Explorations in declarative lighting design

Hai Nam Ha and Patrick Olivier

Informatics Research Institute
University of Newcastle Upon Tyne
Newcastle Upon Tyne, NE1 7RU, UK
h.n.ha@ncl.ac.uk, p.l.olivier@ncl.ac.uk

Abstract. Declarative approaches to lighting design model image quality using an objective function that captures the desired visual properties of an object or scene. The value of the objective function is optimized for a particular camera configuration through the manipulation of the lighting parameters of a scene. We review the notion of declarative lighting design, and introduce LightOp, a tool by which the design of objective functions (the components and settings) and the application of different optimization techniques can be explored. We show how LightOp can be used to explore declarative lighting design through the realization of a number of extensions to existing approaches, including the application and evaluation of stochastic optimization; the use of backlighting to maximize edge enhancement; contrast modeling; and the use of a perceptually uniform color space.

1 Introduction

Lighting design is the problem of finding optimal lighting parameters – positions, directions, colors, and intensities of light sources – in order to maximize the perceptual quality of a specific 3D scene model. Effective lighting can convey much information about a scene and the automatic generation of images with highly discernable visual characteristics is a pressing requirement in application domains such as scientific visualization in which the components of a scene and there relative configuration with respect to the viewers cannot be anticipated in advance. Lighting design is a crucial stage in computer-generated image synthesis and movie making. A lighting design tool is expected to support expert and non-expert users in terms of a reduction in design time and maximization of the visual quality of generated images. Insights from studies of visual cognition are increasingly applied to computer graphics. Indeed, much research has been conducted with the aim of narrowing the gap between the perception of real and computer-generated imagery [McN00][FP04]. Realism in a computer-generated image is not only a matter of physical correctness, but also of perceptual equivalence of the image to the corresponding real world scene. Achieving perceptual equivalence between a computer-generated image and the scene is a significant challenge. Since many perception-based problems must be taken in to consideration. In particular, how to quantify the perceptual quality of an image? Much research has been conducted on the development of perception-based image quality metrics[McN00][RP03][FP04][JAP04]. In this paper, we present a prototype declarative lighting design tool, LightOp, and, motivated by psychological research into human perception, the results of the using LightOp to explore a number of extensions to existing perception-based lighting design algorithms. After a review of existing work in lighting design we extend Shacked and Lischinski's [SL01] approach by applying stochastic optimization, and incorporating new quality components. These extensions aim at revealing features of objects in a scene which in particular support the perception of depth.

2 Previous work

The traditional approach to lighting design for image synthesis is based on direct design methods. Users interactively specify values of lighting parameters, and iteratively render the scene and modify the parameters until the desired visual properties of the scene are achieved. Despite the fact that this can be a tedious and time-consuming process there have been relatively few attempts either to automate or assist the process of lighting design. Schoeneman et al [SDS*93] address lighting design as an inverse problem. Here users set up a set of desired properties that are expected to appear in the final image and the system tries to find a solution (a set of light intensities and colors) whose properties are closest the set of desired properties. Kawai et al. [KPC93] optimize light emission, direction and surface reflectances to obtain the desired illumination for an environment rendered using radiosity-based techniques. In this approach, users have to specify the illumination expected in the final image. Poulin and Fournier [PF92][PRJ97] developed an inverse method for designing light positions through the specification of shadows and highlights in a 3D scene. An interactive sketch-based interface enabled users to sketch desired shadow areas with a mouse pointer. An objective function was defined in a way such that the shadow region for a computed point light (and also some extended light geometries) bounds the sketched regions as tightly as possible. Jolivet et al [JPP02] presented an approach to optimizing light positions in direct lighting using Monte-Carlo methods and reported a declarative paradigm aimed at helping users to specify the lighting goal in an intuitive manner.

3 Perception and lighting

There has been significant effort in recent years in the development of approaches to computer graphics based upon explicit models of a viewer's perception of graphical renderings. Perceptually adaptive approaches have ranged across the entire scope of graphics algorithm and interaction development from schemes for polygon simplification and global illumination that take account of limits on visual attention and acuity, to the design of anthropomorphic animations and gaze-contingent displays [RP03][FP04]. Perception-based lighting design has included implicit approaches that aim to maximize illumination entropy for a fixed viewpoint. Gumhold [Gum02] describes a perceptual illumination entropy approach in which he uses limited user studies to model user preferences in relation to brightness and curvature. In [LHV04] a more explicit model of perceptual preferences is used in the Light Collages framework for which lights are optimized such that the diffuse illumination is proportional to the local curvature, and specular highlights are used only for regions of particularly high curvature.

