabla f$$

Where ∇^2 is the Hessian operator, ∇ is the gradient and f is the objective function, and $\Delta\mathbf{X}$ defines a multi-dimensional search direction, (or vector to the minimum in a quadratic problem).

By inverting the formula we can find the search direction (ΔX) in terms of the inverse Hessian and the current gradient.

$$\Delta\mathbf{X} = -(\nabla^2 f)^{-1} \cdot \nabla f$$

Although Newton's method can have great success, difficulties occur in regions where the Hessian is not positive definite, and because the Hessian itself may be impossible to derive analytically and difficult to compute numerically. A number of Quasi-Newton methods have been used extensively in place of a direct application of Newton's Method. These include the Davidon-Fletcher-Powell (DFP) and Broyden-Fletcher-Goldfarb-Shanno (BFGS) methods. These methods iteratively estimate the inverse Hessian from a series of gradients [23, 22]. In general, they begin with an identity matrix as the Hessian, thus defaulting to a simple gradient descent for the first step and then modifying the inverse Hessian on succeeding iterations. A modification of the BFGS algorithm is used in the work presented here.

Due to non-linearities, the BFGS method does a series of one dimensional line searches until it converges on a local minimum. The inverse of the *Hessian* matrix of second partials is approximated from differences of the gradient and is used, along with the gradient, to select the direction of each line search. The inverse Hessian is updated at each step through the BFGS iteration step. Convergence is achieved when either the gradient vanishes or when two consecutive line searches converge to the same solution.

A final issue which should be addressed in unconstrained optimization methods, is the length of the step in the direction selected. This problem is, in essence, a reduced minimization problem, in which the search space is limited to a line in the hyperspace. The method adopted in this work first brackets the minimum and then estimates the minimum within the local region by a quadratic approximation.

2.2 Radiosity

For an automated illumination design system to be useful, it must be based on a realistic and physically based model of global illumination that takes into account the inter-reflections of light within an environment. For example, indirect light sources illuminate much of the environment only after being reflected off a wall or ceiling. It is important that these effects are captured in the illumination design system. The global illumination problem is still computationally intractable for the fully general case. *Radiosity* methods [4] have developed into an efficient and practical method for the restricted case of diffuse environments.

Radiosity algorithms discretize the environment into a set of elements with an assumed functional form for the radiosity across the surface. The simplest and most common functional form is a constant value, called the radiosity. A balance of energy between elements must exist which imposes a set of interdependent constraints on the element radiosities in the environment:

$$B_i = E_i + \rho_i \sum_{j}^{n} F_{i,j} B_j \tag{2}$$

where B_i is the radiosity of element i,
E_i is the emission of element i,
ρ_i is the reflectivity of element i,
$F_{i,j}$ is the form factor from element i to element j

The form factor is the fraction of light leaving one element (i) that arrives at another (j) and is given by:

$$F_{i,j} = \frac{1}{A_i} \int_{p_i \in A_i} \int_{p_j \in A_j} \delta(p_i, p_j) \frac{cos(\phi_i) cos(\phi_j)}{\pi r_{ij}^2} dA_i dA_j$$

where A_i and A_j are the element surfaces,
p_i and p_j are points on elements i and j respectively,
$\delta(p_i, p_j)$ returns 1 if p_i and p_j are mutually visible and 0 otherwise,
ϕ_i is the angle between the normal vector at p_i and the vector from p_i to p_j,
ϕ_j is the angle between the normal vector at p_j and the vector from p_j to p_i,
r_{ij} is the distance from p_i to p_j.

When all such constraints are considered, a linear system of equations result:

$$\begin{bmatrix} 1 & -\rho F_{1,2} & -\rho F_{1,3} & . & . & F_{1,N} \\ -\rho F_{2,1} & 1 & -\rho F_{2,3} & . & . & F_{2,N} \\ . & & & & & . \\ . & & & & & . \\ -\rho F_{N-1,1} & . & . & . & 1 & F_{N-1,N} \\ -\rho F_{N,1} & . & . & . & F_{N,N-1} & 1 \end{bmatrix} \begin{bmatrix} B_1 \\ B_2 \\ . \\ . \\ B_{N-1} \\ B_N \end{bmatrix} = \begin{bmatrix} E_1 \\ E_2 \\ . \\ . \\ E_{N-1} \\ E_N \end{bmatrix} \tag{3}$$

2.2.1 Gathering, Shooting, and Hierarchy

This system can be solved either by Gauss-Seidel iteration, "gathering" light into each element [4], or by a "shooting" method that distributes the light from the brightness element to the other elements [3]. The solution to this system yields the element radiosities, B_i, for every element in the scene. A final image from any viewpoint can be constructed quickly by projection onto the view plane. The constant elements are most often blended by interpolation of radiosity values to the element vertices and then using Gouraud shading for display. Additional images can be formed from any view point without additional radiosity computations.

A direct solution to the radiosity equation appears to require at least n^2 space and time, given n elements. Early radiosity methods [5] used *substructuring* techniques, decomposing the environment polygons into two levels of hierarchy to alleviate the problems of n^2 time and space. The shooting based progressive radiosity method [3] avoids the space overhead by computing form factors on the fly and never explicitly storing the form factor matrix. More recently, Hanrahan *et al.* have shown that, in fact, the form factor matrix can be stored in $O(n)$ space by exploiting a *block* structure of the matrix yielding a very fast hierarchical radiosity algorithm [16]. Surfaces are hierarchically decomposed into smaller and smaller elements with the entries of the form factor matrix stored as links between nodes in the hierarchical subdivision. The space and time savings results from the fact that most interactions can be represented within a desired tolerance via a single link higher up in the hierarchy rather than many links at the lowest element level.

2.2.2 Spot Lights

Directional lighting effects such as spot lights can be added to the radiosity equation by replacing the $cos(\phi_i)$ term in the form factor equation with a different distribution function:

$$F_{i,j} = \frac{1}{A_i} \int_{p_i \in A_i} \int_{p_j \in A_j} \delta(p_i, p_j) s(\phi_i) \frac{cos(\phi_j)}{\pi r_{ij}^2} dA_i dA_j$$

where $s(\phi_i)$ is the directionality distribution weight for the light source as a function of the angle between the direction vector of the light (element i) and the vector between the points p_i and p_j. We have assumed that the reflectivity of the source, element i, is zero and hence its radiosity is completely dominated by the emittance term. The light source distribution function s could be any radially symmetric function of angle and is commonly represented in the lighting industry by goniometric diagrams plotting energy distribution as a function of angular direction.

