Automating Lighting Design for Interactive Entertainment

Magy Seif El-Nasr
Information Science and Technology
Pennsylvania State University
magy@ist.psu.edu

Ian Hosrwill
Computer Science Department
Northwestern University
ian@cs.northwestern.edu

ABSTRACT

Recent advances in computer graphics, particularly in real-time rendering, have caused major improvements in 3D graphics and rendering techniques used in interactive entertainment. In this paper, we focus on the scene lighting process, which we define as the process of configuring the number of lights used in a scene, their properties (e.g. range and attenuation), positions, angles, and colors. Lighting design is well known among designers, directors, and visual artists for its vital role in influencing viewers' perception by evoking moods, directing their gaze to important areas (i.e. providing visual focus), and conveying visual tension. It is, however, difficult to set positions, angles, or colors for lights within interactive scenes to accommodate these design goals, because an interactive scene's spatial and dramatic configuration, including mood, dramatic intensity, and the relative importance of different characters, change unpredictably in real-time. There are several techniques developed by the game industry that establish spectacular real-time lighting effects within 3-D interactive environments. These techniques are often time and labor intensive. In addition, they are not easily used to dynamically mold the visual design to convey communicative, dramatic, and aesthetic functions as addressed in creative disciplines, such as art, film, and theatre. In this paper, we present a new real-time lighting design model based on cinematic and theatric lighting design theory. The proposed model is designed to automatically, and in real-time, adjust lighting in an interactive scene accommodating the dramatic, aesthetic, and communicative functions described by traditional lighting design theories while accommodating artistic constraints concerning style, visual continuity, and aesthetic function.

1. INTRODUCTION

Filmmakers and animators compose visual images that support and shape the narrative and the dramatic action presented (Knopf 1979; Crowther 1989; Lowell 1992; Alton 1995; Calahan 1996; Gillette 1998; Campbell 1999; Birn 2000; Block 2001; Kidd 2001). Similarly, interactive entertainment should produce visual images that adapt to the narrative content and the dramatic action. However, adapting the visual presentation of an interactive scene to accommodate variations in the narrative and the action is an intractable and daunting problem.

Lighting design is a difficult problem. The role of a lighting designer is to establish a lighting design that serves several goals, enumerated as follows:

Dramatic Goals:

- Adequately light all characters in a shot:
    - for visibility of character's face to show reactions, emotional expressions

    - for visibility of character’s body to emphasize actions and gestures
- Portraying visual tension: use colors to create contrast that parallels tension in the scene
- Provide visual focus to guide the viewer to the dramatic focus of the scene

Lyrical Goals:

- Provide mood

Aesthetic Goals:

- Conform to the lighting style chosen for the piece
- Establish perception of depth
- Establish character modeling

Realistic Goals:

- Establish logical motivation for the direction of light. For example, if a window is present in the scene then the scene has to be lit using the window as the motivational source of light
- Establish a sense of visual continuity between frames

These goals conflict with each other. For instance, if a lighting designer adjusts the lighting to serve realistic goals, dramatic and lyrical goals may be sacrificed. Lighting design, as any design problem, involves trade-offs. Lighting designers favor some goals over others depending on the lighting style of the production and the dramatic situation. A lighting designer, for example, may choose a realistic style, and thus favors realistic goals over dramatic, aesthetic, and lyrical goals.

In addition, interactivity adds another challenge to the problem. In interactive scenes, design parameters, such as character placement and dramatic intensity, cannot be anticipated at design time. The participant is free to move about in the scene, which affects camera angle and position. In addition, characters’ locations and orientations change depending on the participant’s position and orientation, and thus cannot be determined in advance. Furthermore, the participant’s choices and actions affect his/her relationship with other characters, thereby affecting the story, other characters’ behaviors (including orientation and position), and the dramatic intensity of the scene.

It is important to note that we are assuming an adaptive narrative design. While some games and other interactive entertainment productions rely on a linear narrative (e.g. Devil May Cry, Legacy of Kain), others vary the plot and action depending on the interaction (e.g. Silent Hill, Boulders Gate). To design a lighting design model that can accommodate both applications, we are assuming a non-linear narrative, where events, relationships between characters, and the importance of characters change in real-time responding to interaction.

In such an unpredictable environment, it is very hard to adjust the lighting while ensuring that the design adheres to the lighting design communicative and aesthetic goals described by cinematic and theatric designers. Game designers have developed several techniques to produce stunning and spectacular lighting effects. The demo shown for Doom 3, for example (available at doom3.com) shows various impressive lighting effects. However, these effects are still limited in their communicative and aesthetic functions and are typically limited to the underlying narrative structure. For example, assuming a linear narrative, designers can outline lighting changes at design-time to better evoke moods and emotions, which is very difficult to accomplish given an adaptive narrative, where dramatic importance of events is only known at run-time. This is not only due to the unpredictability of the narrative, but also due to the unpredictability of the lighting conditions in a particular situation.

In this paper, we propose a new lighting design model, called ELE (Expressive Lighting Engine), for interactive entertainment. ELE is designed to automatically adjust the lighting in real-time accommodating the situation and satisfying the lighting design goals as documented by film and theatre designers. To guide this process, we have adapted lighting design principles from film and theatre and represented them mathematically within ELE. Recognizing lighting design as an optimization process, ELE uses constraint-based optimization to balance the lighting design goals, whose priorities are set depending on the situation and lighting style chosen for the scene.

ELE also allows artists to override its decisions using a set of high-level constraints, e.g. importance of visibility, importance of conveying tension, and color constraints. Using these parameters, designers can set constraints that specify the perceptual properties and the style of the desired lighting, thus expediting and facilitating the lighting design process.

2. RELATED WORK

Lighting design is a very important element of an interactive 3-D production. The game industry has developed many techniques based on direct adoption of cinematic and theatric design conventions. However, to our knowledge, the adaptation of theories from creative media, such as art, theatre, or film is still at a very early stage. Thus, a theory or model of lighting design for interactive entertainment that adapts creative design elements from film and theatre is needed. In this section, we will review some of the current techniques for lighting design used by the industry and research community.

2.1 Ambient Lighting Design

Ambient lighting is a lighting method where the objects are given constant luminance values (Moller 1999). This is a fast and simple model of lighting in which all objects are equally visible. This type of lighting has been used in interactive entertainment productions, such as the *Sims*, and *Sim City*. Although the technique supplied the desired look and feel for games, like the *Sims* and *Sim City*, it cannot be generalized for use in first person shooters or action games, where a realistic and dramatic lighting style is often more appropriate.

Figure 1. A screenshot from Max Payne

2.2 Realistic Lighting Design

This lighting technique borrows from realistic lighting, whereby lighting is designed to achieve realistic goals, including showing realistic effects, such as shadows cast on character's face to show effect of a light source direction (such as the sun or torch), the reflectance of character's shadows on scene geometry, and more accurate lighting calculations. Designers most often use non-interactive rendering algorithms to generate light maps that provide the lighting required for the scene (Carson 2000; Maattaa 2002). These techniques have been used in many games; examples include Max Payne, Doom 3, and Splinters Cell. Figure 1 shows a screenshot of Max Payne using this lighting technique.