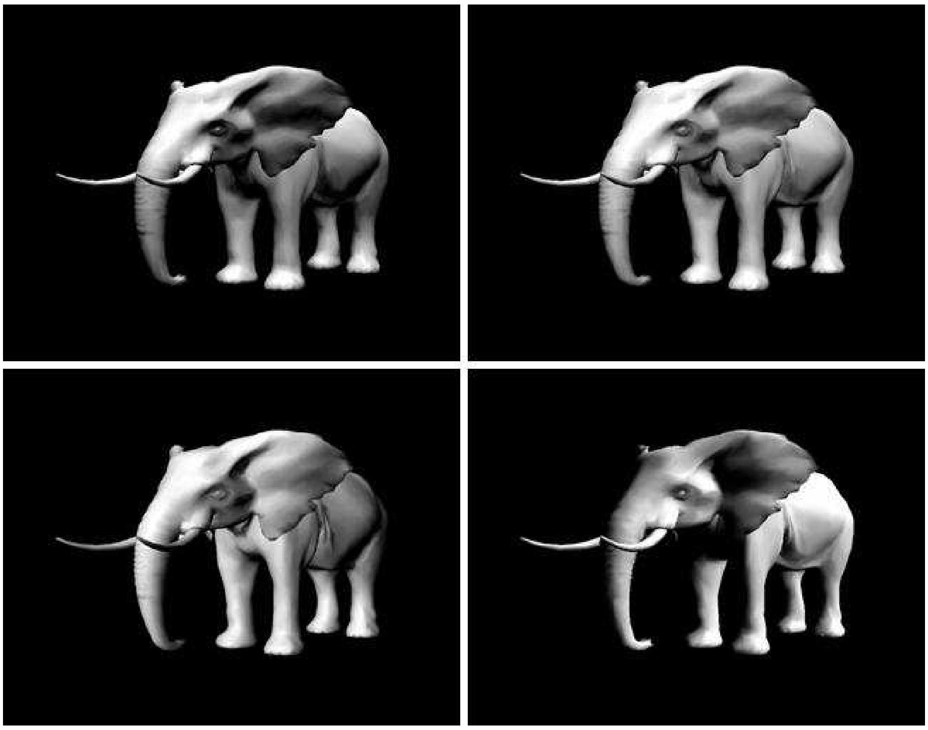

Fig. 1. Shacked and Lischinski base algorithm results (greedy search).

We take as a starting point a declarative lighting approach that maintains an explicit model of object perception due to Shacked and Lischinski [SL01]. In their perception-based lighting design scheme the position and intensity (of specular and diffuse components of a local illumination model) of light sources are optimized using an evaluation function that characterizes separate aspects of low-level processing in the segmentation and recognition of objects. At the heart of this approach is an objective function that is the linear combination of five distinct measures of image quality: edge distinctness (F_{edge}); mean brightness (F_{mean}); mean shading gradient (F_{grad}); intensity range (F_{var}); and evenness of the distribution of intensity in the image (F_{hist}).

$$F(\theta_k, \phi_k, I_{dk}, I_{sk}, R_k) = w_e F_{edge} + w_m F_{mean} + w_g F_{grad} + w_v F_{var} + w_h F_{hist} \quad (1)$$

θ_k: elevation angle of k^{th} light
ϕ_k: azimuth angle of k^{th} light
I_{dk}: diffuse intensity of k^{th} light
I_{sk}: specular intensity of k^{th} light
R_k: distance k^{th} light (fixed for directional lights)
w_e, w_m, w_g, w_v, w_h and w_c: weights for different components in the objective function