Here we restrict ourselves to distributions of the form

$$s_n(\phi) = w(n) \, cos^n(\phi) \tag{4}$$

for values $n >= 1$. This yields a continuously variable range of beam widths. It is useful to be able to change the beam width without affecting the total energy emitted by the light. This requires a normalization factor, $w(n)$, in the emission function s_n. The normalization factor $w(n)$ must be chosen so that the total energy emitted over the hemisphere is constant, independent of n, as the beam width is adjusted. The value of the constant is chosen so that $w(1) = 1$. That is,

$$\int_{\text{hemisphere}} s_n d\omega \;=\; \pi \tag{5}$$

where $d\omega$ is the differential solid angle on the sphere. Carrying out the integration in spherical coordinates yields the normalization weight,

$$w(n) \;=\; \frac{n+1}{2} \tag{6}$$

2.3 Human Perception

2.3.1 Radiance, Luminance, and Brightness

Radiosity methods solve the physics of the global transport of light in terms of the radiometric units of *radiance* (energy per unit time per unit solid angle). Human visual systems are not sensitive to all wavelengths of light and are not equally sensitive across the visible spectrum. A weighting of radiance values by the human *luminous efficiency* curve results in photometric units of *luminance*. The human visual system detects contrasts rather than absolute luminance values. *Brightness* is a measure of the subjective sensation produced by visible light. Brightness, measured in units of brils, relates linearly to human visual response. For example, if two light sources are compared and one appears to be twice as bright as the other, the brightness of the first, in brils, will be twice that of the second.

The human eye is sensitive to a luminance range of approximately ten orders of magnitude. However, at any one time the eye can only detect a brightness range of 100 to 1 with good accuracy. The eye adjusts the iris to open and close, limiting the amount of light entering the eye to seek a state of equilibrium that is appropriate for general brightness conditions. Tumblin and Rushmeier [28] studied work by Stevens [27] who theorized that the adaptation level of a scene can be estimated by as the expected value (mean) of the log_{10} of the luminances visible on the retina:

$$\text{EXP}_{p\in\text{retina}}\{log_{10}(L(p))\} \tag{7}$$

where $L(p)$ is the luminance at a point p on the retina. Miller *et al.* also theorized that differing adaptation of the eye results in a family of curves relating luminance, L, and brightness, P, values in the form

$$log_{10}(P) \;=\; aa * log_{10}(L) + bb \tag{8}$$

where P is the brightness values specified in *brils*
L is the luminance values specified in *nits*
$aa = 0.4 * log_{10}(L_w) + 2.92$
$bb = -0.4 * (log_{10}(L_w))^2 + (-2.584 * log_{10}(L_w)) + 2.0208$
L_w is the white adapting luminance which can be approximated by the equation $log_{10}(L_w) = \text{EXP}\{log_{10}(L_i)\} + 0.84$

This perceptual model accepts luminance values in units of *nits* which in photometric units are related to *lux* by, 1 lux = 1 nit / $10,000$. Thus solving for brils in terms of an element radiosity of B lux is

$$P = 10^{aa*log_{10}(B/10,000)+bb}$$

Since the adaptation of the eye is affected only by what is visible to the retina, perceptual processing is usually done in screen space making the whole process view dependent. This assumes that the viewer adapts to a single view rather than to an entire environment, i.e. the viewer remains transfixed on a single view of an environment long enough to adapt to the lighting level. In practice, we are constantly moving our head and eyes to scan a room and and hence adapt to the overall room lighting rather than to a single view. In our work we proposes a view independent approach to lighting design, therefore, the conversion from luminance units into perceptual units is performed in object space, at the element level. Each element is considered to contribute to the adaptation proportional to its physical size. This neglects the view dependent effects of perspective foreshortening and occlusion but has the advantage that it yields view independent results. We have found that the object space, view independent method gives results that are nearly identically to view dependent screen space methods for typical, single room, architectural models. Clearly if the environment being modeled consists of many separate rooms, only a single local region should be considered. The work of Funkhouser, Sequin, and Teller [12] would be of great value here. In addition to the view independence, calculating perception in object space has the added advantage of faster performance if the number of elements is much smaller than the number of screen pixels.

2.3.2 Subjective Impressions of Illumination

In the 1970's, John Flynn published a series of articles [11, 9, 10], introducing a methodology with which to quantify parameters that elicit a shared human behavioral response and subjective impression. In particular, Flynn examined how non-uniform, peripheral, and bright lighting affects impressions of visual clarity, spaciousness, relaxation, and privacy. Flynn created six different light settings for a conference room and subjectively associated each room with a non-uniform, peripheral, and brightness value so that each room corresponded to a point in a 3 dimensional space of the different lighting characteristics. Flynn also associated a set of semantic differential (SD) rating scales such as large-small and spacious-cramped with each category of impression. Test subjects were then asked to make pair wise comparisons of the differences between each room from the set of SD rating scales where 0 meant no difference and 10 meant a large difference.

The data gathered resulted in a 6x6 symmetric dissimilarity matrix comparing the 6 rooms for each subject tested and each SD comparison made, *e.g.* large-small. A multidimensional scaling program INDSCAL [2, 15], was used to determine how each subjected weighted the non-uniformity, peripheral and brightness values in making each SD comparison. A weighting of each dimension for each subject was determined that best fit the data. The results showed a correlation between the room positions hypothesized by Flynn and the positions computed by INDSCAL, supporting Flynn's hypothesis that brightness, non-uniformity, and peripheral lighting reinforce particular impressions. In addition, there also was a correlation for the weights for each parameter among all the subjects, supporting the concept that particular lighting patterns elicit a shared impression. By this process, Flynn was not only able to demonstrate that there is a definite correlation between the measurable quantities (non-uniform, peripheral, and bright) and the subjective impressions (visual clarity, spaciousness, and relaxation), but was able to quantify how much each of the measurable dimensions affects each subjective impression.

As described shortly, we have adapted this work through an additional level of experimentation in which subjects reported impressions from computer generated images.

3 Problem Formulation

To pose the illumination design task as a constrained optimization problem we must identify:

- the variables involved in the optimization process
- the constraints that must be satisfied,
- and the objective function.

3.1 Optimization Variables

In a normal radiosity based renderer, the element radiosities B_i are the unknowns to be computed in terms of fixed material and light property parameters. In the optimization setting the material and light properties are no longer fixed and must also be considered as variables. Constraints may be imposed on any of these variables and the objective function may involve any or all of them.

In the illumination design problem the optimization variables are light source specification parameters (emissions, spot light directions, spot light focus), element radiosities, B_i, and element reflectivities, ρ_i. We consider two types of light sources. Ordinary diffuse light sources are diffusely emitting elements and are described by a single emissivity parameter E_i. Directional lights are idealized spot lights which are described by a position, direction, and distribution pattern, (cos^n distributions in our system). We assume that the light source position is fixed and only the direction and distribution pattern is allowed to change during optimization.