As described earlier realistic goals are only some of the many important goals of a lighting design. Other important lighting design goals include ensuring visual focus, providing visibility, and paralleling the dramatic tension of the plot/interaction (Knopf 1979; Lowell 1992; Block 2001; Kidd 2001). To achieve these goals designers often employ several tricks, such as halos around objects to direct users to important objects and spot lights that follow characters around to ensure their visibility. These techniques are limited by the underlying narrative structure. For example, if we allow the importance of objects to change unpredictably during interaction, then a halo that suddenly appears around an object may be unrealistic and/or distracting. Theatre and film apply very subtle techniques to achieve these goals. However, such techniques rely on a model of lighting design that coordinates the properties of each light in a scene and integrates their function.

One advantage of using a model based on theatric and cinematic theory is the ability to automatically adapt to the continuous variation in dramatic tension and action. By examining games, such as *Devil May Cry*, it is apparent that the dramatic tension was broken into discrete segments or missions, where an appropriate level is loaded. In some cases, the difference between the levels is only in texture or lighting colors. For example, in the last level in *Devil May Cry* was colored with a distinct saturated red color signifying the climax. There are several problems with this method. First, it is very tedious to redesign and relight each level. Second, the design involves breaking the continuous flow of tension and manually adjusting the textures or lighting to accommodate the tension increase or decrease. This is a cumbersome task that can be avoided by using an appropriate lighting design model.

2.3 Lighting for Emotions

Tomlinson developed a system that changes light colors and camera movements to present the user with an interpretation of the world based on the characters' emotions within the scene (Thompson 2001). He used film grammar to select camera movements and lighting colors that show characters' characteristics or feelings. For example, he used low camera angles to show that a character is powerful, or harsh red light to make a character look demonic.

He categorized lights as global lights (lights lighting the architecture and the environment) and personal lights. Global lights have a default scheme. They are fixed, and are mainly used to provide the key source of illumination, and maintain a basic visual continuity for the scene. Personal lights are fixed on characters. They follow characters around; their colors depend on the emotional state of the characters lit.

Film lighting designers often adjust positions and colors of all lights in the scene, including what Tomlinson refers to as global lights, to accommodate for the camera angle, movement, dramatic intensity, and mood. Tomlinson fixed these lights and changed colors of only personal lights, which restricts the variety of moods that can be produced by the lighting design.

In addition, lighting design, especially choices of color, is most often used to direct viewer's attention to the important characters/objects, convey mood, and provide visual intensity (Millerson 1991; Foss 1992; Lowell 1992; Viera 1993; Block 2001). In Tomlinson's work, lighting is restricted to portraying the emotional states of each character rather than providing a coherent mood or visual focus.

3. ELE – EXPRESSIVE LIGHTING ENGINE

ELE is a lighting design model that is based on visual design theories from theatre and film. ELE is designed to automatically select the number of lights used, their positions, colors, and angles. To accomplish this task, ELE uses lighting design rules represented mathematically in an optimization function. The use of optimization is important to balance the conflicting lighting design goals. While adapting the lighting to the interaction, ELE has to maintain visual continuity, and style. In this section, I will introduce ELE and explain its ability to establish the communicative, dramatic, and aesthetic functions described by film and theatre lighting designers.

I assume that there exists a system that passes several parameters to ELE, including a set of parameters describing style, local light sources, stage configuration and dimensions, characters' dimensions, dramatic focus (the area/characters towards which attention should be directed) and the dramatic intensity of the situation. Using these parameters ELE computes the number of lights used. For each of these lights, it computes the type of instrument (e.g. spot light or point light), color in RGB color space, attenuation, position as a 3D point, orientation including the facing and up vectors, range, masking parameters, and, depending on the light instrument used, the Penumbra and Umbra angles. These parameters are given to a rendering engine to render the frame.

ELE first determines where to direct viewers' attention given the number of characters in the frame and the dramatic importance of their action. I use the term dramatic focus to denote the area where attention should be directed. ELE then dynamically allocates lights to visible areas in the scene. Once lights are allocated to areas, ELE selects angles and colors for each light in the scene, thus forming a light setup. The light setup is then given to the rendering engine to render the frame.

3.1 Selecting Dramatic focus

To achieve visual focus, ELE is designed to differentiate between focus, non-focus, and background areas. A focus area (dramatic focus of the scene) is a character, a group of characters, or an object. ELE selects the dramatic focus as follows: a character/object *c* is the dramatic focus, if:

- The camera is in a close-up, medium close-up, medium, or full shot on *c*
- The only character/object in view is *c*
- Character/object *c* has the most dramatic action (i.e. the action that has the most impact on the plot). Authors can write rules that rate actions on a scale from one to ten (ten being very important action). For example, a running action will be judged as more dramatic than breathing or walking. In addition, ELE uses built-in common sense rules, such as talking is more dramatic than listening.

It is worth noting that the contrast between focus and non-focus areas is the main method used to communicate tension and attention. Artists create visual tension by varying color contrast between lights lighting the focus and non-focus areas. For example, artists increase contrast by increasing warmth or brightness of the lights lighting the dramatic focus, or decrease the brightness or warmth of the lights lighting non-focus and background areas.

3.2 Dynamic Light Allocation

Dynamic lights are a scarce resource and need to be allocated and managed efficiently to comply with the lighting design requirements discussed above while achieving real-time rendering speed. Rendering time is proportional to the number of lights used. Thus, to achieve real-time rendering most rendering engines limit the number of dynamic lights used in a scene, e.g. Wildtangent (a publicly available rendering engine) restricts the number of dynamic lights to eight lights. On the other hand, lighting designers use many lights to gain finer control on the different areas shown in the scene and provide modeling and depth. Lighting designers at Pixar, for example, use eight lights or more to light one character and thirty two lights or more to light a complete scene.

This problem is not new to game design; game designers often have to accommodate different methods for handling dynamic resource allocation including CPU power, number of lights, and audio effects. Through conversations with game designer we established that game designers use scripted rules to shift resources, e.g. dropping background music or sound for emphasizing a dialogue line or an explosion. Dynamically accommodating lights in an interactive scene, however, is a harder problem due to the number of constraints involved, including visual continuity, providing visual focus, and visibility for the action.

To tackle this problem, ELE uses an optimization system to balance constraints and goals. ELE allocates lights only to visible areas, since allocating lights to non-visible areas is a waste of resources. For each shot or camera movement, ELE reallocates the lights. This may create performance problems, however, if the camera is panning or tilting. Given the anticipated camera movements determined by the story engine[1], ELE determines the probability of next camera movement being a pan or a tilt assuming that all camera movements anticipated have equal probability of being fired. It then calculates the visible area accordingly.