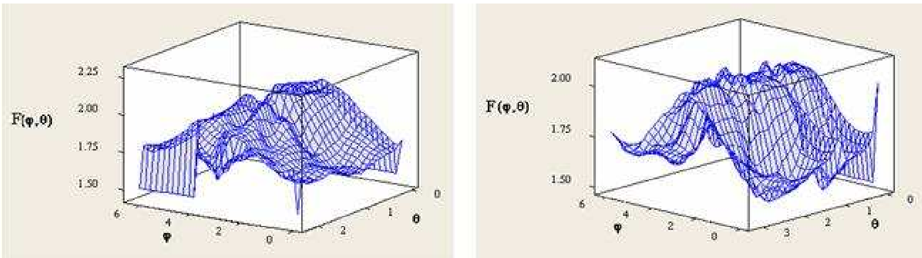

Fig. 2. $F(\theta, \phi)$ for one object (elephant) (left) and $F(\phi, \theta)$ for a two object scene

We have omitted the sixth component of the image quality function used by Shacked and Lischinski which biases the optimization of a key light to a particular elevation and orientation above and in front of the object (relative to the viewpoint). Although this is standard practice in photography and might be justified in terms of evolutionary psychology [Mar94][Mil05] – that our perceptual system evolved for scenes lit by the sun or moon – we take the view that the using the direction of the light as a direct component in an objective function is unjustifiably ad hoc (and procedural). Our objective function is formulated such that lower values of $F(\theta_k, \phi_k, I_{dk}, I_{sk}, R_k)$ correspond to configurations with better visual characteristics and a greedy gradient descent minimization algorithm is utilized in the discovery of appropriate lighting configurations.

4 LightOp: exploring declarative lighting design

LightOp is a tool for the interactive configuration of objective functions and optimization schemes, which we have built to explore the problem of declarative lighting design (see figure 4). A range of optimization techniques have been implemented in LightOp including steepest decent, genetic algorithms, and

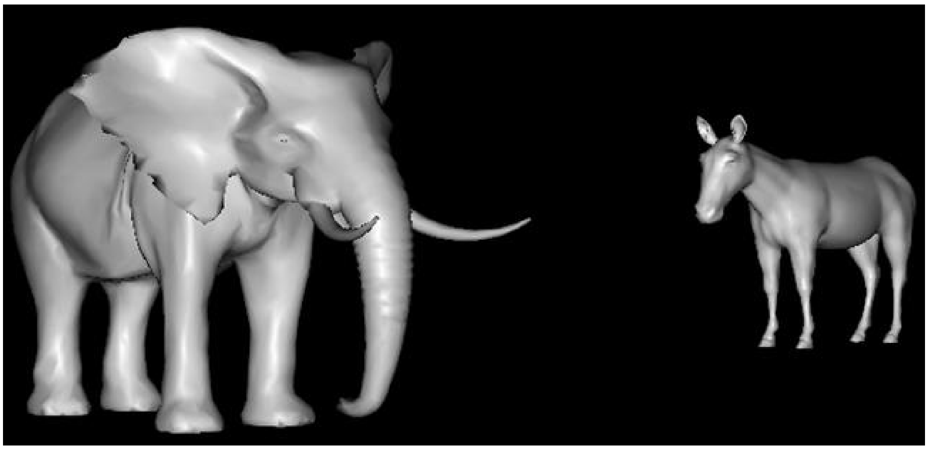

Fig. 3. The simple two object scene used in figure 2 (right).

simulated annealing. Lighting parameters for a scene can be optimized with, and without, shadows, and the number of lights used, the components of the objective function used and nature of the color space can all be interactively specified.

Shacked and Lischinski make no attempt to either characterize the nature of their objective function or assess the suitability of the greedy search employed. From equation (1) it is clear that the multi-objective optimization problem incorporates significant non-linearity, and as the geometric complexity of the scene increases, the number and likelihood of local minima will increase. For a greedy search of a space, such as that illustrated, the solution is highly dependent on the starting condition and figure 1 illustrates the resulting images for four searches conducted from different starting configurations. The nature of the optimization problem suggests a requirement for a more general optimization strategy and we employ a range of stochastic approaches, including a genetic algorithm (GA). The encoding of the lighting problem as a GA is straightforward. The free variables θ_k,ϕ_k, I_{dk},I_{sk},and R_k are encoded directly in the chromosome as real numbered alleles. We have evaluated the optimization problem with varying population sizes and configurations for the GA. From our experiments we established that a population size of 30, 10% elitism and crossover and mutation rates of 80% and 20% respectively were sufficient for our example scenes (figure 8).