Every light source emission, E_i, light direction vectors $\mathbf{V}_i$, and cosine distribution exponent n_i, element radiosity B_i and reflectivity ρ_i have the potential to be a variable in the optimization problem. If all are treated explicitly as domain variables in the optimization an intractably large system will result. The B_i's can be eliminated by direct substitution of the radiosity equation, and typically only a small number of the elements will have variable emission, reflectivity or directionality parameters. These remaining variables are called the "free" variables of the optimization problem.

3.2 Constraints

Constraints fall into three categories.

Physical constraints specify the relationships between light emission and element radiosities that are dictated by the physics of light transport. The constraints are captured in the rendering equation [18]. We assume perfect diffuse surfaces and a discretized environment yielding the radiosity approximation given in equation 2.

Design goals are constraints provided by the user. These may be either equality or inequality constraints and may apply to a single element, or a conglomeration of elements. For example that a particular element's radiosity is a given constant, $B_i = K$ for some constant K is an equality constraint on a single element that expresses a fixed radiosity for the element. Inequality constraints such as $K_{low} \leq B_i \leq K_{high}$ can also be specified (in essence two inequality constraints) requiring the radiosity of element i to stay within the bounds K_{low} and K_{high}. Further, a group of elements, not necessarily from the same patch, may have constraints applied to the maximum, or minimum radiosity of the group.

Barrier constraints are hard bounds on the allowable ranges of the optimization variables that must be satisfied to insure that the model is physically realizable. For example, light emissions must remain positive and element reflectivities must remain in the range $0 <= \rho_i <= 1$. Barrier constraints are conceptually similar to inequality design goals. The main difference is that a barrier constraint *must* be satisfied in order to produce a valid model. Design goals are *desires* that may not need not be satisfied exactly.

3.3 Objective Function

In general, radioptimization problems are under constrained. There may be an infinite number of possible solutions that satisfy the problem constraints. The *objective* function is used to select between the many possible solutions.

Some examples of objective functions are:

- Minimize total energy to save money.
- Desired specific measurable lighting patterns, for example brightness, uniformity, and peripheral vs. central lighting,
- Desired subjective impression of the illumination such as clarity, pleasantness, and privateness.

The simplest, directly measurable objective is the energy,

$$f_{energy} = \sum_i B_i A_i \tag{9}$$

A variation of Flynn's work, described in the previous section was used to develop a way of quantifying subjective impressions. Flynn's experiment was duplicated except, instead of having the subjects judge actual rooms with different lighting characteristics, they were shown rendered images of an identical room with different light patterns (see color plate 1). Once the data was collected, it was processed by INDSCAL with the brightness, non-uniform, and peripheral values for each room computed by the following functions.

$$
\begin{aligned}
f_{brightness}(P,A) &= -\frac{\sum_{i\in all-elements} P_i A_i}{\sum_{i\in all-elements} A_i} \\
f_{non-uniform}(P,A) &= -\left(\frac{\sum_{i\in walls(P_{avg,i}-P_i)^2 A_i}}{\sum_{i\in walls} A_i)}\right)^{\frac{1}{2}} \\
f_{peripheral}(P,A) &= \frac{\sum_{i\in horizontal-plane} P_i A_i}{\sum_{i\in horizontal-plane} A_i} - \frac{\sum_{i\in walls} P_i A_i}{\sum_{i\in walls} A_i}
\end{aligned}
$$

where P_i is the brightness of element i
A_i is the area of element i
$P_{avg,i}$ is the average brightness of the elements around element i

The functions are defined in terms of perceptual values because humans subjectively quantify illumination by brightness not by actual luminance.

The results from INDSCAL showed that there was a correlation among the subjective impressions of visual clarity, privacy, and pleasantness. Taking a linear combination of the average weight of the subjects with similar responses imply that the clarity, privacy, and pleasantness can be roughly evaluated by the following linear combinations of brightness, non-uniform, and peripheral lighting values.

$$
\begin{aligned}
f_{clear}(P,A) &= 0.89963 f_{brightness}(P,A) - 0.38098 f_{non-uniform}(P,A) - 0.58060 f_{peripheral}(P,A) \\
f_{pleasant}(P,A) &= 0.78437 f_{brightness}(P,A) - 0.52679 f_{non-uniform}(P,A) + 0.23984 f_{peripheral}(P,A) \\
f_{private}(P,A) &= 0.89064 f_{brightness}(P,A) + 0.31562 f_{non-uniform}(P,A) - 0.08648 f_{peripheral}(P,A)
\end{aligned}
$$

In theory, any user specified function of the optimization variables could be used as an objective function. An alternative is to provide a fixed library of objective functions and allow the user to construct an objective function via linear combinations of the library functions. Each individual objective function in the library has a well defined and intuitive behavior. The user can then control the weights of the individual objectives to determine the final objective function. This allows the user control without an undo amount of complexity.

The objective functions with which we have experimented is thus the weighted sum

$$
\begin{aligned}
f(B,E,P) = \quad & W_{energy} & f_{energy} & + \\
& W_{brightness} & f_{brightness} & + \\
& W_{non-uniform} & f_{non-uniform} & + \\
& W_{peripheral} & f_{peripheral} & + \\
& W_{clear} & f_{clear} & + \\
& W_{pleasant} & f_{pleasant} & + \\
& W_{private} & f_{private} & +
\end{aligned}
\qquad (10)
$$

3.4 Conversion of the Constrained Problem to an Unconstrained Problem

The design goal constraints can be included in the objective function through the *penalty method* [22] by penalizing deviations from constraints through explicit terms in the objective function. The penalty imposed on the objective is defined as the square of the constraint violation. For example, if the j^{th} constraint, C_j, is an equality constraint specifying a particular radiosity[1] to be a given constant, $(B_{i_j} = K_j)$, this will result in a penalty term f_{C_j} in the cost function given by $f_{C_j} = A_{i_j}(K_j - B_{i_j})^2$. Inequality constraints can be handled through a penalty function that "turns on" when the constraint is not satisfied. For example, the inequality constraint C_j given by $(B_{i_j} < K_j)$ results in a penalty term $f_{C_j} = A_{i_j}(K_j - B_{i_j})^2$ when B_{i_j} is greater than K_j and is zero otherwise.

Barrier Constraints are handled in a similar fashion to impose hard physical restrictions on certain values, for example, the emission variables must always remain positive. Similarly, reflectivities must remain between 0 and 1. A barrier term is added to the objective function for each barrier constraint to avoid violations of these constraints. The barrier constraint G_j given by $(X_j > K_j)$ for some free variable X_j results in a barrier term $f_{G_j} = (X_j - K_j)^{-4}$ for $X_j > K_j$. In addition, the optimization search explicitly enforces the constraint $(X_j > K_j)$ by clamping the X_j to $K_j + \epsilon$ when X_j drops below K_j, where ϵ is a small positive constant. This will yield a large barrier term in the objective function tending to lead the search away from the barrier in the next iteration.

