"Smart" Approaches to Lighting Design

Intelligente Komposition von Lichtquellen in Computergraphiken

Patrick Olivier, Newcastle University, Newcastle Upon Tyne, UK,
Hai NAM Ha, PTIT, Hanoi, Vietnam,
Marc Christie, Université de Nantes 2, LINA, France

Summary The declarative design of graphics lies at the core of the Smart Graphics research agenda. For 3D graphics the number of lights included, and the properties of these lights, has an enormous impact on what a viewer can judge about the content (the objects), properties (the geometric characteristics and spatial relations of the objects) and other aesthetic qualities of a scene. The traditional approach to lighting design for image synthesis is based on manual design methods, whereby users interactively specify values of lighting parameters, render the scene, and modify the lighting parameters until the desired visual properties of the scene are achieved. Non-expert users encounter a number of difficulties in selecting the appropriate lighting parameters, as the process requires both a subtle technical and aesthetic understanding of lighting in computer graphics. We review range of "smart" lighting design and steady slow convergence on ideal lighting approaches which optimise the lighting configuration for a scene with respect to a set of absolute perceptual metrics. More recently perceptual approaches have been combined with aspects of exemplar driven approaches to yield "lighting-by-example" techniques that can replicate the lighting of existing static 2D images and 3D scenes. ▶▶▶ **Zusammenfassung** Die Entwicklung deklarativer Methoden zur Generierung von Graphiken ist eines der Hauptanliegen von Smart Graphics. Die Anzahl und Eigenschaften der Lichtquellen haben einen enormen Einfluss auf die Bildwahrnehmung, insbesondere auf den Bildinhalt (Art der Objekte), und die Bildeigenschaften (die geometrischen und räumlichen Beziehung von Objekten), und bestimmen, welchen ästhetischen Gesamteindruck ein Bild beim Betrachter hinterlässt. Der klassische Ansatz, Lichtquellen richtig auszuwählen und zu platzieren, ist traditionell von einer starken interaktiven, manuellen Vorgehensweise geprägt. Der Benutzer spezifiziert alle wichtigen Parameter, bis der gewünschte Lichteffekt eintritt. Insbesondere Laien haben große Probleme, mit Hilfe dieser Verfahrensweise akzeptable Resultate zu erzielen, da der Vorgang sowohl ein technisches als auch ein ästhetisches Verständnis einfordert. In diesem Artikel untersuchen wir verschiedene Ansätze, dieses Problem zu automatisieren, indem die Anordnung und Parametrisierung der Lichtquellen auf ein bestimmtes absolutes perzeptuelles Zielbild hin optimiert wird. Neuere Ansätze werden präsentiert, die beispielorientiert arbeiten und es Benutzern so ermöglichen, Lichtquellenarrangements zu finden, die klassischen 2D- und 3D-Anordnungen sehr ähnlich sind.

Keywords I.3.7 [Computer Methodologies: Computer Graphics: Three-Dimensional Graphics and Realism]; I.3.6 [Computer Methodologies: Computer Graphics: Methodology and Techniques]; Computer Graphics, Smart Graphics, lighting, optimization, constraints ▶▶▶ **Schlagwörter** Computergraphik, Lichtquellen, Optimierung

1 Introduction

At the heart of the Smart Graphics' research agenda is the goal to develop tools that allow both offline and real-time declarative control of graphics. The idea of being able to allow users free access and interaction with complex graphical content whilst at the same time being able to guarantee aspects of their visual experience is highly desirable but extremely challenging. In recent years we have seen significant advances in the development of declarative approaches to some aspects of what we might call *computational cinematography*, for example, camera control in 3D graphics. Christie et al. [3] review an evolution of declarative approaches to camera control that range from simple ray casting approaches and fixed viewpoint selection, to advanced and expressive constraint satisfaction and optimisation schemes.