To direct viewer's attention to the dramatic focus of the scene, and to obtain finer control on the angles of light on characters' faces, ELE divides the visible area into several areas depending on the maximum number of lights that can be used, the number of characters in the scene, and the dramatic focus computed. It then allocates lights to each area depending on the level of visibility, modeling, and depth needed.

In summary, ELE follows the steps listed below to dynamically allocate lights in a scene:

1. Calculate visible area

   If next likely camera shots (given by story engine) include a pan or a tilt, then increase visible area by some factor ρ which is calculated depending on the anticipated camera speed.

2. Divide visible area into several areas differentiating between focus, non-focus, character, foreground, and background areas.
3. Allocate lights dynamically depending on the lighting design goals and their importance.

3.2.1 Dividing Visible Region to Areas

Given the visible region, ELE creates a number of areas A to cover the background, the foreground, and the characters within the visible region. ELE divides the stage into a number of overlapping areas (called acting areas), where the overlap region is set to a constant *o* (Gillette 1998). ELE then uses a greedy

---

[1] The story engine is a component of an interactive narrative architecture. The story engine determines the next possible camera movements by computing all the possible story events that could fire next, given the story situation and the possible participant's actions.

algorithm to create areas for characters (called character areas), such that all characters are assigned to an area, as follows:

```
Step 1. For each character $c$ create a new area and assign $c$ to it
Step 2. Repeat
            For each area $a$
                if $\exists\ a'$ s.t. $\left|a - a'\right| < \varepsilon$, and both are focus areas (or non-focus)
                    then merge $a$, $a'$
```

Each area *a* is lit within a cylinder cyl(*a*) with center, radius, and height given by:

$$\text{center}(\text{cyl}(a)) = \text{center}(\text{bbox}(a)), \tag{1.1}$$

$$r(\text{cyl}(a)) = \left\|\text{bbox}(a)\right\|_\infty + \max_{\text{character } c \in a} \left\|c\right\|_\infty + s, \tag{1.2}$$

$$h(\text{cyl}(a)) = \max_{\text{object } x \in \text{bbox}(a)} h(x), \tag{1.3}$$

where bbox(*a*) is the bounding box of all characters in area *a*, h(*y*) is the height of some object *y*, and *s* ("slop") is a constant. The notation $\|\mathrm{y}\|_\infty$ is used to denote the maximum dimension of object *y*.

3.2.2 Allocating Lights to Areas

ELE sets a maximum limit to the number of lights that can be assigned to each area. Non-character areas are assigned a maximum of one light. Visible character areas, however, are assigned a maximum of five lights, because character areas may require finer control to establish depth and modeling.

Spot lights are used for character and acting areas. On the other hand, the type of light used to light a background area depends on the practical sources present. For example, point lights are used to simulate light emitted by torches or candles, while spot lights or directional lights are used to simulate the effect of sunlight projected from a window or a door.

Many parameters affect the allocation of lights. ELE allocates lights according to cost parameters associated with visibility, modeling, depth, and visual continuity. We define a light allocation $p : L \to A$ to be an assignment of lights to areas. Note that not all areas will be assigned a light, i.e. *p* may not be onto. We, therefore, define the visibility, V(*p*), of a light allocation to be the percentage of visible areas that are assigned lights by *p*, or:

$$V(p) = \frac{\left|\left\{a \in A \,\middle|\, p^{-1}(a) \neq \varnothing\right\}\right|}{\left|A\right|}, \tag{1.4}$$

We define modeling as the average number of lights assigned to character areas, or:

$$M(p) = \frac{\sum\limits_{a \in A_{character}} \left|p^{-1}(a)\right|}{\left|A_{character}\right|}, \tag{1.5}$$

The depth, D(*p*), of a light allocation *p* is the difference between the number of lights assigned to the background and foreground areas, or:

$$D(p) = \sum_{a \in A_{background}} \left| p^{-1}(a) \right| - \sum_{a \notin A_{background}} \left| p^{-1}(a) \right|. \qquad (1.6)$$

The visual continuity, VC(*p*), of a light allocation *p* is defined as the difference between the configuration being evaluated and the one used in the previous frame:

$$VC(p_t, p_{t-1}) = \frac{1}{|L|} \sum_{a \in A} \left\| p_t^{-1}(a) \right| - \left| p_{t-1}^{-1}(a) \right\|. \qquad (1.7)$$

Hence, given the local light sources, the stage configuration and dimensions, the importance of various lighting goals calculated given the situation using cinematic rules (Campbell 1999), and the formulae defined above, ELE uses a multi-objective function, which is a weighted sum of these formulae, where weights correspond to the importance of lighting design goals, as follows:

$$p_{opt} = \arg\max_{p} \left(\lambda_v V(p) + \lambda_d D(p) + \lambda_m M(p) + \lambda_{vc} VC(p) \right), \qquad (1.8)$$

where λ_v is the importance of visibility, λ_d is the importance of depth, λ_m is the importance of modeling, λ_{vc} is the importance of visual continuity.

We formulated a greedy algorithm that allocates lights to each visible area in the scene, as follows:

1. Each area is assigned the maximum number of lights it can have
2. Remove one light that will incur the smallest loss
3. Repeat step 2 until the number of lights assigned is less than or equal to the maximum

3.3. Selecting Angles

ELE selects an angle for each light in the scene. In this section, we will discuss the method by which ELE selects angles for character areas including key, fill, and backlight angles. The same techniques are used to calculate angles for other areas, and thus will not be repeated.

The system selects an angle for each key[2] light according to several requirements, which include ensuring visual continuity, maintaining the illusion of a practical source, providing mood, and ensuring that all characters are visible.

Cinematic rules used to satisfy these requirements often contradict with one another. Angles used to establish mood, for example, don't usually produce good visibility, e.g. rim or silhouette angles. Thus, we softened these rules into cost functions, where the contradicting requirements, such as mood, visibility, and modeling are controlled using weights that are automatically calculated by ELE given the dramatic situation, or can be overridden by artists depending on style desired. ELE then uses optimization to find the best solution (i.e. solution with minimum cost).

ELE uses the following parameters:

- Cost associated with deviation from angle used in previous frame
- Cost associated with motivation

---

[2] the main source of light that establishes direction and shadows

- Cost of visibility
- Cost associated with mood
- Ideal mood angle (the ideal mood angle; e.g. side-light)

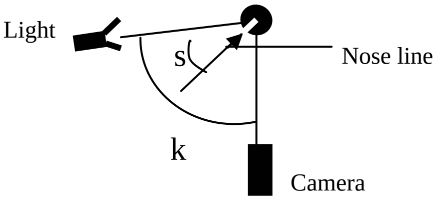

Figure 2. Angles between subject, camera, and light

These parameters are calculated using cinematic rules depending on character relationships, camera shot, and dramatic intensity. For example, if the camera is showing a closeup of character x and character x has negative relationship with the user, the ideal mood angle will be set to underlight. Importance of visibility, motivation, or visual continuity is modulated depending on the style adopted for the scene. For example, if a realistic style is desired, then high importance is given to visual continuity and motivation; while if a dramatic style is desired, then high importance is given to visibility. On the other hand, if a lyrical style is desired, high importance is given to mood. These parameters can also be overridden by artists, if they so desire.