The remaining constraints are the "physical constraints" specified by the radiosity equation (equation 2). These are dealt with by direct substitution. The radiosity equation implicitly defines each B_i in terms of all the $E, \mathbf{V}$, n and ρ's. The B_is are calculated via a radiosity solution algorithm [16]. The values for the P_i's can then be computed directly from the B_i's by equation 2.3.1. The B_i and P_i values can be directly substituted into the objective function. This effectively eliminates all the B_i's and P_i's from the set of optimization domain variables.

Thus the modified optimization problem is given by:

$$
\begin{aligned}
f(\mathbf{X}) = \quad & W_{energy} && f_{energy} && + \\
& W_{brightness} && f_{brightness} && + \\
& W_{non-uniform} && f_{non-uniform} && + \\
& W_{peripheral} && f_{peripheral} && + \\
& W_{clear} && f_{clear} && + \\
& W_{pleasant} && f_{pleasant} && + \\
& W_{private} && f_{private} && + \\
& W_{designgoals} && \textstyle\sum_j f_{C_j} && + \\
& && \textstyle\sum_j f_{G_j} &&
\end{aligned}
\qquad (11)
$$

where $\mathbf{X}$ is a point in the multidimensional space spanned by the remaining free variables, E_i, $\mathbf{V}_i$, n_i, and ρ_i.

Through the use of the penalty method, barrier functions, and substitution of physics constraints, the optimization problem can now be stated as a simple unconstrained, multidimensional minimization problem. Let $\mathbf{X}$ be a multidimensional vector in the "design space", the space spanned

[1] B_{i_j} indicates the radiosity of the i^{th} element, where i was selected by the j^{th} constraint.

by the free variables in the design. We must identify a point in the design space, $\mathbf{X}^\star$ such that the objective function $f(\mathbf{X}^\star)$ is (at least locally) minimized. There are many solution methods for such a minimization problem. We use the well known BFGS method described above [23].

4 Implementation

4.1 Overview

The ideas discussed above have been implemented on SGI IRIS and IBM RS6000 workstations. The radiosity analysis portion of the work is based on Hanrahan *et al.*'s hierarchical radiosity code. The user provides an initial model which is rendered as is to provide a baseline rendering. The user can select elements interactively from an image generated from the baseline solution to specify the free variables in the optimization process. These may be light source emissions, element reflectivities and spot light directionality parameters. The user can also specify the objective function weights $W_{energy}, W_{brightness}, W_{non-uniform}$, etc. to direct the optimization process. After all the design goals and objective weights are specified, the optimization process is run until convergence is achieved.

This process can be described in Pseudo code by:

Compute baseline rendering.
Establish constraints and objectives.
REPEAT
    Evaluate partial derivatives.
    Compute search direction $\Delta\mathbf{X}$ *using BFGS.*
    Perform line search in the direction $\Delta\mathbf{X}$*.*
    Display results, and allow user to modify constraints and objectives
UNTIL convergence.

4.2 Baseline rendering

The initial model is rendered with the hierarchical radiosity algorithm. During baseline rendering, the input model is subdivided into a hierarchical structure and links are established between nodes in the hierarchy to establish the block structured form factor matrix as described in [16]. Element radiosities are computed and an image is displayed, with interactive user control over the viewpoint and view direction.

4.3 Establishing Constraints and Objectives

Once an image is displayed the user can select elements directly from the screen with the mouse and set constraints via the user interface shown in color plate 2. In this example, the desk top has been selected as indicated by the green outline. Current illumination information for the selected element is displayed in the lower right corner of the interface. Through a set of buttons in the interface, the user can elect to impose a constraint on the element radiosity, and/or specify that the element reflectivity or emission should be a free variable in the optimization process. Spot lights are handled with a similar interface that allows the light direction vector and/or distribution parameter n to be marked as free variables in the optimization. The objective function weights can also be adjusted with slider bars in this interface.

4.4 Partial Derivative Estimation

Evaluation of partial derivatives of the modified objective with respect to the free variables is required by the optimization process. Rapid derivative evaluation is critical to an efficient solution to the optimization problem. For each free optimization variable we must be able to evaluate the partial derivative of the modified objective function relative to the free variable. For example, to compute the partial derivative of the objective function with respect to a light emission, E_k, we must evaluate:

$$\begin{array}{rlll} \partial f/\partial E_k = & W_{energy} & \sum_i \partial E_i/\partial E_k \; A_i & + \\ & W_{brightness} & \partial f_{brightness}/\partial E_k & + \\ & W_{non-uniform} & \partial f_{non-uniform}/\partial E_k & + \\ & W_{peripheral} & \partial f_{peripheral}/\partial E_k & + \\ & W_{clear} & \partial f_{clear}/\partial E_k & + \\ & W_{pleasant} & \partial f_{pleasant}/\partial E_k & + \\ & W_{private} & \partial f_{private}/\partial E_k & + \\ & W_{design} & \partial \sum_j f_{C_j}/\partial E_k & + \\ & & \partial \sum_j f_{G_j}/\partial E_k & \end{array} \quad (12)$$

The partial of the constraint function f_{C_j} for an equality constraint $C_j : (B_{i_j} = K_j)$ is:

$$\partial f_{C_j}/\partial E_k = -2A_{i_j} \cdot (K_j - B_{i_j})\partial B_{i_j}/\partial E_k \quad (13)$$

For an inequality constraint, the partial $\partial f_{C_j}/\partial E_k$ is zero when the constraint is satisfied and is given by equation 13 otherwise.

The partial of a barrier functions f_{G_j} can also be expressed directly as:

$$\partial f_{G_j}/\partial E_k = -4(E_j - G_j)^{-5}\partial E_i/\partial E_k \quad (14)$$

The partials of the form $\partial E_i/\partial E_k$ are 1 if $i = k$ and zero otherwise. The partials in the form $\partial B_i/\partial E_k$ represent the "influence" that the free variable E_k has on each element radiosity B_i. These *influence factors* are equivalent to entries of the inverse of the form factor matrix. Once the influence factors are known, the scene can be rerendering with new light source emissivities without resolving the radiosity equations. Besides providing the partial derivatives necessary for the optimization process, explicit storage of the influence factors also allows interactive, near real time, user adjustments to the lighting.

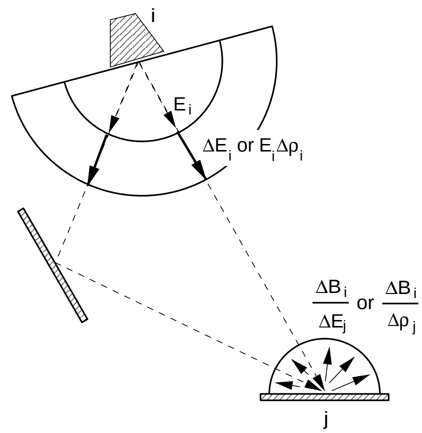

Figure 1: Estimation of $\partial B_i/\partial E_k$ by shooting a delta emission from source k.