A related problem that has received significantly less attention is the automatic configuration of the lighting for 3D scenes. The number of lights included, and the properties of these lights, has an enormous impact on what a viewer can judge about the content (the objects), properties (the geometric characteristics and spatial relations of the objects) and other aesthetic qualities of a scene. In a number of design scenarios, the viewpoint on a scene is normally known prior to the design of the lighting. The traditional approach to lighting design for image synthesis is based on manual design methods where users interactively specify the values of lighting parameters, render the scene, and modify the lighting parameters until the desired visual properties such as brightness, shading and contrast of the scene are achieved. Non-expert users are likely to encounter difficulties in adjusting lighting parameters to achieve their desired effects. Indeed, adjusting lighting parameters is not a random process but one driven by both a technical and aesthetic appreciation of lighting in computer graphics. Even experienced artists can find that manual lighting design is a tedious and time-consuming process.

By analogy to camera control, automated approaches to lighting design should be able to exploit the fact that we know the precise location of all the edges of an object – this includes both silhouette edges, edges between object parts, and self-occluding edges –, we also have complete information about which pixels in the image correspond to which geometric features of the objects. Indeed previous approaches have adopted similar notions to camera control, such as the use of augmented interactive control and models of our perception of graphics. However, there have been relatively few attempts either to automate or assist the process of lighting design and of those approaches we can broadly divide the solutions into approaches that treat lighting design as an inverse design problem, as a perceptual graphics problem and one as an example-based authoring problem.

2 Inverse Lighting Design

Schoeneman et al. [24] address lighting design as an inverse problem. Here, users set up a set of desired lighting effects that are expected to appear in the final image; the system tries to find a solution in which the lighting effects are closest to the set of desired lighting effects. The lighting parameters optimised in this approach are light intensities and colours. It is claimed that the reason why light positions themselves are not optimised is that users often know where to put the lighting, but not how to combine them in terms of setting up intensities and colours for lights. Optimal lighting parameters are found at each step on the basis of minimizing an objective function, which is defined as an error function between a set of target functions and a set of approximations of functions.

Such functions can be ray traced images [30] of each scene from the same view point or radiance functions over surfaces in the scene computed by radiosity method [2]. Directly painting on the surfaces of the rendered scene causes a change in the surface radiance functions, and these after-painted surface radiance functions are used as target values in the optimisation process that follows. Painted surfaces in the rendered image are given more weight, which biases the optimisation towards solutions with properties that best match the painted surfaces. In this approach, painted surfaces can be considered as examples affecting the target radiance surface functions. Solutions are limited to direct illumination from lights in order to achieve appropriate interactive time with users. The system employed the method of finding quickly the patches being affected by the brush developed by [14] in order to achieve interactive speeds. This approach has a potential for a lighting design tool which assists users to set up intensities and colours of lights. One of the drawbacks of this "painting with light" approach is that there is always the possibility that a user will paint on surfaces that cannot be illuminated by the preconfigured lights. Although some constraint-based approach was employed, this still appears to be an incomplete solution.

Kawai et al. [17] proposed a goal-based approach to lighting design, which can also be considered an inverse approach. The lighting parameters that were optimised in this approach were light source emissivity, spotlight direction, spot light focus, element radiosity and element reflectivity; again the light positions were fixed in this approach. The objective function was based on an energy calculated using the radiosity method (and a set of components loosely motivated by research in visual perception). Perceptual components in the objective function were motivated by Flynn [6; 7] and which try to quantify parameters that elicit a shared human behavioural response and a subjective impression. In particular, the impact of non-uniform, peripheral and bright light, on impressions of clarity, spaciousness, relaxation and privacy, were examined. Once the objective function was defined, the system tried to minimize the objective function in an optimisation phase. In this approach, the optimisation problem was constrained by goals defined by users that were either physical constraints, design goals or barrier constraints. Physical constraints related light emission and element radiosities as determined by the physical nature of light transport. Design goals were subdivided into equality and inequality constraints on radiosity values. Barrier constraints were bound conditions on optimisation variables that must be satisfied. One of the drawbacks of this approach is that it is not suitable for novice users who would normally not be familiar with terms such as clarity, spaciousness, relaxation and privacy, which themselves should be defined by users via an interactive interface of the system. Another drawback of this approach is that the perceptual components in the objective function seem to be specifically defined for indoor room environments.