The parameters are used as weights in the following cost function:

$$\mathrm{cost}(k, s, k^{-}, m) = \lambda_v (1 - V(k, s)) + \lambda_{-} \left| k - k^{-} \right| + \lambda_m \left| k - m \right| + \lambda_l \min_i \left| k - l_i \right|, \qquad (1.9)$$

where k and s are defined as the key light azimuth angle relative to the camera and the subject angle relative to the key light, respectively, as shown in figure 3, k^{-} is the key light azimuth angle from the previous frame, λ_{-} is the cost of changing the key light angle over time (to enforce visual continuity), λ_m is the cost of deviation from the mood azimuth angle, m is the mood azimuth angle suggested by the artist, λ_l is the cost of azimuth angle deviation from a practical source direction, l_i is the azimuth angle of light emitted by the practical source i, and λ_v is the cost of deviation from an orientation of light that establishes best visibility.

Based on Millerson's documented rules (Millerson 1991), we formulated the following equation to evaluate the visibility and modeling of a given key light azimuth angle:

$$V(k, s) = \sin(k)\cos(s). \qquad (1.10)$$

Millerson recommended an elevation angle between $\pi/6$ and $\pi/3$.

ELE uses a non-linear optimization system based on hill climbing to select an angle for each key light that minimizes the cost function above.

ELE uses rules based on Millerson's (Millerson 1991) guidelines to select fill and backlight azimuth angles depending on the value of the key light angle. According to Millerson's guidelines (Millerson 1991), fill light azimuth and elevation angles are calculated to be the mirror image of the key light angle. We define backlight azimuth angle as:

$$b = (k + \pi) \bmod 2\pi. \qquad (1.11)$$

3.4 Color

Colors of lights in a scene compose the contrast and feeling of the entire image. Game designers often use contrast to create mood; example of such effects can be seen in many games including *Silent Hill*, *Resident Evil*, and *Splinters Cell*. In these games, however, the color parameters are set by a level designer at the beginning of the scene.

Our goal is to use ELE to automatically adapt the lighting to the interaction, while allowing designers to write high-level rules setting constraints, including color constraints and lighting style desired. This is beneficial to designers, not only because it provides an easier, faster, and more design or art-oriented approach to lighting, but it also represents a more appropriate design method especially assuming non-linear or unpredictable environments, where narrative details, lighting conditions, and camera properties are not known and could be predicted at design time.

Designing such an automatic adaptive color subsystem is difficult, because a change of one light's color may affect the entire image. Such change affects not only the visual tension portrayed, but visual continuity, visual focus, and visibility goals as well. Therefore, ELE manipulates the colors by using an optimization algorithm that searches for the best color that achieves the desired effect on the entire image projected and conforms to artist's style.

ELE evaluates color using several parameters, including visual tension level and style, importance of portraying visual focus, emphasis on depth, importance of visual continuity, color palette constraints, as well as ideal values for the saturation, lightness, warmth/coolness for each light, and costs describing the importance of adhering to these values. Using default rules formulated based on film and theatre lighting theory, ELE determines default values for these parameters given the style, dramatic intensity, and dramatic focus. These parameters can also be overridden by artists, if they so desire.

Using the ideal values and their associated costs, ELE uses nonlinear optimization to search through a nine-dimension space of RGB values differentiating between colors of focus, non-focus, and background areas to select a color for each individual light in the scene evaluating this color using a multi-objective cost function, where each objective evaluates the color against the lighting design goals, including establishing depth, conforming to color style and constraints, paralleling dramatic tension, adhering to desired hue, saturation, and lightness, and maintaining visual continuity. The equation is defined as follows:

$$\begin{aligned} cost(c^t, c^{t-1}) = {} & \lambda_d \left(D(c^t) - d\right)^2 + \lambda_c \left(\text{contrast}_\phi(c^t) - \delta\right)^2 + v(x) + \\ & \sum_{i \in \{f,n,b\}} P(c_i^t, c_i^{t-1}), \end{aligned} \qquad (1.12)$$

$$\text{where} \quad p(c_i^t, c_i^{t-1}) = \begin{aligned} & \lambda_{s_i} \left(S(c_i^t) - s_i\right)^2 + \lambda_{h_i} \left(H(c_i^t) - h_i\right)^2 + \\ & \lambda_{l_i} \left(L(c_i^t) - l_i\right)^2 + \lambda_{w_i} \left(W(c_i^t) - w_i\right)^2 + \lambda_{ch} E(c_i^t, c_i^{t-1}), \end{aligned}$$

and where c^t is a vector of light colors for focus, *f*, non-focus, *n*, and background, *b*, areas at frame *t*. Color c_i^t is represented in RGB color space, S(*c*) denotes the saturation of color *c*, H(*c*) denotes the hue of color *c*, L(*c*) denotes lightness of color *c* (in RGB color space).

ELE uses CIEDE2000 – a well known formula for measuring color difference (Hill 1997; Luo 2000) as follows:

$$E = \sqrt{\left(\frac{\Delta L}{k_L S_L}\right)^2 + \left(\frac{\Delta C}{k_C S_C}\right)^2 + \left(\frac{\Delta H}{k_H S_H}\right)^2 + \Delta R}, \qquad (1.13)$$

where $\Delta R = R_T f(\Delta C \Delta H)$, and ΔL, ΔC, ΔH are CIELAB metric lightness, chroma, and hue differences respectively, S_L, S_C, S_H are weighting functions for the lightness, chroma and hue components, and k_L, k_C, k_H are parameters to be adjusted depending on model material information.

The depth, *D*(*c*), of a color vector, *c*, is defined as the color difference between colors of lights lighting the background areas and those lighting other areas, formulated as follows:

$$D(c) = \sum_{b \in B} \sum_{n \in NB} E(c_b, c_n), \qquad (1.14)$$

where *B* are the indices of the background lights and *NB* are the indices of the non-background lights, and where E is color difference defined above.

Following traditional lighting design theory, contrast is used to establish visual tension paralleling the dramatic tension in the scene (Block 2001). Using the guidelines documented by filmmakers and designers, contrast is defined relative to the dramatic focus of the scene (Block 2001) as the difference between colors of lights lighting the focus area and a weighted sum of the colors lighting the other areas as follows:

$$\text{contrast}_{\phi}(c) = \sum_{i \neq focus} w_i \left|\phi(c_{focus}) - \phi(c_i)\right|, \qquad (1.15)$$

where ϕ is the color component (lightness, warmth, or saturation) over which we're computing contrast, which largely depend on the style of the scene, *c* is a vector of the light colors, c_i is a color for an area type *i*, where $i \in \{focus, non-focus, background\}$ and *focus* is the index of the dramatic focus area.