Rather than perform an explicit inversion of the block structured system, the partial derivatives can be estimated by finite differences. A small "delta" emission, ΔE_j is shot from the variable emission light source as indicated in figure 1 and allowed to interreflect. The iterative shooting operation are very rapid since the form factor matrix was already computed during the baseline rendering and is stored as a network of links between elements in the hierarchy.

The result of shooting a small amount of energy through the network of links results in its effect on each element radiosity, ΔB_i, thus providing all the derivative estimates $\Delta B_i/\Delta E_j$. If the only free variables in the optimization are light emissions, these influence factors need only be evaluated once, due to linearity. On the other hand, if any element reflectance is allowed to be variable, light emission influence factors must be updated each time one or more element reflectance is changed.

The partial derivative of the objective with respect to a variable element reflectivity is handled in a similar fashion. The element reflectivity ρ_k is adjusted by a small delta $\Delta\rho_k$. The effect on all other elements can be evaluated by "shooting" the unshot radiosity due to the change in reflectivity: $B_k\Delta\rho_k$. As for light sources, several shooting iterations may be necessary in order to account for multiple bounce effects. Once convergence has been achieved, the effect of $\Delta\rho_k$ on element radiosity ΔB_i is available and the influence factor estimate $\Delta B_i/\Delta\rho_k$ can be recorded.

Influence factors for spot light directionality variables, $\mathbf{V}_k$ and n_k are also approximated through finite differences. For example, a small change can be made to the direction vector $\Delta\mathbf{V}_k$ and the effect on each element radiosity can be determined by a series of shooting steps. The first shooting step, illustrated in figure 2, shoots a delta emission from the modified spot light to all other elements. The delta emission is determined according to the change in the directionality parameter, in this case, $E_{i_k}\frac{(n+1)}{2}(cos^n(\phi_{\mathrm{v}+\Delta\mathrm{v}})-cos^n(\phi_{\mathrm{v}}))$ where ϕ_{v} is the angle between the original direction vector of the light and the direction of the element and $\phi_{\mathrm{v}+\Delta\mathrm{v}}$ is the angle between the *new* spot light direction vector and the direction of the element. Note that this delta emission

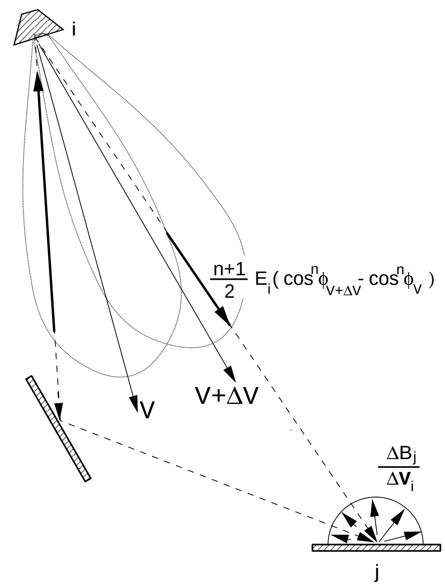

Figure 2: Estimation of $\partial B_i/\partial \mathbf{V}_k$ by shooting a delta emission from source k.

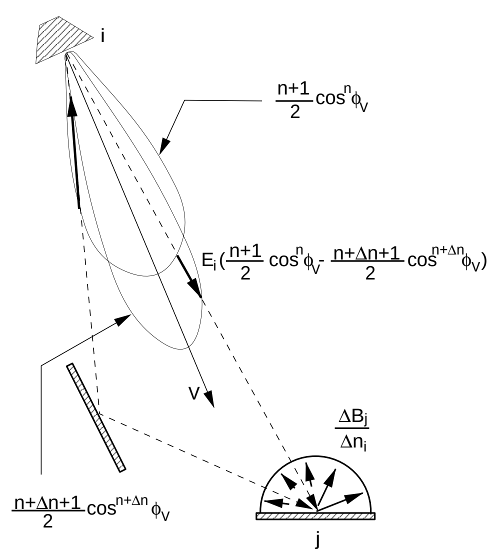

Figure 3: Estimation of $\partial B_i/\partial n_k$ by shooting a delta emission from source k.

will be negative in some directions. Subsequent shooting steps proceed in the normal fashion in order to handle multiple bounce effects. The same technique can be used when the distribution pattern parameter n is changed as illustrated in figure 3. In this case the directionality weight is $E_{i_k}(\frac{(n+1)}{2}cos^n(\phi_v) - \frac{(n+\Delta n+1)}{2}cos^n(\phi_v))$.

The cost functions that measure patterns of light or subjective impressions are defined in terms of perception. The partial derivatives of the functions examining lighting patterns with respect to light emission, E_k are:

$$\begin{aligned}
\frac{\partial f_{brightness}}{\partial E_k} &= -W_{brightness}\frac{\sum_i \frac{\partial P_i}{\partial E_i} A_i}{\sum_i A_i} \\
\frac{\partial f_{non-uniform}}{\partial E_k} &= -W_{non-uniform}\left[\frac{\sum_i (P_{avg,i} - P_i)^2 A_i}{\sum_i A_i}\right]^{-\frac{1}{2}} \left[\frac{\sum_i (P_{avg,i} - P_i)\left(\frac{\partial P_{avg,i}}{\partial E_k} - \frac{\partial P_i}{\partial E_k}\right)}{\sum_i A_i}\right] \\
\frac{\partial f_{peripheral}}{\partial E_k} &= W_{peripheral}\left(\frac{\sum_i \frac{\partial P_i}{\partial E_k} A_i}{\sum_i A_i} - \frac{\sum_j \frac{\partial P_j}{\partial E_k} A_j}{\sum_j A_j}\right)
\end{aligned}$$

The partials of the subjective impressions are just a linear combination of the partial derivatives of $f_{brightness}$, $f_{non-uniform}$, and $f_{peripheral}$.

$$\begin{aligned}
\frac{\partial f_{clear}}{\partial E_k} &= 0.89963\frac{\partial f_{brightness}}{\partial E_k} - 0.38098\frac{\partial f_{non-uniform}}{\partial E_k} - 0.58060\frac{\partial f_{peripheral}}{\partial E_k} \\
\frac{\partial f_{pleasant}}{\partial E_k} &= 0.78437\frac{\partial f_{brightness}}{\partial E_k} - 0.52679\frac{\partial f_{non-uniform}}{\partial E_k} + 0.23984\frac{\partial f_{peripheral}}{\partial E_k} \\
\frac{\partial f_{private}}{\partial E_k} &= 0.89064\frac{\partial f_{brightness}}{\partial E_k} + 0.31562\frac{\partial f_{non-uniform}}{\partial E_k} - 0.08648\frac{\partial f_{peripheral}}{\partial E_k}
\end{aligned}$$

