

1-

FILM GRAIN SIMULATION METHOD BASED ON PRE-COMPUTED TRANSFORM COEFFICIENTS

CROSS-REFERENCE TO RELATED APPLICATIONS

5

This application claims priority under 35 U.S.C. 119(e) to U.S. Provisional Patent Application Serial No 60/630640, filed November 23, 2004, the teachings of which are incorporated herein.

10 TECHNICAL FIELD

This invention relates to a technique for simulating film grain in an image.

BACKGROUND OF THE INVENTION

15

Motion picture films comprise silver-halide crystals dispersed in an emulsion, coated in thin layers on a film base. The exposure and development of these crystals form the photographic image consisting of discrete tiny particles of silver. In color negatives, the silver undergoes chemical removal after development and tiny blobs of dye occur on the sites where the silver crystals form. These small specks of dye are commonly called 'grain' in color film. Grain appears randomly distributed on the resulting image because of the random formation of silver crystals on the original emulsion. Within a uniformly exposed area, some crystals develop after exposure while others do not.

20

Grain varies in size and shape. The faster the film, the larger the clumps of silver formed and blobs of dye generated, and the more they tend to group together in random patterns. The grain pattern is typically known as 'granularity'. The naked eye cannot distinguish individual grains, which vary from 0.0002 mm to about 0.002 mm. Instead, the eye resolves groups of grains, referred to as blobs. A viewer identifies these groups of blobs as film grain. As the image resolution becomes larger, the perception of the film grain becomes higher. Film grain becomes clearly noticeable on cinema and high-definition images, whereas film grain progressively loses importance in SDTV and becomes imperceptible in smaller formats.

30

Motion picture film typically contains image-dependent noise resulting either from the physical process of exposure and development of the photographic film or from the subsequent editing of the images. The photographic film possesses a characteristic quasi-

35

2-

random pattern, or texture, resulting from physical granularity of the photographic emulsion. Alternatively, a similar pattern can be simulated over computed-generated images in order to blend them with photographic film. In both cases, this image-dependent noise is referred to as grain. Quite often, moderate grain texture presents a desirable feature in motion pictures. In some instances, the film grain provides visual cues that facilitate the correct perception of two-dimensional pictures. Film grain is often varied within a single film to provide various clues as to time reference, point of view, etc. Many other technical and artistic uses exist for controlling grain texture in the motion picture industry. Therefore, preserving the grainy appearance of images throughout image processing and delivery chain has become a requirement in the motion picture industry.

Several commercially available products have the capability of simulating film grain, often for blending a computer-generated object into a natural scene. Cineon® from Eastman Kodak Co, Rochester New York, one of the first digital film applications to implement grain simulation, produces very realistic results for many grain types. However, the Cineon® application does not yield good performance for many high-speed films because of the noticeable diagonal stripes the application produces for high grain size settings. Further, the Cineon® application fails to simulate grain with adequate fidelity when images are subject to previous processing, for example, such as when the images are copied or digitally processed.

Another commercial product that simulates film grain is *Grain Surgery*™ from Visual Infinity Inc., which is used as a plug-in of Adobe® After Effects®. The *Grain Surgery*™ product appears to generate synthetic grain by filtering a set of random numbers. This approach suffers from disadvantage of a high computational complexity.

None of these past schemes solves the problem of restoring film grain in compressed video. Film grain constitutes a high frequency quasi-random phenomenon that typically cannot undergo compression using conventional spatial and temporal methods that take advantage of redundancies in the video sequences. Attempts to process film-originated images using MPEG-2 or ITU-T/ISO H.264 compression techniques usually result either in an unacceptably low degree of compression or complete loss of the grain texture.

Thus, there exists a need for a technique simulating film grain, especially a technique that affords relatively low complexity.

BRIEF SUMMARY OF THE INVENTION

In accordance with a first aspect of the present invention, there is provided a method for creating a block of film grain, comprising the steps of (a) obtaining a pre-computed block of transformed coefficients; (b) frequency filtering the transformed coefficients responsive to a frequency range that characterize a desired pattern of film grain; and (c) performing an inverse transform on the frequency-filtered transformed coefficients.

In accordance with a second aspect of the present invention, there is provided an apparatus comprising: a first memory for storing at least one a pre-computed block of transformed coefficients; at least one of a programmed processor, and dedicated logic circuit for simulating film grain by (a) obtaining a pre-computed block of transformed coefficients from the memory; (b) frequency filtering the transformed coefficients responsive to a frequency range that characterize a desired pattern of film grain; and (c) performing an inverse transform on the frequency-filtered transformed coefficients.