Shadows tell us a great deal about the depth and shape of objects and their spatial relations to light sources. Object-light relations can be extracted from the shadows and highlights themselves. Poulin and Fournier [21] developed an inverse method for designing light positions through the specification of shadows and highlights in a 3D scene. In this approach, shadows and highlights were used in the specification of the shape and position of a light source. Highlights are strongly related to the specular element of a reflection model, and a user can define a highlight on the surface of an object by positioning two points that control the roughness coefficients of the surface. The first defines the maximum intensity point of the highlight, while the second determines the boundary of the highlight. When the specular term of Phong's shading model reaches a predefined threshold, the surface roughness coefficient is calculated. These two points are sufficient to determine both the direction of a directional light and the surface roughness coefficient. A boundary fill algorithm was developed to determine the shape of the highlight. Highlights in the scene must be recomputed if the camera position is changed, as highlight information is strongly dependent on camera position.

This approach also developed an algorithm for determining light positions and direction from information extracted from the properties of shadows, in which a shadow was defined by manipulating a curve using so-called pivot points. Light positions found by using shadow information are, of course, independent of view point. One of the drawbacks of using highlight information to determining light positions is that highlights just contain information about the direction of light sources. For light types other than directional light sources, more constraints are needed to properly interpret the highlight information.

In an extension to the approach proposed in [21], Poulin et al. [22] developed an interactive sketch-based interface that enabled users to sketch desired shadow areas with a mouse pointer. The sketching process starts with selections of an object to cast a shadow and a light. A user sketches strokes on a 2D image plane projected directly from a 3D scene and the light position is recalculated at every stroke. An objective function was defined in such a way that the shadow region for a computed point light (and also some extended light geometries) bounds the sketched regions as tightly as possible. The objective function is based on distance between sketch points and the perspective position of a light, which is maximized in the optimisation process.

Finally, Jolivet et al. [15] presented an approach to optimising light positions in direct lighting using Monte-Carlo ray casting techniques [27]. The core of the approach was to stochastically shoot rays at patches to be illuminated to find out patches that are directly visible. When a ray was shot, the first patch reached by this ray would be added to a candidate list as a candidate for the light position. The ray shooting process stopped when no patch was hit after a predefined number of rays were shot. The form factors were then calculated for every candidate patch in the candidate list on the basis of the ratio of the number of rays hitting the patch to total number of rays cast. The main limitation was the use of just a single quadrilateral emitting patch, built from a candidate list of patches using an algorithm developed in [26]. This approach also proposed a lighting design paradigm based on declarative modelling approach developed in order to help users specify the lighting goal in intuitive-using linguistic descriptions. The declarative model developed also used fuzzy logic [4] to allow the use of a set of modifiers to enrich the linguistic description of the lighting goal. Again, the main limitation is that other light properties, such as intensities, were not considered, and that it was not clear how to choose properly representative patches from surfaces to be lit.

3 Semi-Interactive Approaches

Design Galleries [19] was one of the first practical interactive tools for lighting design, although it can also be considered to have aspect of an examples-based approach (despite the fact that there is no specific example). Through its generation of a wide range of clustered images as examples, Design Galleries presents sets of exemplars for users to perform selection on as part of a render-selection loop. Thus, there is no information about what effects users want to have in the final images, but the user has the opportunity to select likely candidates. Marks et al. try to build a mapping function between an input vector that contains light position, light type, and light direction and an output vector that contains a set of values that summarizes the perceptual qualities of the final image. Possible light positions are located at so-called light hook surfaces defined by the user. During the optimisation process lights are moved from one predefined position to another. At each position, a light type is selected from a set of light types, which describes attributes of a light such as its basic class (e. g., point light, area light, or spotlight), characteristics of a shadow cast by that light, its fall-off behaviour and class specific parameters. Corresponding to a light position and a chosen light type, an image is generated and an output vector created. An output vector is calculated on the basis of luminance of images at different low resolutions that are differently weighted.