The partials $\partial P_i/\partial E_k$ are derived by differentiating equation 2.3.1.

$$\begin{aligned}
P_i &= 10^{aa*log_{10}(B_i/10,000)+bb} \\
\frac{\partial P_i}{\partial E_k} &= 10^{aa*log_{10}(B_i/10,000)+bb}\left[\frac{aa}{B_i}\frac{\partial B_i}{\partial E_k} + \frac{\partial \alpha}{\partial E_k}(0.4log_{10}(B_i/10,000) - ln(10)(0.8\alpha + 2.584))\right]
\end{aligned}$$

where α = adaption level = $EXP\{log_{10}(L_i)\} = \frac{\sum_i log_{10}(B_i/10,000)A_i}{\sum_i A_i}$

If the partials assume that the adaptation level is constant with regard to a change in emission E_k, $\partial\alpha/\partial E_k = 0$. If the change in adaptation is taken into account then differentiating alpha with respect to E_k gives

$$\frac{\partial \alpha}{\partial E_k} = \frac{A_i}{B_i ln(10)\sum_j (A_j)}\frac{\partial B_i}{\partial E_k}$$

4.5 Optimization

The optimization process uses the BFGS algorithm. The BFGS iteration step evaluates the objective function and its gradient at the current location in the design space. The Newton step provides the search direction, $\Delta\mathbf{X}$, and a line search is performed in this direction. The line search first brackets the function minima in the direction $\Delta\mathbf{X}$ then converges to the solution by a sequence of quadratic fit steps. Each step in the line search involves another evaluation of the objective function and hence a reevaluation of the element radiosities given the current position in the design space. Again, these evaluations are rapid since the form factor matrix has already been computed. At each BFGS step, the element radiosities are evaluated and a new image is presented to the user. This allows the user to watch the optimization as it progresses.

5 Experiences and Results

Our first implementation of the Radioptimization system allowed an objective function based only on photometeric measures and did not take into account the psychophysical properties of lighting. The system could sucessfully optimize lighting but required quite a bit of unintuitive "tweaking" of the objective function weights in order to achieve lighting that had the right subjective appearance. These early experiences led us to investigate psychophysical objective functions.

Color plate 3 shows the effects that the subjective impressions have on an optimization. The left image constrains the table to have a small amount of illumination while preserving energy and creating an overall impression of visual clarity. To improve efficiency the optimization was run at a low resolution on a simplified model, without the chairs and television set. The optimization process took 1 minute and 21 seconds on an IBM Model 550 RISC System 6000. The image on the right has the same design goals as the left image except that it tries to elicit an impression of privateness. This optimization took 2 minutes and 11 seconds.

It took two or three hours of performing design iterations before developing an intuitive "feel" of the optimization process and the effects of the weights on the objective function. One of the problems with the design cycle is that there may be local minimums of the specified objective that are visually unattractive. For example in addition to the design goals mentioned above for color plate 3, we needed to add an addition constraint limiting the illumination of the ceiling because pointing the lights directly at the ceiling was an optimal way of increasing the overall brightness of the room.

One drawback of the system at this point is that it is not fast enough to allow a highly interactive feedback cycle for complex models. However since the system allows a designer to think in terms of their own design goals, it requires fewer design iterations to achieve the desired result.

6 Conclusions

This paper has presented a new method of designing illumination in a computer simulated environment, based on teleological or goal directed modeling. We use a library of functions that approximate a room's success in meeting certain lighting design goals such as minimizing energy or evoking an impression of privacy. In order to develop functions that evaluate the impression that a room evokes, we had a number of subjects order a set of images according to a particular impression. Processing this data with INDSCAL, showed that there was a correlation among the subjects of what lighting patterns they considered to be visually clear, pleasant, and private. Once the lighting design goals have been set, the software system searches the space of lighting configurations for the illumination pattern that "best" meets the design specifications. The system absorbs much of the burden for search the design space allowing the user to focus on the goals of the illumination design rather than the intricate details of a complete illumination specification.

Our system explores one possible path in the application of optimization techniques to image synthesis design problems. Many other possibilities remain for future work. Constrained optimization techniques may be more suitable than the unconstrained penalty method technique used here when the weight of the design goals must be satisfied preciously. Discrete optimization methods may be appropriate in some instances, for example when emissivities are constrained to a finite set, *e.g.* $\{60$ Watts $, 100$ Watts $, \cdots\}$. Geometric properties of the model, such as the position of the lights or the size and position of the windows, could be allowed as free variables. More general image synthesis methods could be applied to account for non-diffuse effects such as glare.

Acknowledgments

Pat Hanrahan, Larry Aupperle and David Salzman provided the radiosity software used as a basis for this work. Greg Ward offered suggestions on useful objective functions for lighting design. Shinichi Kasahara participated in discussions about the work as it developed.

The first and second author's work was supported by NFS (CCR-9210587). All opinions, findings, conclusions, or recommendations expressed in this document are those of the authors and do not necessarily reflect the views of the sponsoring agencies.

