Creation and Rendering of Stochastic Dynamic Light Effects

D. Sekulovski[1], R. Clout[1], B. Kater[2], and T. Overbeek[1]; [1] Philips Research Europe, Eindhoven, The Netherlands; [2] G2 Speech B.V., Nuenen, The Netherlands

Abstract

We present a new method of dynamic light effect generation using stochastic models. Similar to dynamic lighting scenes in nature, the resulting light effects are unpredictable, yet recognizable. Furthermore, we present a method to learn the stochastic models from a video source of a natural scene. The method extracts the representative colors from the video and subsequently learns the typical transitions between the colors. After the model has been learned, the rendering of the effects has low memory and processing requirements, making it suitable for implementation even on embedded platforms. The recognition of the produced light effects was tested using a large user base and three automatically created models and a hand crafted one. The results show the suitability of the method for dynamic atmosphere creation, but also a high appreciation of the produced light effects.

Introduction

Recent advances in lighting introduced a revolution in the capabilities of lighting systems and elevate light to a status of a new medium. The main drivers of the advance, solid state light sources, enable lighting environments with a high spatial resolution, with fully controllable color and a wide range of dynamic capabilities.

One of the characteristics of modern lighting systems which is a large differentiator are the dynamic capabilities of the light sources used. The normal operating mode of most of general use traditional light sources, disco lights being one of the few exceptions, is static and they only have an on/off functionality or a limited dimming range. This is contrary to lighting conditions in nature, which are inherently dynamic, from the slow change of the intensity and the color temperature of daylight during one day, to the fast flashes of lightning in a thunderstorm. Furthermore, most of the light effects we experience in nature are unpredictable on a certain timescale.

The new capabilities also introduce new challenges as the number of controllable parameters is much higher compared to traditional lighting systems. As a result, the standard control paradigms used in traditional lighting systems become ineffective and a novel representation of the problem is needed to tackle the complexities.

One of the ways to simplify the control of modern lighting systems is to use a different medium, for example text, images, or video, as a representation of the desired ambiance and translation algorithms to translate it to the available light sources. This enables the users of the system to control it by simple to find examples. Due to the fact that the user directly selects the final effect and consequently a set of algorithms compute the control values of the light sources to realize the effect, this control paradigm is an example of the so called effect driven control.

In this paper, we propose a new method for generation of light effects which are locally unpredictable and non repeating, but resemble a natural light effect. The generation is done by simulating the execution of a stochastic process. The use of a simple stochastic process, a first order Markov chain, simplifies the creation of the stochastic models without a significant sacrifice on the expressive power. Furthermore, we present an unsupervised learning algorithm that produces a model based on a video of a natural light effect. To measure the recognizability and desirability of the produced light effects, a large scale user test was carried out, the results of which are presented.

Stochastic models for light effects generation

A model of a stochastic (random) process X, $X = \{X_1, X_2, ..., X_N\}$, is a set of rules that characterize the joint probability distribution between all its random variables. The simplest model is the one where the random variables $X_1, X_2, ..., X_N$ are independent and identically-distributed (i.i.d.). For many natural processes, however, the assumptions, specially the assumption of independence of the random variables, are unrealistic. An often used generalization from the i.i.d. processes comes in the class of Markov processes.

The probabilistic behavior of a Markov process is determined only by the dependencies between a subset of successive random variables. In the case of a first order Markov process the behavior is determined by the dependencies of immediate successors - between X_1 and X_2 , between X_2 and X_3 , etc. Despite their apparent simplicity and restrictions, Markov chains are rich in behavior, amenable to analysis, and adaptable to many applications, from weather to baseball prediction. They are centrally important to applied and theoretical probability. In the context of this work, we concentrate only on Markov models with discrete times and a finite number of states, or finite state space Markov chains.

Notably, one of the uses of Markov chains close to the one presented in this paper is for generation of media. The application that motivated the development of Markov models, text generation based on rules, has been used in a wide array of artificial intelligence applications. The rules in the model are in the form of conditional probabilities on the succession of words, learned from a large corpus of example content. Testing the performance of text processing algorithms often uses text content generated using Markov chain Monte Carlo [1]. An interesting and humorous example of the recognizability of the produced text by a random process related to Markov chains is the random scientific paper generator [2], which produced a "scientific paper" that was accepted on a conference with a questionable review process. Markov chains have been used for modeling music [4] and for automatic generation of music [3].