Based on the results collected by Wooten and described in (Katra 1995), we used a multiple linear regression method to formulate color warmth in RGB color space, as follows:

$$warmth\left(\begin{bmatrix} R \\ G \\ B \end{bmatrix}\right) = \begin{bmatrix} 0.008 \\ 0.0006 \\ -0.0105 \end{bmatrix}^T \begin{bmatrix} R \\ G \\ B \end{bmatrix} - 0.422. \qquad (1.16)$$

The optimization problem discussed above is a constraint-based optimization problem, where the color, *c*, is constrained to a specific space of values defined by style (e.g. realistic style restricts the values of saturation or hue). ELE uses a boundary method to bind the feasible solutions using a barrier function *v*(*x*), such that $v(x) \to \infty$ as x approaches the boundary defined by the feasibility region. ELE uses the following formula for *v*(*x*):

$$\varepsilon \sum_{j}^{\rho} \log(-g_j(x)). \qquad (1.17)$$

Although gradient descent has major drawbacks, including occurrence of oscillations and being easily stuck in a local minimum, ELE uses gradient descent for several reasons. First, it provides a fast and simple solution. Second, a local minimum in this case is preferable because it provides a solution closer to the older one, thus ensuring visual continuity. Third, alternative methods rely on the existence of a second derivative, which is not necessarily true in this case.

3.5 Maintaining Visual Continuity and Style

ELE adapts the lighting automatically to the scene's dramatic action and tension maintaining the established style and visual continuity and accommodating the lighting design goals, including visual attention, visual tension, and depth. Even though the cost functions above include a term for visual continuity, sometimes suggested changes need to occur and can be very distracting. ELE uses camera cuts to hide such distracting lighting changes. Neuroscience and psychology literature have found that human vision does not perceive changes made between camera cuts (Hollingworth in press). If ELE selects a lighting configuration with a high cost (i.e. higher than a specific threshold) and if the dominating factor is the visual continuity cost, then ELE sends a message to the camera system asking for a cut. ELE interacts with the camera to synchronize the changes within the cuts.

ELE also ensures that the lighting accommodates the scene's dramatic development as well as the established style. ELE uses rules to manipulate the parameters used by the cost functions to ensure that established style is adequately maintained. ELE keeps track of established style by maintaining a style state that is represented in terms of three styles: realistic, dramatic, and expressionistic, and where these styles are represented as a number 0-100. ELE uses rules to manipulate these values as the colors change or as authored rules are fired. It, thus, adapts its style state to the changes that occur in the scene. Depending on this style state, ELE manipulates the lighting design cost parameters, including contrast, color warmth, and mood angles.

4. RESULTS

ELE has been implemented and tested in five interactive scenes from *Mirage*. In this section, we will review some results comparing ELE's performance to other techniques used in interactive entertainment. It is important to note that there are several issues that readers should be aware of while evaluating ELE using images. First, ELE is an optimization system that balances off many goals including visual continuity, and thus it is better demonstrated through a movie instead of static images. Therefore, readers are referred to the movie accompanying this paper for a better representation of ELE at work. Second, lighting itself is hard to show, because good lighting should be subtle and unnoticeable. The video accompanying the paper shows several styles of lighting; these styles were intentionally chosen to be expressive and exaggerated to show the effect of lighting and its role in supporting and emphasizing the narrative. In normal situations, especially with a realistic lighting style, these effects will not be noticeable.

As stated above, the technique of dynamic light allocation has several advantages over static techniques. First, it can dynamically adapt to unpredictable changes in camera angle or character actions, thus providing better visibility. Second, it can dynamically accommodate variations in the dramatic focus of the scene.