With a large number of light positions and light types, large sets of input and output vectors are created. A dispersion algorithm was developed in order to achieve a set of characteristic input vectors which creates a set of output vectors that covers the pre-dispersed set of output vectors. Dispersion results in a much smaller set of input vectors, which are a drastic reduction in the computational complexity. The dispersion step is based on distances between images, and a luminance-based distance metric was proposed for measuring the distance between two images. Design Galleries provide users with a large number of images to select from. Of course, the more images that are presented, the more difficult it is for users to select a particular image, so a graph-based partition scheme based on [16] was applied to group images which have similar illumination effects. The edge costs of the graph used in the partitioning scheme are based on the inverse of the distance metric used in the dispersion step.

One of the limitations of Design Galleries is that the number of images presented to users is large, even with the application of the partition scheme. In addition, the distance metric used in dispersion and partition phases is simply constructed on the basis of the luminance distance between images, which is hard to justify in terms of our visual perception of image difference. Finally, the definition of light hook surfaces as well as light positions on those surfaces appears arbitrary and is not well motivated. Non-photorealistic lighting is not constrained by physical correctness. Gooch et al. [8] proposed a lighting model that uses luminance and changes in hue to convey surface orientation, edges and highlights. By recognizing that, in technical illustrations, the communication of shapes and forms is more important than realism, they use shading to convey direction and the shape of the surfaces.

A shading technique based on cool-to-warm tones was developed to shade surfaces with a limited dynamic range of colours that are visually distinct from colours of edge line. Anderson and Levoy [1] proposed an approach to enhancing visualization of cuneiform tablets using curvature and accessibility-based shading. In this approach, standard illumination was replaced by non-photorealistic rendering. A model of the tablets was first constructed using a high-resolution 3D range scanner. In the unwrapping stage, irregular meshes were partitioned into rectangular patches with four connected curves. Each of the rectangular patches was then fitted to a grid of springs by iteratively relaxing and subdividing the spring grid in order to create flat rectangles. Finally, curvature and accessibility-based shadings were applied to the unwrapped text. Vicinity shading for volumetric data, which was claimed to enhance the perception of surfaces within the volume, as proposed by Steward [28], has extended the idea of accessibility shading by incorporating uniform diffuse illumination, which arrives equally from all directions at each surface point in the volume, and uses occlusion by local occluders. In simple terms, each surface point is shaded on the basis of uniform diffuse lighting that is blocked only in the vicinity of the surface point. Vicinity shading was claimed to provide better perceptual cues than a conventional diffuse-plus-specular illumination model (e. g., Phong shading).

4 Perception-Based Design

In recent years, a number of approaches to computer graphics have been proposed that are based upon explicit models of a viewer's perception of graphical renderings. Viewer 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 [5; 23].

Perception-based lighting design has included implicit approaches that aim to maximize illumination entropy for a fixed viewpoint. Gumhold [9] describes a perceptual illumination entropy approach in which he uses limited user studies to model user preferences in relation to brightness and curvature. The maximum entropy-based approach has been widely used in a variety of fields such as communication theory computer vision, and recently in computer graphics. This approach assumes that all light sources have white colour and that hue and saturation are not important in the lighting design process that is, brightness is the only component used for calculating illumination entropy.

In this approach, colour values are first converted to brightness values ranging from 0 to 255. Brightness values are then mapped to 30 bins for the calculation of the probabilities that will be used for the illumination entropy computation. A fast lighting approach is also proposed based on the brightness dependent characteristic of the approach, while a hierarchical representation of different resolutions of the image is used for the calculations. The optimisation phase, using local and global schemes, tries to maximize the illumination entropy. In a simple experiment, they asked users to set up a directional light source for a fixed view point such that they could best perceive the three dimensional shapes. On the basis of the experiment results, a perceptual illumination entropy computation was proposed. To avoid regions that are too dark or too bright, the bins of brightness values were changed by choosing larger intervals for the bins of low and high brightness than for those of intermediate brightness (using a non-linear function for mapping brightness values to the bins).

Another perceptual feature, which was added to the entropy computation, was a measure of curvature importance. It was argued that regions with high curvature carry much more information than those with low curvature. With this assumption, each pixel was given an importance weight. An importance weight of a pixel was estimated on the basis of gradient calculation of the normal at that pixel from normals of adjacent pixels. Although the entropy-based approach is strongly motivated by information theory, it is less well grounded in theories of visual perception. Indeed, there is no proven relation between information calculated using information theory and the 3D information that the human visual system perceives and process. Many features crucial to the process of visual perception are not considered, such as shading gradient information, the average brightness of the image, and the brightness adaptation mechanism of the human visual system.