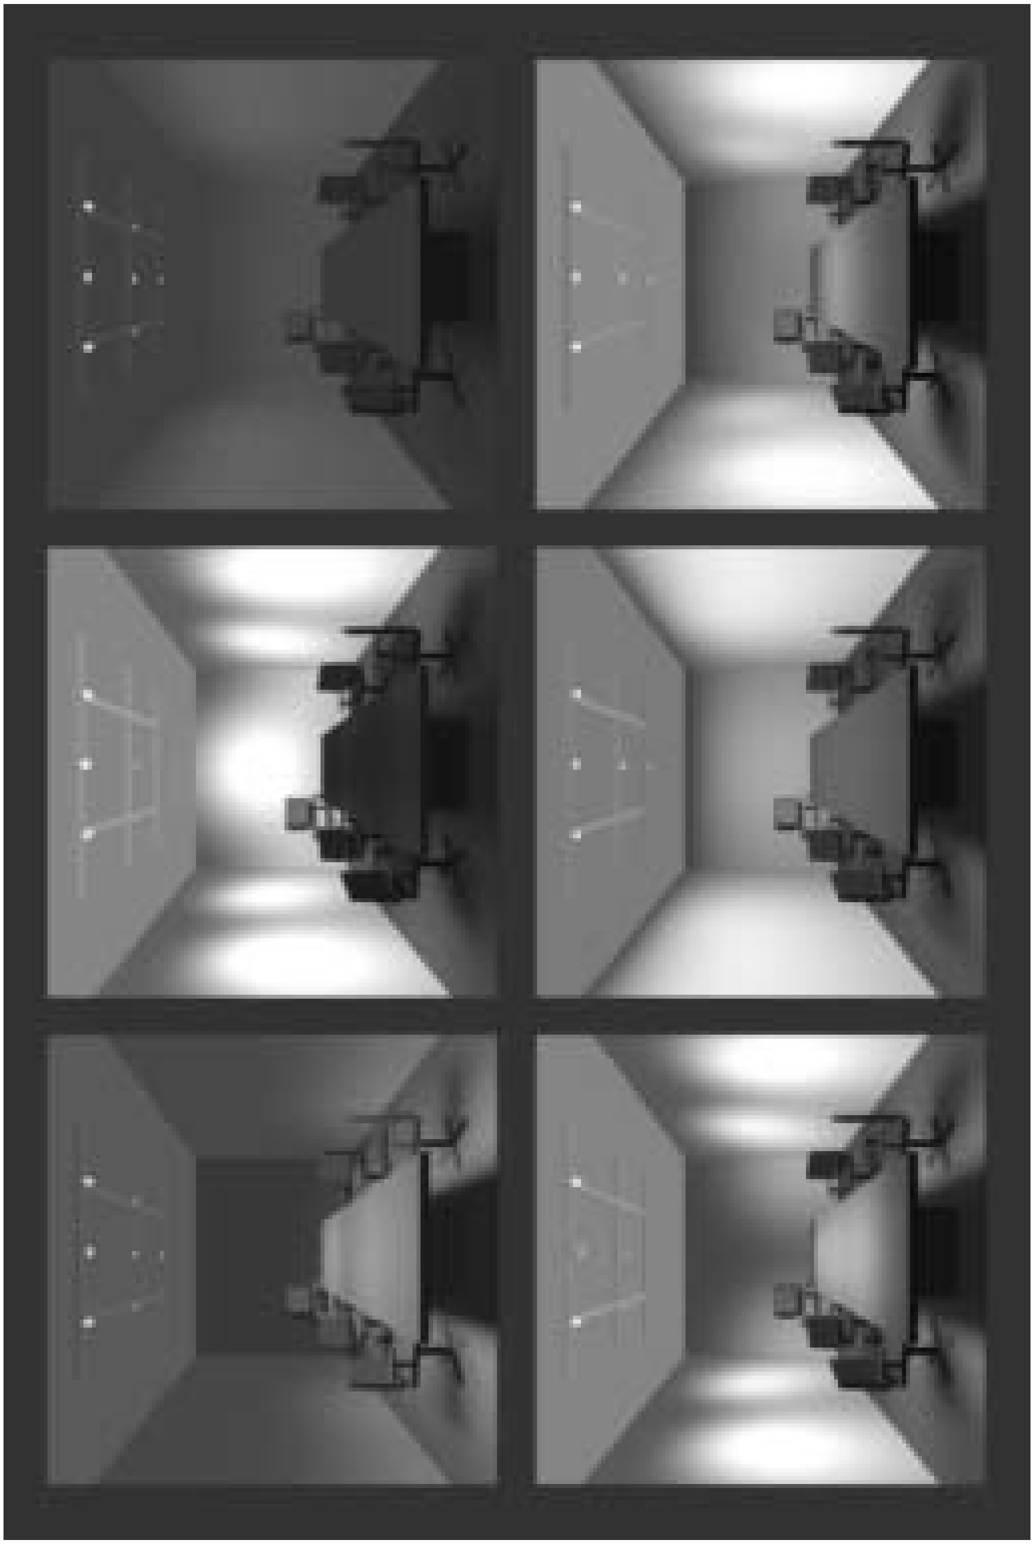

Color Plate 1.