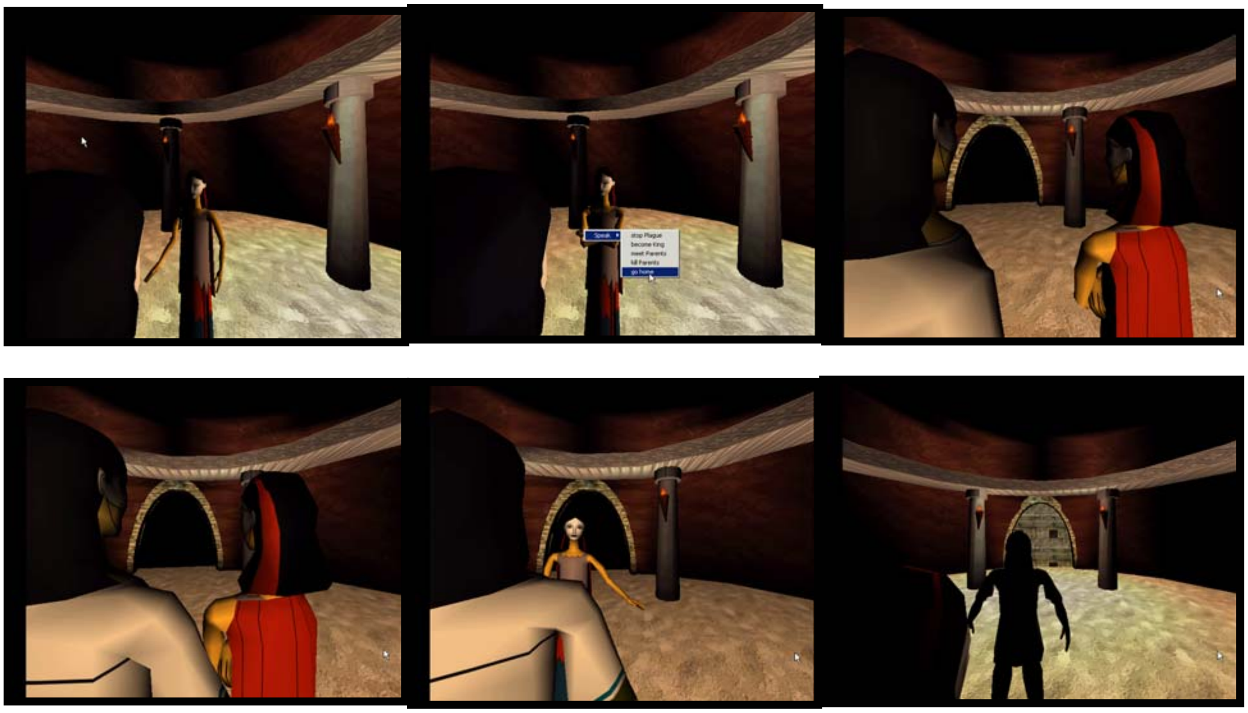

Figure 3. Sequence from scene 8 using static lighting

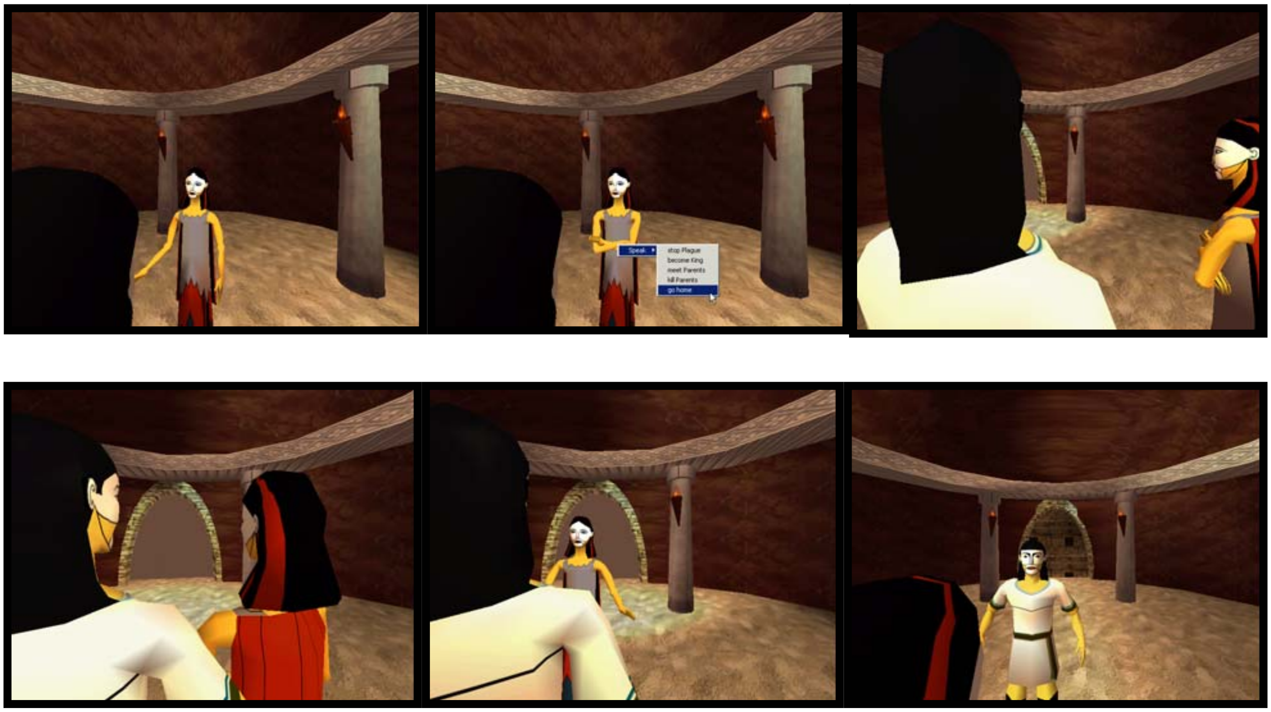

Figure 4. Sequence from scene 8 using dynamic light allocation

These advantages can be better portrayed by juxtaposing a scene rendered by both a dynamic and a static lighting technique (where artists manually position lights at design time), as shown in figures 3 (static) and 4 (dynamic). Figure 3 shows a scene rendered using static lighting, which, as depicted, suffers from partially lit characters due to (1) restriction of the number of lights used and (2) unanticipated camera and character movements. Additionally, as shown the visual focus does not match the dramatic focus of the scene. On the other hand, figure 4 shows the same scene rendered using ELE. As shown, ELE achieves better visibility and visual focus.

As described above, ELE can be adjusted to give more importance to depth than other lighting design goals. In such a situation, ELE may sacrifice realistic, dramatic, and visibility goals for depth. Figure 5 (right) shows a screenshot of a scene rendered using ELE, where the depth priority was emphasized. The figure also shows the effectiveness of ELE in portraying perceptual depth compared to other techniques, such as camera fill. In the figure lit by ELE (shown on the right), depth is much more pronounced because ELE assigned different lights to the background and the foreground and because the lights lighting the background were given lower intensity than those lighting the foreground. The camera fill approach renders a flatter image, as shown.

Figure 6 presents a sequence of screenshots showing ELE in action lighting a scene. The figure shows changes in lighting as the dramatic tension and action changes in the scene. ELE adapts the lighting colors by manipulating the color contrast to variation in dramatic tension, as depicted. These changes were made gradually and between camera cuts. For example, in the moment after the choice point, ELE accommodated the steep rise in dramatic tension by first requesting a cut from the camera system, and then changing the color contrast within the cut, thus maintaining visual continuity. Comparing the contrast before the choice point to the contrast after the choice point, one can see the dramatic difference. However, viewers were not distracted by this steep change, since it was done within a camera cut (Hollingworth in press).

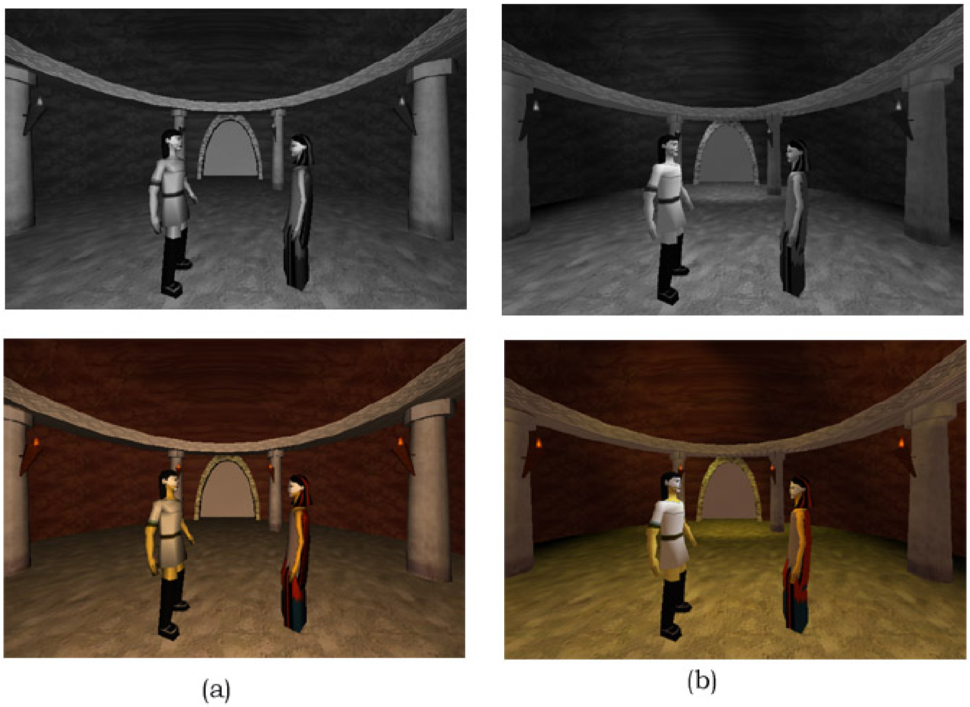

Figure 5. (a) shows a screenshot from scene 8 from *Mirage* where camera fill is used, (b) shows the same moment as (a) but with the lighting system where depth is given a high priority pictures on top portion of the image are b/w versions of the colored images shown