In accordance with a third aspect of the present invention, there is provided an apparatus comprising: a first memory for storing at least one a pre-computed block of transformed coefficients; means for obtaining a pre-computed block of transformed coefficients from the memory; means for frequency filtering the transformed coefficients responsive to a frequency range that characterize a desired pattern of film grain; and means for performing an inverse transform on the frequency-filtered transformed coefficients.

DETAILED DESCRIPTION OF THE DRAWINGS

FIGURE 1 depicts a block schematic diagram of the combination of a transmitter and receiver in a film grain processing chain useful for practicing the technique of the present principles;

FIGURE 2 depicts, in flow chart form, the steps of a first method for creating a block of film grain using pre-computed coefficients;

FIGURE 3 depicts, in flow chart form, the steps of a method for film grain pattern creation using pre-computed Discrete Cosine Transformation (DCT) coefficients of a single image of Gaussian Noise; and

FIGURE 4 depicts, in flow chart form, the steps of a method for film grain pattern creation using pre-computed Discrete Cosine Transformation (DCT) coefficients of several images of Gaussian Noise.

DETAILED DESCRIPTION

To understand the technique of the present principles for simulating film grain using a set of pre-computed transformed coefficients, a brief overview of film grain simulation will prove helpful. FIGURE 1 depicts a block schematic diagram of a transmitter 10, which receives an input video signal and, in turn, generates a compressed video stream at its output. In addition, the transmitter 10 also generates information indicative of the film grain (if any)

4-

present in the sample. In practice, the transmitter 10 could comprises part of a head-end array of a cable television system, or other such system that distributes compressed video to one or more downstream receivers 11, only one of which is shown in FIG. 1. The transmitter 10 could also take the form of encoder that presents media like DVDs. The receiver 11 decodes the coded video stream and simulates film grain in accordance with the film grain information and decoded video, both received from the transmitter 10 or directly from the media itself in the case of a DVD or the like, to yield an output video stream that has simulated film grain. The receiver 11 can take the form of a set-top box or other such mechanism that serves to decode compressed video and simulate film grain in that video.

The overall management of film grain requires the transmitter 10 (i.e., the encoder) provide information with respect to the film grain in the incoming video. In other words, the transmitter 10 "models" the film grain. Further the receiver 11 (i.e., decoder) simulates the film grain according to the film grain information received from the transmitter 10. The transmitter 10 enhances the quality of the compressed video by enabling the receiver 11 to simulate film grain in the video signal when difficulty exists in retaining the film grain during the video coding process.

In the illustrated embodiment of FIG. 1, the transmitter 10 includes a video encoder 12 which encodes the video stream using any of the well known video compression techniques such as the ITU-T Rec. H.264 | ISO/IEC 14496-10 video compression standard. Optionally, a film grain remover 14, in the form of a filter or the like depicted in dashed lines in FIG. 1, could exist upstream of the encoder 12 to remove any film grain in the incoming video stream prior to encoding. To the extent that the incoming video contains no film grain, no need would exist for the film grain remover 14.

A film grain modeler 16 accepts the input video stream, as well as the output signal of the film grain remover 14 (when present). Using such input information, the film grain modeler 16 establishes the film grain in the incoming video signal. In its simplest form, the film grain modeler 16 could comprise a look up table containing film grain models for different film stocks. Information in the incoming video signal would specify the particular film stock originally used to record the image prior to conversion into a video signal, thus allowing the film grain modeler 16 to select the appropriate film grain model for such film stock. Alternatively, the film grain modeler 16 could comprise a processor or dedicated logic

5-

circuit that would execute one or more algorithms to sample the incoming video and determine the film grain pattern that is present.

The receiver 11 typically includes a video decoder 18 that serves to decode the compressed video stream received from the transmitter 10. The structure of the decoder 18 will depend on the type of compression performed by the encoder 12 within the transmitter 10. Thus, for example, the use within the transmitter 10 of an encoder 12 that employs the ITU-T Rec. H.264 | ISO/IEC 14496-10 video compression standard to compress outgoing video will dictate the need for an H.264-compliant decoder 18. Within the receiver 11, a film grain simulator 20 receives the film grain information from the film grain model 16. The film grain simulator 20 can take the form of a programmed processor, or dedicated logic circuit having the capability of simulating film grain for combination via a combiner 22 with the decoded video stream.