Note that the process of optimizing a single (key) light takes place under the constraint that the light position is limited in the values of elevation and azimuth angles (e.g. the position of a light is limited to a quarter sphere in front and above the center of the scene). This constraint must be respected in all stages of the optimization, but particularly at initialization and during the generation of random mutations (members of the population whose values do not reside in

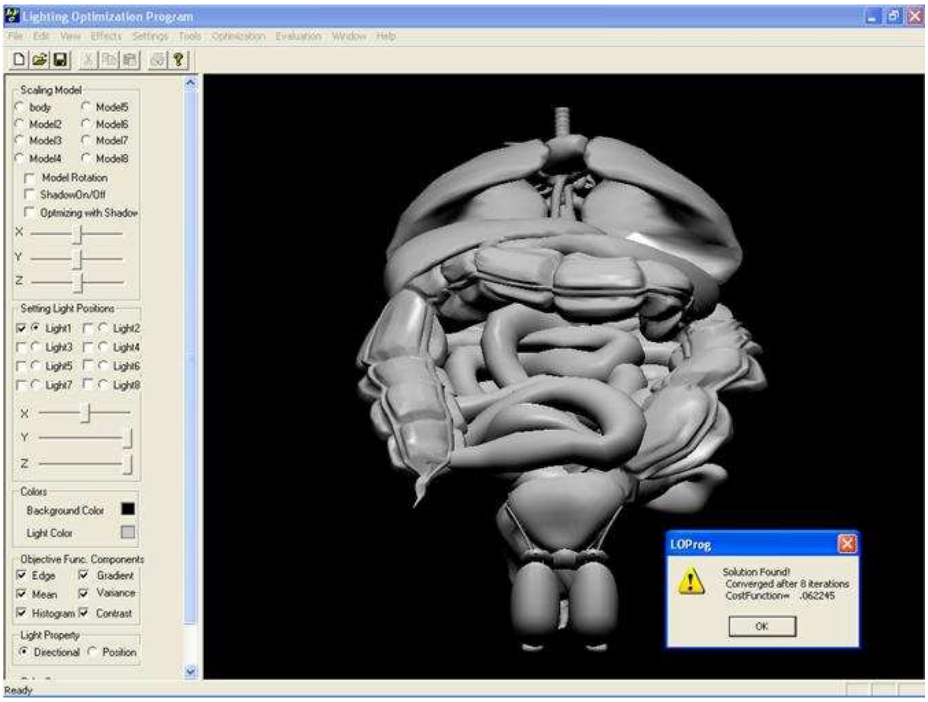

Fig. 4. User interface for LightOp showing the objective function & light parameter controls and the result of an optimization.

constrained ranges are rejected). Note that as the constrained region is convex, crossover cannot yield offspring that violate the constraint. The GA exhibited consistently better results than the greedy search both in terms of the final value of the objective function and the visual quality of the solution. Figure 6 shows a direct comparison of the greedy and GA optimisation results for the elephant model.

5 Enhancing perception-based design

Components of perception-based lighting approaches are motivated directly from the results of studies of human perception, in particular, object recognition [Mar94]. For example, Shacked and Lischinski's edge enhancement criteria relates directly to theories of object segmentation and neuropsychological findings as to the nature of retinal processes [Mar94]. In the course of extending Shacked and Lischinski's approach we identified a number of features currently not addressed in perception-based lighting design: (a) *contrast* : differences in luminance between different surfaces of an object have been shown to convey significant information about the shape and depth of objects [SJ90]; (b) $back{-}lighting$: lighting