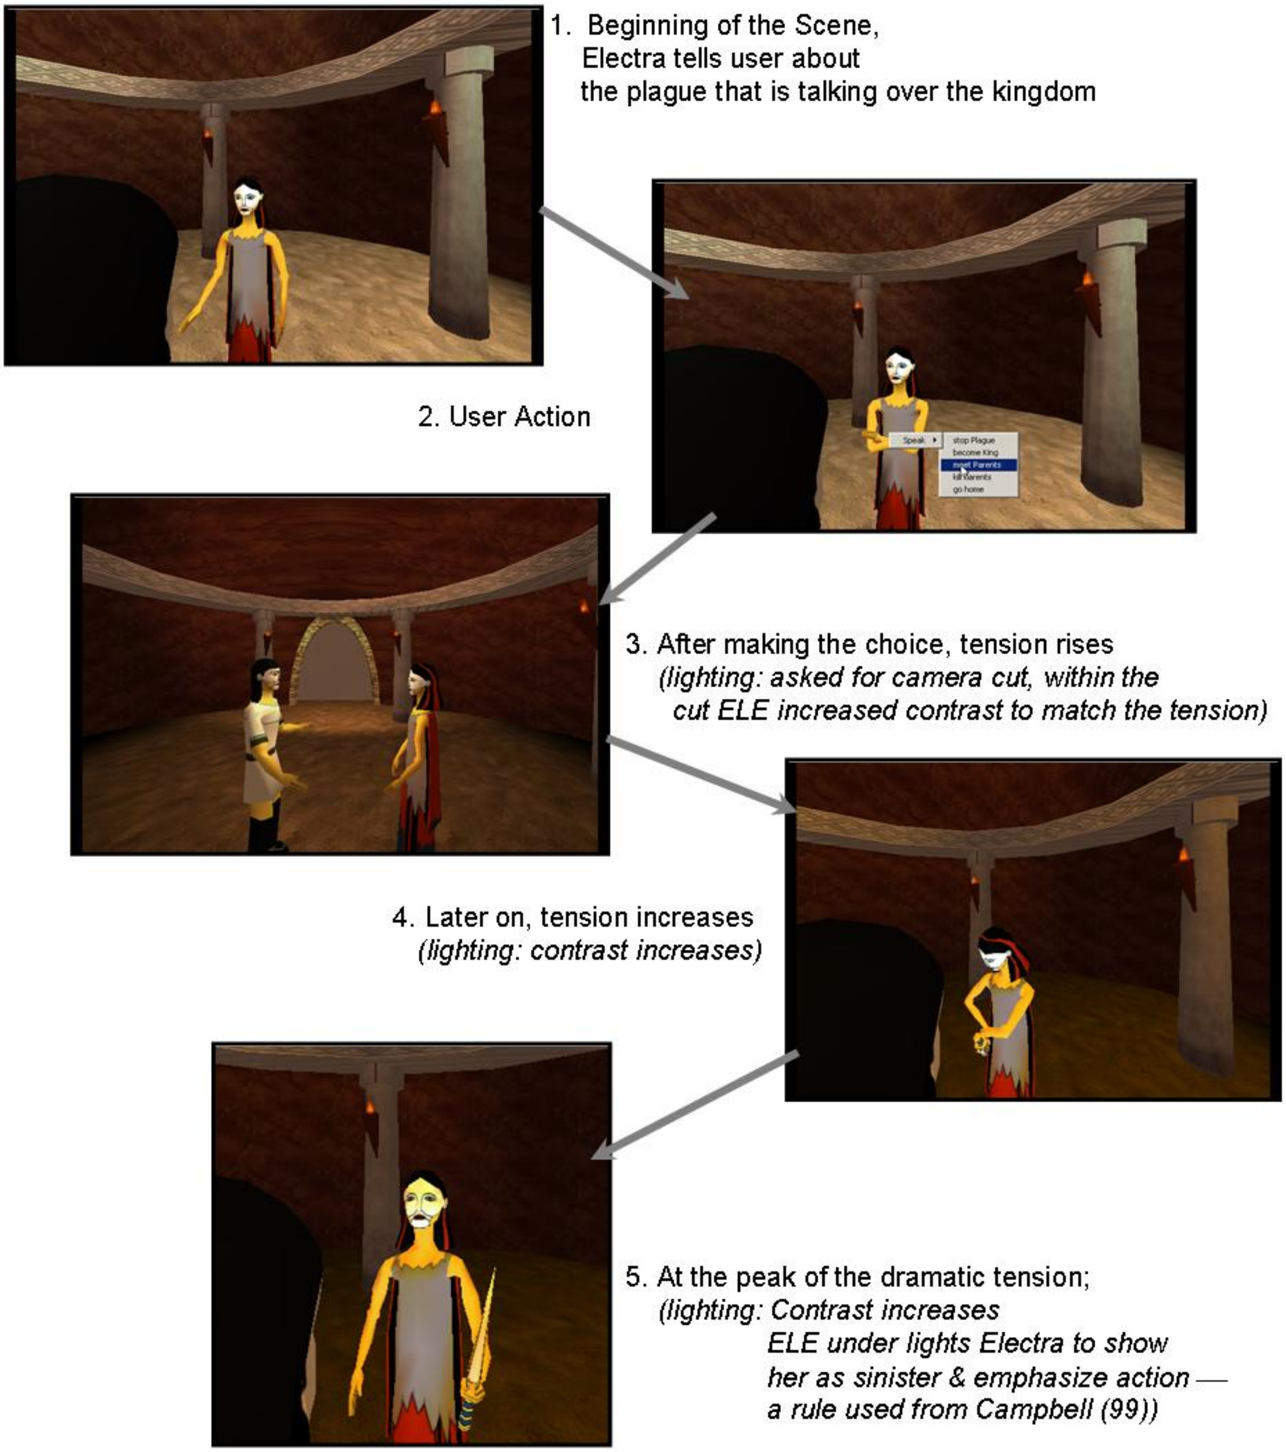

1. Beginning of the Scene,
Electra tells user about
the plague that is talking over the kingdom

2. User Action

3. After making the choice, tension rises
*(lighting: asked for camera cut, within the
cut ELE increased contrast to match the tension)*

4. Later on, tension increases
(lighting: contrast increases)

5. At the peak of the dramatic tension;
*(lighting: Contrast increases
ELE under lights Electra to show
her as sinister & emphasize action —
a rule used from Campbell (99))*

Figure 6. ELE lighting scene 8 from *Mirage*

5. DISCUSSION

Lighting design can be viewed as one of the important yet difficult processes of visual design because it involves balancing several intricate and delicate design goals. Under simple circumstances, one can adopt a lighting model that presents a solution to one or two of these goals. In this paper, we presented a lighting design model, ELE, that operates in an unpredictable and complex environment, where lighting conditions, camera properties, and character actions are not known at design time. ELE uses optimization to balance the many intricate and delicate visual design goals suggested by cinematic and theatric lighting design theory to accommodate for the unpredictability of the interactive environment.