In Markov chain applications, the index t in X_t is usually thought of as a time index. X_t represents the state of the Markov chain model of the process at "time" t. Formally, a Markov chain is a sequence $X = \{X_t\}_{t\geq 1} = \{X_1, X_2, ...\}$ of random variables taking values in a discrete set $\mathscr{E}$, and a matrix of conditional probabilities $A = (p_{ij})$, called the transition probability matrix. The elements of $\mathscr{E}$ are called states. The elements of the transition probability matrix p_{ij} give the probability that the process will be in a state j at "time" $t+1$, given that it was in state i at "time" t, $p_{ij} = \{X_{t+1} = j | X_t = i\}$.

To use a Markov chain as a model of a process, the process has to be represented as a time varying sequence of a finite number of unique states. In the case of light effect creation, this representation is a natural one. In this case, the states can be mapped to single colors on a particular luminaire or a probability distribution in a color space.

The set of states of a simple example light effect, a night sky with lighting, would have two colors, a dark blue one and a warm white one. The transition probability matrix of such an example light effect is also simple. The dark blue color state would have a high probability of not changing state and a very small probability of changing state to the warm white state, while the warm white state would have a very small probability of not changing the state and a large probability of transition to the dark blue state. In more complex examples the state space of the process can be a set of probability distributions instead of single colors. In the above example of a stormy night, the warm white state can be substituted with a three dimensional Gaussian distribution in a color space centered around a warm white color.

Given a Markov chain with a set of states and the transition probability matrix, a light effect can be generated by simulating the Markov process. The starting state of the system can be a user set state, a random state or the most probable state of the Markov chain. At regular time intervals, a new state of the system, and thus a new color of the luminaire is computed by sampling from the distribution given by the transition probability matrix. As the new state only depends on the current state of the system, the sampling of the distribution is carried out by generating a uniform $U([0,1))$ random number and based on that number a search through the cumulative distribution of the transition probabilities for the given current state. The implementation of this operation is straightforward and computationally cheap, making it ideal for implementation on embedded platforms. The suitability for embedded use is boosted by the small memory requirement of the generated model, that is quadratic to the number of states. Additionally, if a central control is used to control multiple light sources, only one copy of the model is needed for all of the light sources.

Learning

Given a video of a light effect, the following method for learning of a stochastic model of the light effect is proposed. As the input to the learning algorithm, a video depicting the target light effect is used. The learning of the model constitutes of three main steps.

The first step is to extract representative colors from each video frame. This is done using a central tendency estimator, for example the sample mean or the sample color median, of the colors present in the video frame. The first step transforms the video into a discrete time sequence of colors representing the video frames.

The second step is the clustering of the representative colors from all the frames in the video into a small number of classes, whose centroids will represent the states in the stochastic model. The clustering is done by using the blurring mean shift algorithm [5, 6] on the representative colors represented in a nearly perceptually uniform color space ($CIEL^*a^*b^*$). The blurring mean shift is a density based algorithm that takes into account the local structure of the distribution in the clustering. The advantage of the mean shift algorithm over standard clustering algorithms like k-means is that the input parameter to the method, the size of the kernel, can be set using a meaningful criterion, the minimum distance between two cluster centroids that are not merged during the clustering. After the clustering, the colors in the time sequence are substituted with their respective class representatives producing a quantized time sequence. In case of complex states, the probability distribution of every state can be estimated from the colors that were clustered together to form the state.

Using the time sequence of class representatives, the state transition probabilities are estimated using the frequency of transitions between consecutive states in the source material. Assuming the time sequence is generated from a Markov chain, the frequency of state transitions between consecutive states is a maximum likelihood estimator of the transitions probabilities.

User study

To validate the recognition rate and the desirability of the produced light effects, a user study was conducted during a corporate research exhibition. The study was designed with an application of home scene setting in mind, which influenced the setup and the method used.

Setup

Two luminaires were used, each having three independent RGB LED light sources. The participant could only observe the light reflected from a white surface and couldn't see the light sources or the encasing luminaires.

Stimuli

Four stimuli were used, three automatically created and one hand crafted.

The first stimulus, *fire* was created from a video of a beach fire. The second stimulus, *underwater*, was created from a low resolution representation of an underwater scene. The transition probabilities of the model were computed using the spatial neighbor probabilities, resulting in an effect equivalent to the one produced from a video of a camera randomly moving over the scene with a constant speed. The third stimulus, *fireworks*, was a manually built impression of multicolored fireworks with periods of faster and slower dynamics. The fourth stimulus, *clouds*, was created from a time lapse video of a cloudy sky.