Lee et al. [20] proposed a more explicit model of perceptual preferences in Light Collages, in which lights were optimised such that the diffuse illumination is proportional to the local curvature and specular highlights were used only for regions of particularly high curvature. The key idea was to map a set of lights to a set of surface patches. Therefore, each light would light a subset of surface patches. In this approach, the three dimensional mesh was segmented into patches on the basis of surface curvature. Segmentation was performed by first calculating the average curvatures at each vertex of the mesh [29] and then applying the watershed method [18]. Sharp visual discontinuities across the patch boundaries were addressed by blending illumination from adjacent patches. Blending weights for a vertex were calculated on the basis of the distance from that vertex to the boundaries of the patches in which the vertex does not reside. These weights were then used to weight the illumination contribution of light sources to that vertex. Only white directional lights were used in this approach and the directions were computed by optimising a so-called light placement function. The objective function was based on two sub-functions, the specular weight function and the diffuse weight function, which formulated on the basis of the mean curvature at the vertices, the view direction, and the light directions. Lights were sequentially added to the scene in the direction that maximizes the objective function. In this approach, Lee et al. also explored the use of practical lighting techniques, including the use of backlighting technique to enhance the silhouette visibility of the objects in the scene. One of the drawbacks of Light Collages is that the intensities of lights were not included in the optimisation framework, and scenes of only one object could be lit.

Much research has been carried out aimed at narrowing the gap between the perception between real images and computer-generated images. Realism of a computer-generated image is not only a matter of physical correctness, but also of perceptual equivalence of the image to the scene that it is intended to represent. Achieving perceptual equivalence between a computer-generated image and the scene is significantly challenging, since many perception-based problems must be taken into consideration. In particular, how to quantify the perceptual quality of an image? Much research has been carried out to find efficient perception-based image quality metrics. Notably, Shacked and Lischinsky [25] proposed an approach to solve lighting design problems using a perceptual quality metric. At the heart of Shacked & Lischinsky's research is an objective function that is motivated by the human perception of 3D scenes. The first component taken into consideration in their approach is the magnitude of shading gradient of the rendered image. Psychological research shows that variations in shading provide important information about shape, depth and the impression of realism. The second component of Shacked & Lischinsky's objective function is the measure of edge contrast. It is well established that edges convey information about shape and they should be prominently displayed by the image. In the objective function the prominence of edges is determined by the ratio between the number of pixels detected by a pixel-based edge detection operator (applied to the rendered image) and the maximum number of edge pixels that could appear. Visible edges must therefore be determined in order to calculate this edge term.

The third component used in the objective function is the variance of the luminance pixels in an image. The variance of the pixel luminance is useful because of its importance in our brightness adaptation mechanism. In any situation, the human visual system (HVS) adapts to a particular light intensity, which is called the brightness adaptation level, and the HVS is most sensitive to intensities around this level. Research on adaptation levels proposes that the adaptation level can be estimated by the log of the luminance visible on the retina. The HVS is not sensitive to intensities at values far below this level. This means that we should not create an image with a large range of intensities, and the variance of the luminance of pixels should be minimized. Whilst this seems to conflict with our goal of maximising the shading gradient component, we can moderate the tension between the average shading gradient component and the variance of luminance component by the selection of a target value.

The mean of all pixel luminance in an image determines the overall brightness of an image, and our perception of images is negatively affected if the brightness is too high or too low, since overly dark or light images tend to weaken the effects of shading and may obscure features in the scene. To address this problem, empirically established target brightnesses were used. In addition, the histogram of the objected function prevents the production of large areas with uniform shading. Related to the mean luminance is the luminance histogram of an image, which is the distribution of pixels over different grey-scale levels. An equalized histogram was used as a target distribution and the difference between the actual histogram and this target equalized histogram was minimized (a technique commonly exploited in image processing). A final component of the objective function is the direction of a light source. This component is motivated by psychophysical studies, indicating that it is easier for the HVS to interpret 3D shape when illuminated from above. The direction of light sources were constrained to come from above with a particular target direction; angular deviations from this were penalized.