Film grain simulation aims to synthesize film grain samples that simulate the look of the original film content. As described, film grain modeling occurs at the transmitter 10 of FIG. 1, whereas film grain simulation occurs at the receiver 11. In particular, film grain simulation occurs in the receiver 11 along with the decoding the incoming video stream from the transmitter 10 upstream of the output of the decoded video stream. Note that the decoding process that occurs in the receiver 11 makes no use of images with added film grain. Rather, film grain simulation constitutes a post-processing method for synthesizing simulated film grain in the decoded images for display. For that reason, the ITU-T Rec. H.264 | ISO/IEC 14496-10 video compression standard contains no specifications regarding the film grain simulation process. However, film grain simulation requires information concerning the grain pattern in the incoming video signal, which information typically undergoes transmission in a film grain characteristics Supplemental Enhancement Information (SEI) message when using the ITU-T Rec. H.264 | ISO/IEC 14496-10 video compression standard as specified by the Amendment 1 (Fidelity Range Extensions) of that compression standard.

Film grain simulation in accordance with the present principles makes use of a pre-computed set of transformed coefficients. In other words, the simulation process begins with a block, typically, although not necessarily, of $N \times N$ size, whose coefficients have been transformed in advance of simulation, typically, but not necessarily, using a Discrete Cosine Transformation. Other transformations are possible. Bit-accurate implementation of the

6-

simulation method of the present principles occurs by performing an integer inverse transform on the block of pre-computed values.

The film grain simulation method of the present principles presents an interesting tradeoff between complexity and memory requirements when compared to previously disclosed methods. On the one hand, the present method reduces the complexity of the transform-based approaches, by avoiding computation of the direct transform. On the other hand, the method reduces the memory requirements of database-based approaches by storing transformed coefficients instead of film grain patterns. The proposed method could be applied to HD DVD systems, BD ROM systems as well as satellite broadcasting just to cite a few examples.

FIGURE 2 illustrates in flow-chart form, the steps of the method of the present principles, for simulating film grain using pre-computed coefficients. The method of FIG. 2 commences upon execution of the start step 100 during which initialization typically happens, although such initialization need not necessarily to occur. Thereafter, step 102 occurs during which a block of pre-computed transformed coefficients, typically, although not necessarily, $N \times N$ in size, is read from a memory 103. Many techniques exist for creating the set of pre-computed coefficients in the memory 103 of FIG. 2. For example, the transformed coefficients could be pre-computed using a Discrete Cosine Transformation (DCT) on a set of random values, as described in detail with respect to FIGS. 3 and 4. Other techniques exist for computing a set of transformed coefficients in advance of film grain simulation. The transformed coefficients then undergo frequency filtering during step 105 using a predefined set of cut frequencies f_{HL} , f_{VL} , f_{HH} , and f_{VH} that represent cut-off frequencies (in two dimensions) of a filter that characterizes the desired film grain pattern. During step 106, the frequency filtered block of transformed coefficients undergoes an inverse transformation, typically, although not necessarily, an Inverse Discrete Cosine Transformation (IDCT), to yield a block of film grain, before the process ends with step 108. Under some circumstances, scaling of the inversely transformed block of coefficient following step 106 will prove useful prior to ending of the process.

FIGURE 3 illustrates in flow chart form a method for film grain pattern simulation using pre-computed DCT coefficients of a single image of Gaussian random noise. The method of FIG. 3 commences upon execution of the start step 300 during which initialization occurs, although such initialization need not necessarily happen. Thereafter, step 202 occurs

7-

during which a block of pre-computed transformed coefficients, typically, although not necessarily $N \times N$ in size, is read from a memory 203. The block of coefficients read from the memory 203 during step 202 typically has been created by applying a DCT transform to an $N \times N$ image of Gaussian random values.

5 Following step 202, step 204 occurs, initiating entry into a loop that repeats for all possible film grain sizes and shapes. Upon entry into the loop, the transformed coefficients undergo frequency filtering during step 205 using a predefined set of cut frequencies f_{HL} , f_{VL} , f_{HH} , and f_{VH} that represent cut-off frequencies (in two dimensions) of a filter that characterizes the desired film grain pattern. Next step 206 occurs during which the frequency filtered block
10 of transformed coefficients undergoes an inverse transformation, typically, although not necessarily, an Inverse Discrete Cosine Transformation (IDCT), to yield a block of film grain. Under some circumstances, scaling of the inversely transformed block of coefficients following or preceding step 206 will prove useful.