We can summarize the contributions of ELE to the following. ELE proposes a new model of lighting design for interactive media based on an adaptation of cinematic and theatric lighting design theory. Game designers have long recognized the importance of cinematic and theatric lighting. They have already adopted several techniques from cinematic and theatric lighting, e.g. the use of low key lighting and color contrast to stimulate emotions, which can be seen in many horror games, e.g. *Silent Hill*. ELE expands on this idea by addressing several issues. We identify the need to automatically adapt the lighting in real-time to the changes in the situation, and thus ELE automatically adapts the lighting to suit the situation and enhance the tension and emotional stimulation. We also recognize that there are many functions of lighting and that balancing these functions is key to a good lighting design. Thus, ELE automatically manipulates the lighting while balancing its function subject to the desired style and the level of visual continuity required. We also acknowledge that the lighting design process is a very intricate, time and labor intensive design process. ELE presents a new approach to lighting design where designers can control lighting at a high-level by establishing constraints that specify style, aesthetic, and communicative functions of light. This approach can significantly facilitate and accelerate the lighting design process. We believe that by allowing designers to manipulate high-level design parameters for lights rather than low-level properties, e.g. positions and angles, this approach will transform the lighting process into a lighting design process.

ELE is just a small step towards adapting visual design theory from film and theater. ELE has several limitations. Shadows are a very important elements of lighting design. In fact, most of cinematic and theatric lighting design principles are designed based on shadows. There are two types of shadows: cast shadows of a character or an object on another, and shadows created by lack of lights hitting the surface of an object or area. ELE incorporates the later. Cast shadows was not incorporated in ELE due to the constraints of using real-time rendering algorithms that did not (at the time of development) incorporate cast shadows and self-cast shadows. However, the hypothesis is that with advances in real-time rendering these types of shadows will be incorporated, and ELE will need to be re-evaluated.

Lighting designers in film often use exposure techniques to manipulate the amount of light projected in some areas in an image (Viera 1993). Sometimes the lighting in an image appears different than expected due to interaction of light colors, and properties of the materials used (Birn 2000). Graphics designers often use a histogram to judge if a picture has some exposure-related problems, such as overexposure or underexposure (Birn 2000). A histogram is defined as a chart that shows tone frequency in an image (Birn 2000). We didn’t incorporate this technique in ELE. As a result, one can see overexposed images. One future direction to enhance lighting appearance is to use techniques such as histograms to ensure that the image portrays the qualities that the designer intended.

One limitation of the images shown above is the rendering engine used. This is not a limitation of ELE, but of the rendering engine chosen for ELE’s implementation. We didn’t address the light rendering question in this paper. The paper addresses an intelligent adaptive system for lighting design. It is abstracted from the rendering model used. However, its appearance will depend on the rendering engine used. We currently use Wildtangent as the rendering engine for Mirage (shown above). Wildtangent does not account for cast shadows or inter-reflections of light. Images produced by ELE could have been greatly enhanced had this feature been allowed. Thus, for future research, we will explore the possibility of adding shadows and calculating one or two level reflection. We will also explore the utility of using existing interactive global illumination methods, such as the work described by Tole et al.'s at Cornell (Tole 2002).

6. CONCLUSION

In this paper, we outlined a new model of lighting design for interactive scenes. The model, ELE (Expressive Lighting Engine), is grounded on many theories, including design theories from creative disciplines, such as film and theatre, and principles of vision from psychology, neuroscience, and perception. We have shown the utility of ELE as a model that automatically and in real-time adapts the lighting to the continually changing and unpredictable situation, while (a) satisfying visual design goals, including supplying visual focus, paralleling dramatic tension, and providing mood, and (b) maintaining style and visual continuity. We have discussed the principles, algorithms, and rules used by ELE to establish such adaptation. We have also shown the use of ELE within a scene from *Mirage*. The results show a great promise for the success of the approach. We have identified and discussed several limitations and directions for future research. Lighting is only one of many visual design elements; thus we view ELE as a step in accomplishing better and more adaptable visual designs that can morph to the narrative and enhance the overall interactive experience. In the future, we would like to extend our architecture to integrate camera, staging, and performance gestures.

7. REFERENCES

Alton, J. (1995). Painting with Light. Berkeley, University of California Press.

Birn, J., Ed. (2000). Digital Lighting & Rendering. Indianapolis, New Riders.

Block, B. (2001). The Visual Story: Seeing the Structure of Film, TV, and New Media. New York, Focal Press.

Calahan, S. (1996). Storytelling through lighting: a computer graphics perspective. Siggraph Course Notes.

Campbell, D. (1999). Technical Theatre for Non-technical People, Allworth Press.

Carson, D. (2000). Environmental Storytelling: Creating Immersive 3D Worlds Using Lessons Learned from the Theme Park Industry. Gamasutra.

Crowther, B. (1989). Film Noir: Reflections in a Dark Mirror. New York, Continuum.

Foss, B. (1992). Filmmaking: Narrative & Structural Techniques. Los Angeles, Silman-James Press.

Gillette, J. M. (1998). Designing with Light. Mountain View, CA, Mayfield.

Hill, B. H., Roger, T. , and Vorhagen, F. W. (1997). "Comparative Analysis of the Quantization of Color Spaces on the Basis of the CIELAB Color-Difference Formula." ACM Transactions on Graphics **16**(2): 109-154.

Hollingworth, H. a. (in press). "Global transsaccadic change blindness during scene preception." Psychological Science.

Katra, E. a. W., B. R. (1995). Preceived lightness/darkness and warmth/coolness in chromatic experience.

Kidd, W. (2001). Vermeer: Master of Light Video. J. J. Krakora. Washington, National Gallery of Art.

Knopf, D. C. a. A. (1979). The Book of Movie Photography. London, Alfred Knopf, Inc.

Lowell, R. (1992). Matters of Light and Depth. New York, Lowel-light Manufacturing, Inc.

Luo, M. R., Cul, G., and Rigg, B. (2000). The Development of CIE 2000 Colour Difference Formula: CIEDE2000, CIELAB. **2000**.

Maattaa, A. (2002). GDC 2002:Realistic Level Design for Max Payne. Gamasutra.

Millerson, G. (1991). The Technique of Lighting for Telivision and Film. Oxford, Focus Press.

Moller, T., and Haines, E. (1999). Real-time Rendering. MA, A K Peters, Ltd.

Thompson, D. B. a. K. (2001). Film Art: An Introduction. New York, Mc Graw Hill.

Tole, P., Pellacini, F., Walter, B., and Greenberg, D. P. (2002). Interactive Global Illumination in Dynamic Scene. Siggraph 2002, San Antonio, Texas, ACM.

Viera, D. (1993). Lighting for Film and Electronic Cinematography. Belmont, Wadsworth Publishing Company.