5 Towards Lighting-by-Example

Perception-based lighting approaches generally use local illumination information and pixel luminance statistics, but their main shortcoming is that users have to set up target values for the specific viewing parameters. Furthermore, important characteristics such as the contrast between objects, curvature of surfaces, and perceptually non-uniform characteristics of different colour systems (on which the metrics were based) are also not addressed. Recently, lighting-by-example approaches have been formulated that seek to combine the advantage of inverse design and interactivity and extend perception-based approaches. Ha & Olivier combined elements of all three approaches in a set of tools that could use the luminance of 3D [11] and even 2D [12] exemplars. For 3D exemplars [11], the appearance of objects in the source scene (the scene to be copied) was evaluated using an extended set of objective functions [11] based on [25]. These values were then used to set the target values for these functions in the stochastic optimisation of the lights in for a target scene (the scene to be lit). Viewer similarity studies demonstrated people's judgement of the proximity of the resulting illumination to the source scenes (as compared to differently lit examples) [13].

Using a perception-based approach with a 3D exemplar assumes that a credible mapping can be made between the different objects in the target and source scenes. Furthermore, there is no readily available set of appropriately lit 3D scenes to use as targets. As an alternative Ha & Olivier used photographs as exemplars, and instead of evaluating properties based on perceptual characteristics that be computed from a fully segmented model of the scene, they used a wavelet characterization of the photograph's luminance. By using optimising the light properties so as minimize the difference between the wavelet model of the source and target images this technique manages to replicate the spatial frequency distribution of the luminance if the source (2D) image in the target (3D image). Again the veracity of the approach was demonstrated in viewer judgement studies [13].

Although light-by-example would appear to address many of the shortcoming of previous approaches many open problems for automatic lighting design remain:

Dynamic Real-Time Lighting. In applications that include dynamic scenes with moving objects, a dynamic real-time lighting framework is highly desirable. In such applications, changes of scene structure from frame-to-frame results mean that lighting parameters could be re-optimised in order to express the intent and purpose of current scene structure. Current lighting optimisation processes are resource intensive and so we need to develop alternative real-time approaches, for example using caching and light mapping techniques, constraining the locations of lights in an appropriate manner.

Implementation of Multi-Light Optimization. In theory, the approaches developed, such as ideal perception-based lighting design and lighting-by-example, can support any number of lights. Due to our particular research focus being on the discovery of different lighting design algorithms (rather than performance and applications in their own right), our current lighting optimisation framework currently supports a maximum of eight lights for each scene, while each optimisation loop only supports two lights. In terms of practical application, it would be desirable to realise these techniques in the form of a plug-in for graphics tool such as 3D-Max and Light Wave.

Developing the Perceptual Metrics. Wavelet-based lighting-by-example is in one sense a step back from perceptually driven approaches, better computational models of viewers perception of graphics can only help to inform our understanding of lighting design. There is significant scope for the development of further perceptual metrics, both in terms of the low-level components of a perceptual model of lighting (i. e., for ideal lighting models) and also in more holistic metrics such the distance between source and target components in the objective function, which would enhance lighting-by-example directly.

Lighting for Non-Photorealistic Scenes. Nonphoto-realistic lighting is not constrained to physically correct lighting and shading and is instead application dependent. For example, supporting the communication of form and character means that many of the conventions of classical perception (incorporated in an objective function) would have to take second place to the stylised effects of lighting realised in cartoons and other non-photorealistic media. It is likely that we would need to develop an alternative ontology of lighting effects for such domains, but the underlying notions of lighting-by-example would still apply.

A Language for Lighting. The fact that lighting effects can expressively convey mood and emotion has inspired our lighting-by-example approach. Within different visual cultures, there is some consistency in the representation of moods/emotions using different lighting effects and lighting properties. The integration of emotion/mood descriptors into the lighting optimisation process has the potential to increase the flexibility of both our lighting-by-example and direct lighting approaches by enabling users to communicate explicit goals with the system. Emotion/mood descriptors might be used either as a criteria to search for an appropriate exemplar or directly as a target for the lighting optimisation.