Two of the stimuli, *fire* and *fireworks* had fast dynamics, contrary to the other two, which had long term, smooth dynamics. The stimulus *fire* had predominantly warm colors, the stimuli *underwater* and *clouds* had cold colors and the stimulus *fireworks* didn't use a specific selection of colors.

Method

The desirability of the produced light effects was measured using a seven point Likert scale [7]. For each of the stimuli, the

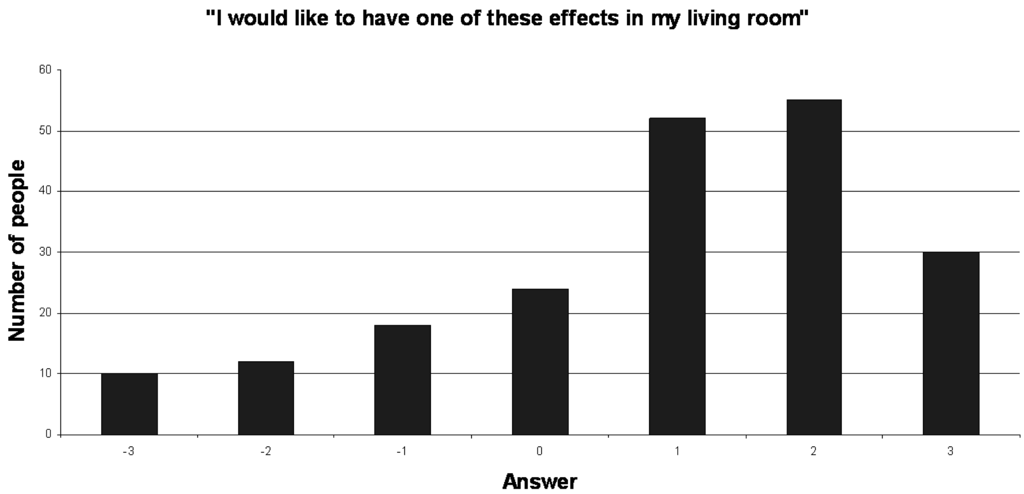

Figure 1. *Histogram of the maximum score over the four stimuly on the question "I would like to have this light effect in my living room".*

participants answered the question "I would like to have this light effect in my living room" on a scale from "Not at all" (-3) to "Very much" (3). Additionally, for each stimulus a set of questions measuring the recognition of the stimulus were asked. Four recognition questions were asked for each stimulus, "This effect looks like a fire", "This effect looks like an underwater scene", "This effect looks like fireworks", and "This effect looks like a cloudy sky". The possible answers ranged from "Not at all" (-3) to "Very much" (3).

The authors are aware of the possible bias on the recognition rate produced by naming the effects that were presented, but considering the application of home scene setting, where the user already knows the target effect she picked, this was not considered a problem in the context of this study. Furthermore, including all four effects in the questionnaire for each stimulus enabled a computation of a confusion matrix, given in the results section. Using another method as for example free association, would have resulted in a considerably longer testing time and would require subjective evaluation of the results.

The order of presentation of the stimuli was balanced over the participants. Additional to the questions, participants could provide additional comments.

Participants

The study was conducted with 202 participants, 155 of which male, with a minimum age of 24, a maximum age of 59 and a median age of 35.5. 64 of the participants had experience working with atmosphere providing light sources.

Results

Results on suitability of the produced light effects for use in a living room environment show that the most desired stimulus was *clouds*, with a median result of 2, followed by *underwater*, with a median result of 1, *fire* with 0 and *fireworks*, which was scored as highly unsuitable for the context and had a median score of -2. As the stimulus *fireworks* was very dynamic, this result is not surprising. It is unclear, though, both from the answers to the questions and the additional comments given, what the reason for the difference of desirability of the other three stimuli is. Some people mention the color temperature as one of the reasons they scored a certain light effect high or low, while others give the dynamics as the primary reason. As there was no slow warm stimulus, a conclusion on the relative importance of these factors cannot be given.

To judge the overall desirability of the produced light effects, the maximum score over the four stimuli was computed. Figure 1 shows a histogram of the resulting scores. As can be seen from the histogram, the overall desirability was scored high, with a median score of 2. This shows, together with the comments, that people think that dynamic light effects can be suitable for use in their living room, but only if the effects are very localized like the *fire*, or very slow and subtle as the *clouds* and *underwater* stimuli.