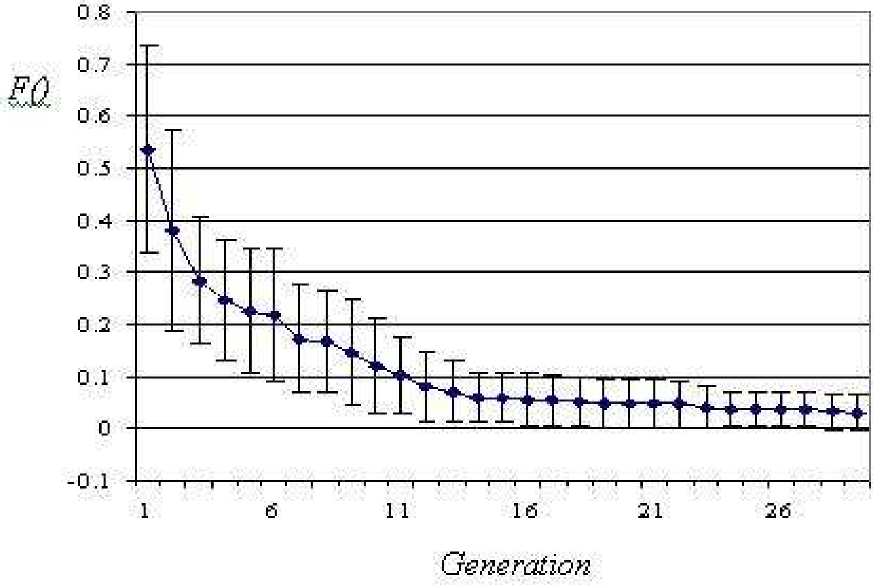

Fig. 5. Averaged results for 10 runs of a 30 member population, the value of F() for the best member and standard deviation.

an object from behind (where the background is dark) is a well established feature of cinematic and photographic practice aimed at maximizing edge enhancement of the silhouette of an object [Mil05]; (c) $perceptually uniform color space$: standard approaches in lighting design implement image quality metrics with respect to RGB (or similar) color spaces, despite the fact that such spaces are highly non-uniform with respect to human judgments of color difference [Mar94].

5.1 Contrast enhancement

Empirical studies of visual cognition have demonstrated that object perception depends on both absolute luminance levels and differences in an object's luminance from its background [8]. We have included this notion through the provision of a means of evaluating differences in luminance between adjacent parts of an object and incorporating this in our objective function. The contrast between two parts of an object is given by:

$$C_{i,j} = \frac{(Y_i - Y_j)}{Y_i} \tag{2}$$

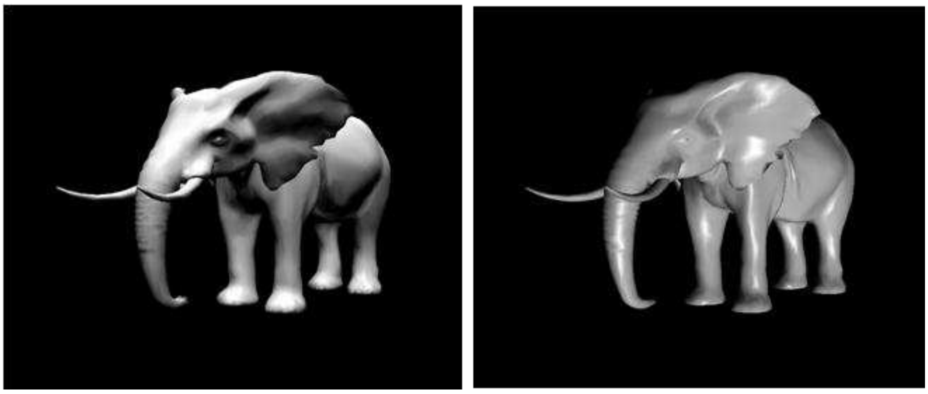

Fig. 6. Greedy optimization result for the Shacked & Lischinski [1] objective function (left), and genetic algorithm optimization result for the Shacked & Lischinski [1] objective function (right).

$C_{i,j}$: contrast between part i and part j Y_i: the mean luminance of part i Y_j: the mean luminance of part j The mean luminance of a part is calculated as follows:

$$Y_i = \frac{1}{N_i} \sum_{I(x,y) \in P_i} I(x,y) \tag{3}$$

A pixel type map of objects in a 3D scene is extracted by applying an edge detector operator to the depth buffer [ST90]. Edges in the pixel type map correspond to boundaries between parts of an object. With this assumption, we developed an algorithm to calculate the contrast between adjacent parts of an object using the pixel type map. The algorithm, which is supported by definition 1, can be described in pseudo-code as follows:

Definition 1. *$p(x_i, y_i)$ is a run-length from $p(x_j, y_j)$ if $p(x_i, y_i)$ is the same pixel type as $p(x_j, y_j)$ and all the pixels, which belong to a horizontally continuous connection between $p(x_i, y_i)$ and $p(x_j, y_j)$, are the same pixel type as $p(x_j, y_j)$.*

```
For each pixel p(x,y) in the pixel type map
    If (p(x,y) is an EDGE pixel)
        Begin
            x_r = x
            Repeat
                x_r = xr +1
                Calculate mean_right_in
            Until (p(x_r,y)  is BACKGROUND pixel) OR
                     ((p(x_r,y) is EDGE pixel)) AND
                      (p(x_r,y) not RUN-LENGTH from  p(x,y))
```

```
      x_l = x
      Repeat
          x_l = x_l -1
          Calculate mean_left_in
    Until (p(x_l,y)  is BACKGROUND pixel) OR
               ((p(x_l,y) is an EDGE pixel)) AND
                (p(x_l,y) not RUN-LENGTH from p(x,y))
          Calculate mean_contrast from mean_right_in and mean_left_in
          Add mean_contrast to mean_global_contrast
  End
```

Figure 7 (top-left) shows the effect of the contrast component on the example. As can be seen in the other examples shown in figure 8 (top), it is apparent that the contrast component has a tendency to increase the apparent shininess of objects and yields a contrast between adjacent parts of an object that leads to better cues as to depth and the relative positions of parts of an object.

5.2 Edge enhancement & back-lighting

In practical contexts, such as a photographer's studio or a film set, backlighting is a useful technique that often makes a valuable contribution to pictorial lighting. Backlighting aids enhancement of external edges and facilitates a viewer in segmenting objects from the background. This is especially true where the surfaces of the object are dark, likewise, backlighting catches the sharply-folded contours and gives them shape and solidity. In our approach, the initial position of the backlight is calculated as a centroid of a set of vertices of objects in a 3D scene. During the optimization process the azimuth angle of the backlight is fixed. Since backlighting primarily effects external edges, we only use edge component in the objective function when optimizing parameters of the backlight. In reality, a backlight is always positioned above and behind objects, so that the elevation angle of the backlight is confined to a predefined range, and we implement this constraint in our optimization process. Backlight optimization can be viewed as a distinct process from establishing the position and intensity of the key light. We implement it as a second optimization stage using a greedy algorithm. The impact of adding backlighting to the contrast enhanced objective function can be seen in figures 7 (top-right) and 7 (bottom-left). Figure 7 (top-right) shows the result due to the back light alone, and figure 7 (bottom-left) shows the image resulting from the addition of the backlight in which the geometry of the feet is more clearly discernable as are the top of the head and the top of the right ear.

5.3 Perceptually uniform color space

Despite the fact that perception-based graphics algorithms attempt to model and quantify human responses to visual stimuli, they routinely use standard RGB (or