Lighting-by-example is in fact a step away from perceptual fidelity, and although this may seem unusual direction for a field of Smart Graphics, as with all artistic endeavors the wider goal is not to produce a cognitively accurate (and empirically verified) objective function (for ideally illuminated scenes) but a framework for the specification of appropriately expressive tools (e. g., lighting using example 3D scenes and photographs). Such approaches do still presume the ability to model target scenes in the form of *meaningful* objective functions, and to optimise source scenes using these objectives. Lighting-by-example approaches actually refer back to users (the viewers of images) to both configure the parameters of our objective function and to verify the results of our approaches to lighting design, and perhaps provide the most likely avenue for fruitful research and beneficial tools. The credibility of lighting design research is dependent on whether a user’s judgement of the results generated demonstrates the improvements predicted. Though the results to date are not elegant integrations of methods and tools, the foundations of the declarative approaches required have been laid.

References

[1] S. Anderson and M. Levoy. Unwrapping and visualizing cuneiform tablets. In: *IEEE Computer Graphics and Applications* 22(6):82–88, 2002.

[2] M. G. Cindy, E. T. Kenneth, P. G. Donald, and B. Battaile. Modelling the interaction of light between diffuse surfaces. In: *Proc. of the 11th Annual Int’l Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 213–222, 1984.

[3] M. Christie, P. Olivier, and J.-M. Normand. Camera Control in Computer Graphics. In: *Computer Graphics Forum* 27(8):2197–2218, 2008.

[4] E. Desmontils. Expressing constraint satisfaction problems in declarative modeling using natural language and fuzzy sets. In: *Computers & Graphics* 24:555–568, 2000.

[5] J. P. Farugia and B. Peroche. A progressive rendering algorithm using an adaptive perceptually based image metric. In: *Computer Graphics Forum* 23(3):605–616, 2004.

[6] J. E. Flynn. A study of subjective responses to low energy and non-uniform lighting systems. In: *Lighting Design and Application* 7(1):167–179, 1977.

[7] J. E. Flynn, C. Hendrick, T. J. Spencer, and O. Martyniuk. A guide to methodology procedures for measuring subjective impressions in lighting. In: *Journal of the Illuminating Engineering Society* 8:95–110, 1979.

[8] A. Gooch, B. Gooch, P. Shirley, and E. Cohen. A non-photorealistic lighting model for automatic technical illustration. In: *Proc. of the 25th Annual Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 447–452, 1998.

[9] S. Gumhold. Maximum entropy light source placement. In: *Proc. of IEEE Visualization*, pp. 275–282, 2002.

[10] H.-N. Ha and P. Olivier. Explorations in declarative lighting design. In: *Proc. of the 6th Int’l Symp. on Smart Graphics*, Vancouver, Canada. LNCS 4073, Springer-Verlag, Berlin-Heidelberg, 2006.

[11] H.-N. Ha and P. Olivier. Lighting-by-Example with Wavelets. In: *Proc. of the 7th Int’l Symp. on Smart Graphics*, Kyoto, Japan. LNCS 4569, pp. 110–123, Springer-Verlag, Berlin-Heidelberg, 2007.

[12] H.-N. Ha and P. Olivier. Perception-Based Lighting-by-Example. In: *Theory and Practice of Computer Graphics*, pp. 61–68, University of Bangor, UK, 2007.

[13] H.-N. Ha. *Automatic Lighting Design*. Ph.D. Thesis, School of Computing Science, Newcastle University, UK, 2008.

[14] P. Hanrahan and P. Haeberli. Direct WYSIWYG painting and texturing on 3D shapes. In: *Proc. of the 17th Annual Int’l Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 215–223, 1990.

[15] V. Jolivet, D. Plemenos, and P. I. Poulingeas. Inverse direct lighting with a Monte Carlo method and declarative modelling. In: *Proc of the Int’l Conf. on Computational Science (ICCS)*, pp. 3–12, 2002.

[16] G. Karypis and V. Kumar. Multilevel k-way partitioning scheme for irregular graphs. In: *Journal of Parallel and Distributed Computing* 48:96–129, 1995.