Color Plate 2

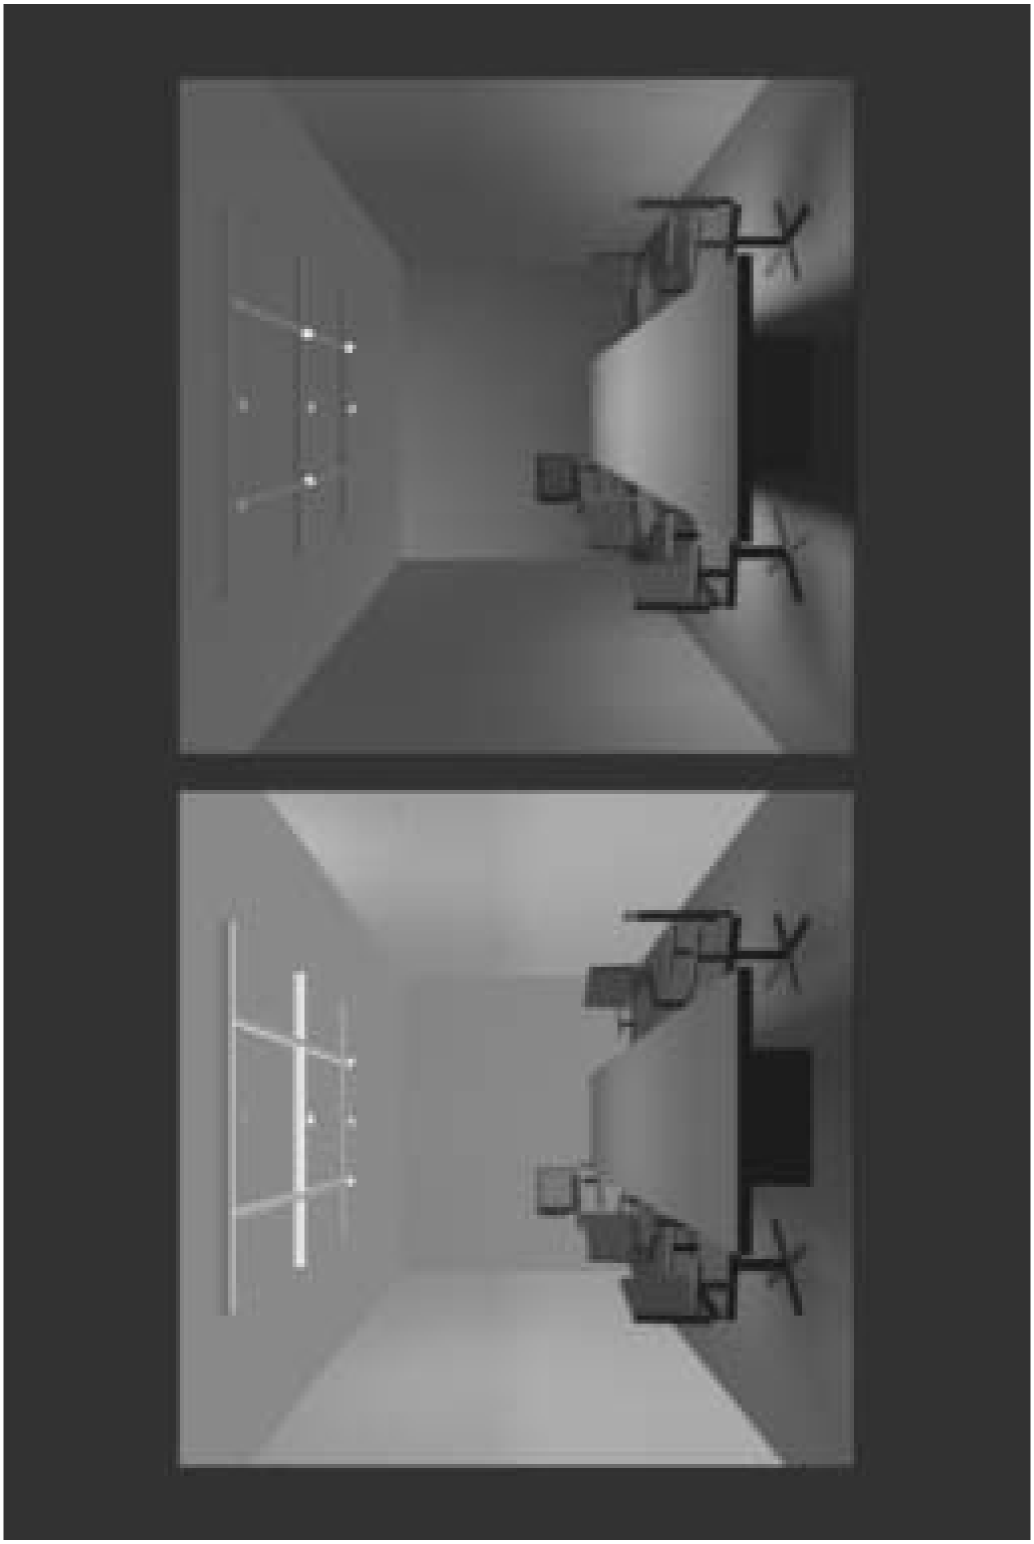

Color Plate 3

References

[1] Barr, A. Teleological modeling. *Siggraph Course Notes #30* (1989), B1–B7.

[2] Chang, J. J., and Carroll, J. D. How to use INDSCAL: a computer program for canonical decomposition of N-way tables and individual differences in multidimensional scaling. Tech. rep., Bell Telephone Laboratories, 1972.

[3] Cohen, M., Chen, S., Wallace, J., and Greenberg, D. A progressive refinement approach to fast radiosity image generation. *Computer Graphics (SIGGRAPH '88 Proceedings) 22*, 4 (July 1988), 75–82.

[4] Cohen, M., and Greenberg, D. The hemi-cube: A radiosity for complex environments. *Computer Graphics (SIGGRAPH '85 Proceedings) 19*, 3 (July 1985), 31–40.

[5] Cohen, M., Greenberg, D., Immel, D., and Brock, P. An efficient radiosity approach for realistic image synthesis. *IEEE Computer Graphics and Applications 6*, 3 (Mar. 1986), 26–35.

[6] Dano, S. *Nonlinear and Dynamic Programming.* Springer-Verlag, New York, 1975.

[7] Ecker, J. G., and Kupferschmid, M. *Introduction to Operations Research.* John Wiley and Sons, New York, 1988.

[8] Fletcher, R. *Practical Methods of Optimization, Vol. 1 and 2.* John Wiley and Sons, 1980.

[9] Flynn, J. E. A study of subjective responses to low energy and nonuniform lighting systems. *Lighting Design and Application* (Feb. 1977).

[10] Flynn, J. E., Hendrick, C., Specer, T. J., and Martyniuk, O. A guide to methodology procedures for measuring subjective impressions in lighting. *Journal of the IES* (Jan. 1979).

[11] Flynn, J. E., Spencer, T. J., and Osyp Martyniuk, C. H. Interim study of procedures for investigating the effect of light on impression and behavior. *Journal of the IES* (Oct. 1973).

[12] Funkhouser, T. A., Sequin, C. H., and Teller, S. J. Management of large amounts of data in interactive building walkthroughs. In *1992 Symposium on Interactive Graphics* (March 1992), SIGGRAPH-ACM, pp. 11–20.

[13] Gill, P. E., and Murray, W. *Numerical Methods for Constrained Optimization.* Academic Press, 1974.

[14] Goral, C., Torrence, K., Greenberg, D., and Battaile, B. Modeling the interaction of light between diffuse surfaces. *Computer Graphics (SIGGRAPH '84 Proceedings) 18*, 3 (July 1984), 213–222.

[15] Green, P. E., Jr, F. J. C., and Smith, S. M. *Multidimensional Scaling Concepts and Applications.* Smith, Allyn, and Bacon, 1989.

[16] Hanrahan, P., Salzman, D., and Aupperle, L. A Rapid Hierarchical Radiosity Algorithm. *Computer Graphics (SIGGRAPH '91 Proceedings) 25*, 4 (July 1991), 197–206.

[17] Himmelblau, D. M. *Applied Nonlinear Programming.* McGraw-Hill, New York, 1972.

[18] Kajiya, J. The rendering equation. *Computer Graphics (SIGGRAPH '86 Proceedings) 20*, 4 (Aug. 1986), 143–150.

[19] Kirk, D. E. *Optimal Control Theory.* Prentice-Hall, Englewood Cliffs, NJ, 1970.

[20] Marimont, R. Model for visual respose to contrast. *Journal of the Optical Society of America 52* (1962).

[21] McKay, H. N. Energy optimization and quality lighting design. *Lighting Design and Application* (Mar. 1986).

[22] Papalambros, P. Y., and Wilde, D. J. *Principles of Optimal Design.* Cambridge University Press, Cambridge, England, 1988.

[23] Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. *Numerical Recipes.* Cambridge University Press, New York, 1986.

[24] Rosen, J. B. The gradient projection method for non-linear programming, part i: Linear constraints. *SIAM 8* (1960), 181–217.

[25] Rosen, J. B. The gradient projection method for non-linear programming, part ii: Non-linear constraints. *SIAM 9* (1961), 514–532.

[26] Sorcar, P. C. *Architectural Lighting for Commercial Interiors.* John Wiley and Sons Inc., 1987.

[27] Stevens, S. S., and Stevens, J. C. Brightness function: Effects of adaptation. *Journal of the Optical Society of America 53*, 3 (Mar. 1963).

[28] Tumblin, J., and Rushmeier, H. Tone reproductions for realistic computer generated images. Tech. Rep. GIT-GVU-91-13, Graphics, Visualization, and Usability Center, Georgia Institute of Technology, July 1991.