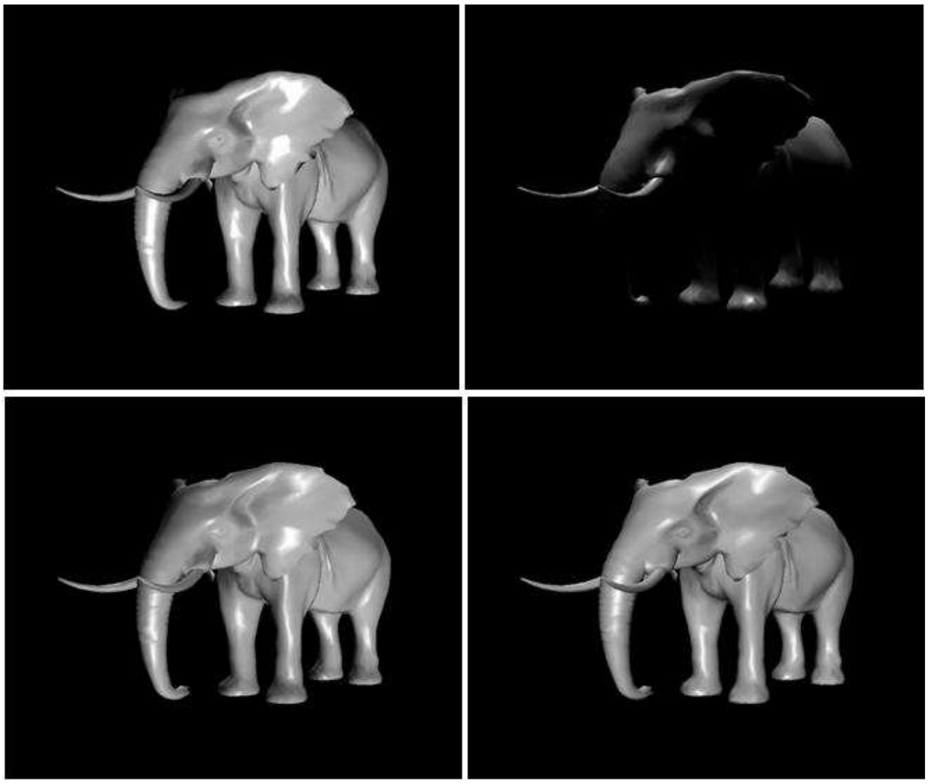

Fig. 7. From top-left clockwise: contrast; effect of the back-light; contrast and a back-light; and contrast, back-light, & perceptually uniform color space.

related) color spaces, and ignore the fact that such color spaces are highly non-linear with respect to our judgments of color difference. As an alternative, we transform RGB color values to the CIE L*, a*, b* color space – which is approximately uniform [Hal93]. In computing components of the objective function in a perceptually uniform color space the goal is to obtain edge contrasts, gradients and variances that are perceptually distinct. Even for target-based components such as Fmean the error function in the optimization will be enhanced due to a linear relationship between computed (from the color space) and perceived differences. Figure 7 (bottom-right) illustrates a result of incorporating the perceptually uniform color space.

6 Discussion

We have used LightOp, our prototype declarative lighting tool, to explore a number of enhancements to the perception-based lighting design approach pro-

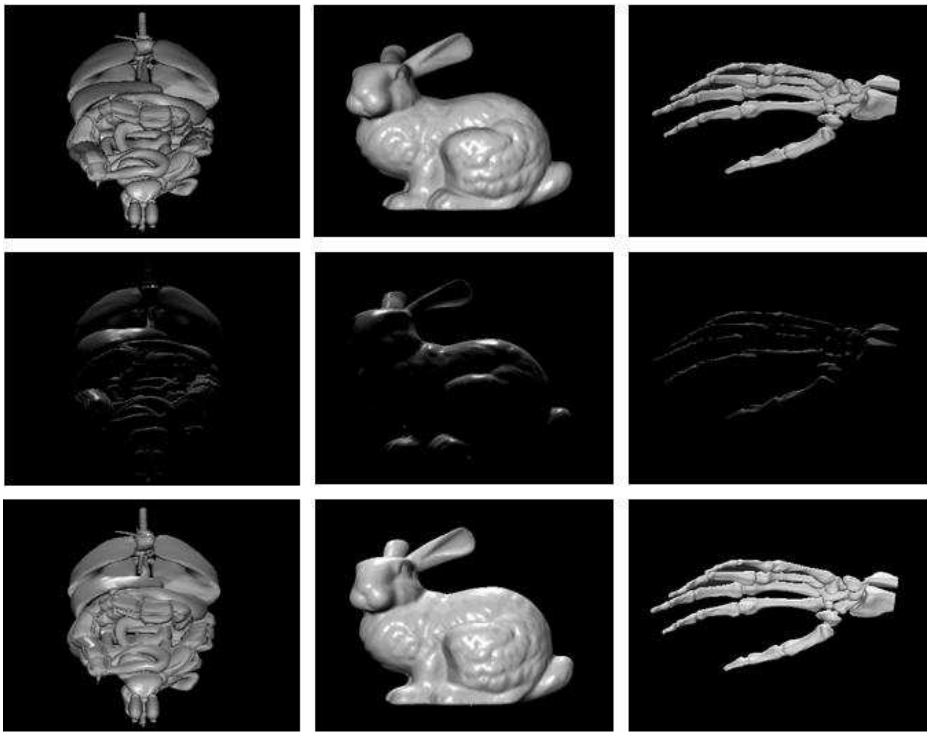

Fig. 8. Sample results for three models: contrast (top); effect of back-light (middle); contrast, back-light and perceptually uniform color space (bottom).

posed by Shacked and Lischinski [SL01]. Much work remains to be done and it is our intention that LightOp will serve as the vehicle for a thorough investigation both of the appropriate nature of the objective function, and detailed empirical studies by which users' responses to differently lit scenes are more rigourously investigated. For example, numerous standard tasks in the study of spatial cognition, such as mental rotation and object recognition, rely on the ability of users to identify key aspects of the geometric nature of objects. The performance of users on such tasks is well understood and we intend to leverage such methodologies in charactersing objective functions that are in some way optimal with respect to these tasks.