[17] J. K. Kawai, J. S. Painter, and M. F. Cohen. Radioptimization-goal based rendering. In: *Proc. of the 20th Annual Int’l Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 147–154, 1993.

[18] A. P. Mangan and R. T. Whitaker. Partitioning 3d surface meshes using watershed segmentation. In: *IEEE Transactions on Visualization and Computer Graphics* 5(4):308–321, 1999.

[19] J. Marks, B. Andalman, P. A. Beardsley, W. Freeman, S. Gibson, J. 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: *Proc. of the 24th Annual Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 389–400, 1997.

[20] C. H. Lee, X. Hao, and A. Varshney. Light Collages: Lighting design for effective visualization. In: *Proc. of IEEE Visualization*, pp. 281–288, 2004.

[21] P. Poulin and A. Fournier. Lights from highlights and shadows. In: *Proc. of the Symp. on Interactive 3D Graphics*, pp. 31–38, 1992.

[22] P. Poulin, K. Ratib, and M. Jacques. Sketching shadows and highlights to position lights. In: *Proc. of the Conf. on Computer Graphics International (CGI)*, pp. 56–63, 1997.

[23] P. S. A. Reitsma and N. S. Pollard. Perceptual metrics for character animation: sensitivity to errors in ballistic motion. In: *ACM Transactions on Graphics* 22(3):537–542, 2003.

[24] C. Schoeneman, J. Dorsey, B. Smits, J. Arvo, and D. Greenburg. Painting with Light. In: *Proc. of the 20th Annual Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 143–146, 1993.

[25] R. Shacked and D. Lischinski. Automatic Lighting Design using a perceptual quality metric. In: *Computer Graphics Forum* 20(3):215–226, 2001.

[26] R. L. Shirley. An efficient algorithm for determining the convex hull of a finite set of points in the plane. In: *Information Processing Letters* 1:132–133, 1972.

[27] P. Shirley. Radiosity via ray tracing. In: *Graphics Gems II*, J. Arvo (Ed.), pp. 306–310, Academic Press, San Diego 1991.

[28] A. J. Stewart. Vicinity shading for enhanced perception of volumetric data. In: *Proc. of IEEE Visualization*, pp. 355–362, 2003.

[29] G. Taubin. Estimating the tensor of curvature of a surface from a polyhedral approximation. In: *Proc. of the 5th International Conference on Computer Vision (ICCV)*, pp. 902–907, 1995.

[30] T. Whitted. An improved illumination model for shaded display. In: *Communications of the ACM* 23(6):343–349, 1980.

Received: December 1, 2008

Dr. Patrick Olivier is a Reader in Human-Computer Interaction in the School of Computing Science, Newcastle University. Based in Culture Lab, Newcastle's interdisciplinary centre for digital practice, Dr. Olivier is exploring a number of key elements of the smart graphics research agenda, including multimodal and cross-modal interaction, computational cinematography and the use of creative practice in interface design.

Address: Newcastle University, NE1 7RU Newcastle Upon Tyne, United Kingdom,
e-mail: p.l.olivier@ncl.ac.uk

Dr. Hai NAM Ha is an Assistant Professor at the Post and Telecoms Institute of Technology in Hanoi, Vietnam. Dr. Ha studied for a PhD in Computer Science at Newcastle University where he explored the application of optimisation techniques to the declarative design of lighting in computer graphics applications.

Address: PTIT, Nguyen Trai Road KM10, Hanoi, Vietnam, e-mail: nam.hahai@googlemail.com

Dr. Marc Christie is a visiting researcher in the Bunraku team at INRIA Rennes. Since 2003, he has been an Assistant Professor at Nantes University. His research is located at the intersection of constraint solving techniques and computer graphics. Many complex problems in modelling and animation can be expressed and solved efficiently by interval constraint techniques, in particular in the domain of virtual camera control. He has been publishing contributions in this field in both the constraints and graphics communities.

Address: LINA, Laboratoire d'Informatique de Nantes Atlantique Université de Nantes 2, rue de la Houssinière BP 92208, 44322 Nantes Cedex 03, France, e-mail: marc.christie@irisa.fr