Figure 2 shows the median result on the answers from the recognition questions. The images in the row header depict the stimulus that was presented to the participant. The median of the scored similarity to the effects given in the column header is given in the figure. The general recognition rate of all the effects is high, with most of them having a median score 2 for the matching stimulus. The most easily recognizable effect wes the *fireworks*, while the most confused one was the *underwater*. An interesting effect can be seen in the confusion of the *underwater* and the *clouds* stimuli. While when presented with the *underwater* stimulus, participants scored equally high for both *underwater* and *clouds*, when presented with the *clouds* stimulus, the score for *clouds* was significantly higher than the score for *underwater*. It was also observed that the effect persists over different orderings of the stimuli. This additionally shows that different light effects have a different range in which they can be recognized and for some of them that range can be very large.

Conclusions

A new method of dynamic light effect generation using stochastic models was presented. Similar to dynamic lighting scenes in nature, the resulting light effects are unpredictable, yet

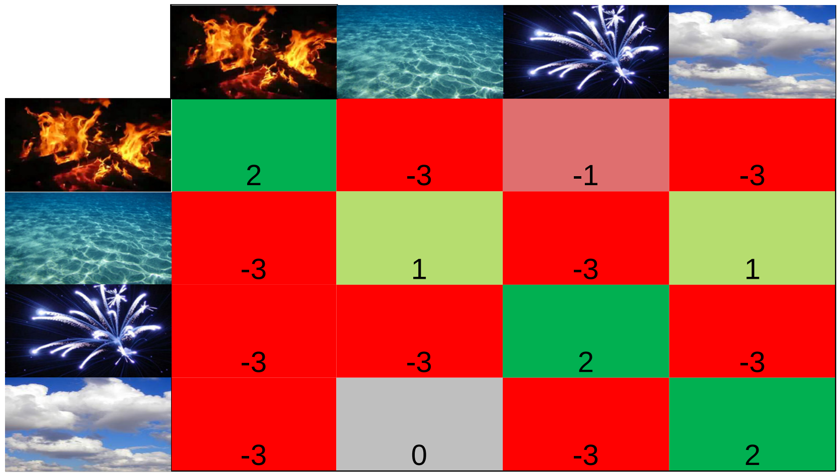

Figure 2. *The median response of the participants to the question "This effect looks like ..." for all stimuli.*

recognizable. Next, a method to learn the stochastic models from a video source of a natural scene was shown. The method extracts the representative colors from the video and subsequently learns the typical transitions between the colors. After the model has been learned, the rendering of the effects has low memory and processing requirements, making it suitable for implementation even on embedded platforms. The recognition of the produced light effects was tested using a large user base and three automatically created models and a hand crafted one. The results show the suitability of the method for dynamic atmosphere creation, but also a high appreciation of the produced light effects.

References

[1] Dani Gamerman, Hedibert Freitas Lopes, *Markov chain Monte Carlo: stochastic simulation for Bayesian inference*, CRC Press, 2006.

[2] Jeremy Stribling, Max Krohn, Dan Aguayo *SciGen* http://pdos.csail.mit.edu/scigen/.

[3] Lejaren A. Hiller, L. M. Isaacson, *Experimental Music-Composition with an Electronic Computer*, New York: McGraw-Hill, 1959.

[4] David M. Franz, *Markov Chains as Tool for Jazz Improvisation Analysis*, thesis, Faculty of the Virginia Polytechnic, Institute and State University, Blackburg, Virginia.

[5] K. Fukunaga, L. Hostetler *The estimation of the gradient of a density function, with applications in pattern recognition*, IEEE Transactions on Information Theory, 1975, 21 - 1, pp 32–40.

[6] D. Comaniciu, P. Meer, *Robust Analysis of Feature Spaces: Color Image Segmentation*, IEEE Conf. on Comp. Vis. and Pattern Recognition, 1997, pp 750–755.

[7] Rensis Likert, *A Technique for the Measurement of Attitudes*, Archives of Psychology 140: pp 1–55, 1932.

Author Biography

Dragan Sekulovski is a Research Scientist at the Philips Research Europe laboratory in Eindhoven, The Netherlands, where he works on topics related to models of visual perception and their use in lighting and display applications. He has a master of Computer Science from the university "St. Cyril and Methodius", Skopje and is currently pursuing a PhD on the topic of algorithms for ambient intelligent lighting. His interests include intuitive and perception based user interaction with lighting systems.