One should note, however, that this notion of *ideal lighting* is only appropriate with respect to a minority of real-world graphics applications. For example, in some domains, in the automatic lighting of 3D visualizations where the number, position and orientation of 3D glyphs or other objects cannot be predicted in advance (and are typically not textured), the user's ability to recognise and reason about the external representations requires the lighting to be *ideal*. However, in it is more typical for lighting configurations to be used to convey more aesthetic characteristics of a scene such as mood, emotion, and factors other than the purely geometrical. The investigation of aesthetic responses to lighting is a considerable enterprise in itself, though we believe that tools such as LightOp, and a rigorous understanding of perceptually adaptive lighting, are a necessary prerequisite for any such undertaking.

References

[Bar97] Barzel, R.: Lighting Controls for Computer Cinematography. Journal of Graphics Tool, vol. 2, no. 1, pp 1-20, 1997.

[FP04] Farugia, J. P., Peroche, B.: A progressive rendering algorithm using an adaptive perceptually based image metric. In Eurographics conference proceedings (2004).

[Gum02] Gumhold, S.: Maximum entropy light source placement. IEEE Visualization 2002, pp. 275-282.

[Hal93] Hall, R.: Illumination and Color in Computer Generated Imagery. Monographs in Visual Communication. New York, Springer, 1993.

[JAP04] Jean, P. F., Albin, S., Peroche, B. A perceptual Adaptive Image Metric for Computer Graphics. In Proc. International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision WSCG04, 2004.

[JPP02] Jolivet, V., Plemenos, D., Poulingeas, P.: Inverse Direct Lighting with a Monte Carlo Method and Declarative Modelling. CG&GM'02, 2002.

[KPC93] Kawai, J., Painter, J., Cohen, M.: Radioptimization – Goal Based Rendering. Proc. SIGGRAPH '93, pp. 147-154, 1993.

[LHV04] Lee, C. H., Hao, X., Varshney, A., Light Collages: Lighting Design for Effective Visualization. IEEE Visualization 2004, Austin, October 2004.

[MAB*97] Marks, J., Andalman, B., Beardsley, P. A., Freeman, W., Gibson, S., Hodgins, J., Kang, T., Mirtich, B., Pfister, H., Ruml, W., Ryall, K., Seims, J., and Shieber, S.: Design galleries: a general approach to setting parameters for computer graphics and animation. Proc. SIGGRAPH, pp. 389-400, 1997.

[Mar94] Markus G., Visual Computing, Springer-Verlag, 1994.

[Mil05] Millerson, G.: Lighting for Television and Film, 3rd Edition., Focal Press, 2005.

[McN00] McNamara, A.: Comparing real and synthetic scenes using human judgments of lightness. doctoral dissertation , Univ. Bristol , 2000.

[PF92] Poulin, P., Fournier, A.: Lights from highlights and shadows. In Proceedings of the 1992 Symposium on interactive 3D Graphics (Cambridge, Massachusetts, United States). SI3D '92. ACM Press, New York, NY, 31-38, 1992.

[PRJ97] Poulin, P., Ratib, K., Jacques, M.: Sketching Shadows and Highlights to Position Lights. cgi, p. 56, Computer Graphics International 1997 (CGI'97), 1997.

[RP03] Reitsma, P. S. A., Pollard, N.S.: Perceptual metrics for character animation: sensitivity to errors in ballistic motion. ACM Transactions on Graph. 22:3, 537-542, 2003.

[SDS*93] Schoeneman, C., Dorsey, J., Smits, B., Arvo, J., and Greenburg, D.: Painting with light. Proc. SIGGRAPH. 143-146. 1993.

[SJ90] Spillmann, L., John, S. W.: Visual Perception the Neurophisiological Foundations. Academic Press Inc., 1990.

[SL01] Shacked, R.,Lischinski, D.:Automatic Lighting Design using a perceptual quality metric. EuroGraphics, vol. 20, no.3, 2001.

[ST90] Saito, T., Takahashi, T.: Comprehensible rendering of 3-D shapes. Proc. SIGGRAPH, pp. 197-206, 1990.