 Thereafter, the block of film grain that results from the inverse transformation
15 performed during step 206 (or after scaling if such scaling occurs) undergoes storage in a database 209 during step 208. The loop initiated during step 204 undergoes repeating (i.e., steps 205 and 206 are repeated) for all possible film grain sizes and shapes, whereupon loop execution ends during step 210, after which the method ends during step 212. As can be appreciated from the foregoing description, the single image of Gaussian random values read
20 from the memory 203 during step 202 of FIG. 3 serves as the kernel for each of the film grain patterns stored in the database 210.

 FIGURE 4 illustrates in flow chart form a method for film grain pattern simulation using pre-computed DCT coefficients of multiple images of Gaussian random noise. The method of FIG. 4 commences upon execution of the start step 300 during which initialization
25 occurs, although such initialization need not necessarily happen. Next step 301 occurs, initiating entry into a loop that repeats for all possible film grain size and shape. Upon entry into the loop, step 302 occurs and a block of pre-computed transformed coefficients, typically, although not necessarily, $N \times N$ in size, is read from a memory 303. The block of coefficients read from the memory 303 during step 302 typically comprises a database of DCT transforms
30 of $N \times N$ images of Gaussian random values.

 Following step 302, step 305 occurs and the $N \times N$ image of DCT coefficients read from the memory 303 undergoes frequency filtering using a predefined set of cut frequencies

8-

f_{HL} , f_{VL} , f_{HH} , and f_{VH} that represent cut-off frequencies (in two dimensions) of a filter that characterizes the desired film grain pattern. Next step 306 occurs during which the frequency filtered block of transformed coefficients undergoes an inverse transformation, typically, although not necessarily, an Inverse Discrete Cosine Transformation (IDCT), to yield a block of film grain. Under some circumstances, scaling of the inversely transformed block of coefficients following step 306 will prove useful.

Thereafter, step 308 occurs and the block of film grain that results from the inverse transformation performed (and scaling if performed) undergoes storage in a database 309 during step 308.

The loop initiated during step 301 undergoes repeating (i.e., steps 302-308 are repeated) for all possible film grain sizes and shapes, whereupon loop execution ends during step 310, after which the method ends during step 312. Unlike the method of FIG. 3 that makes use of a single DCT image of Gaussian random values for all possible film grain shapes and sizes, the method of FIG. 4 makes use of a separate DCT image for each different film grain shape and size. Both methods allow creating a database of film grain values upon system initialization or reset for storage in conventional memory such as an SRAM memory.

The foregoing describes a technique for simulating film grain using pre-computed transformed coefficients, thereby reducing complexity.

9-

CLAIMS

- 1 1. A method for creating a block of film grain, comprising the steps of
2 (a) obtaining a pre-computed block of transformed coefficients;
3 (b) frequency filtering the transformed coefficients responsive to a frequency range
4 that characterize a desired pattern of film grain; and
5 (c) performing an inverse transform on the frequency-filtered transformed coefficients.
- 1 2. The method according to claim 1 wherein the step of obtaining the pre-
2 computed block of transformed coefficients further comprises the step of reading at least one
3 pre-computed block of transformed coefficients from a memory.
- 1 3. The method according to claim 1 wherein the step of frequency filtering further
2 comprises the step of filtering the transformed coefficients according to a set of cut
3 frequencies f_{HL} , f_{VL} , f_{HH} and f_{VH} of a two-dimensional filter that characterizes the desired film
4 grain pattern.
- 1 4. The method according to claim 1 wherein the step of performing an inverse
2 transform coefficients further comprises the step of performing an inverse discrete cosine
3 transform.
- 1 5. The method according to claim 1 further comprising the step of scaling the
2 inversely transformed block of coefficients.
- 1 6. The method according to claim 1 wherein steps (b) and (c) are repeated for all
2 possible film grain sizes and shapes to obtain a plurality of film grain patterns all derived from
3 one pre-computed block of transformed coefficients.
- 1 7. The method according to claim 1 wherein steps (a), (b), and (c) are repeated for
2 all possible film grain sizes and shapes to obtain a plurality of film grain patterns, each
3 derived from a separate pre-computed block of transformed coefficients.

1 8. The method according to claim 6 further comprising the step of storing each
2 inversely transformed set of filtered coefficients in a memory upon each repeated execution of
3 step (c).

1 9. The method according to claim 7 further comprising the step of storing each
2 inversely transformed set of filtered coefficients in a memory upon each repeated execution of
3 step (c).

1 10. The method according to claim 1 further comprising the step of scaling the
2 inversely transformed set of filtered coefficients before the coefficients undergo inverse
3 transformation.

1 11. Apparatus comprising:
2 a first memory for storing at least one a pre-computed block of transformed
3 coefficients;
4 at least one of a programmed processor, and dedicated logic circuit for simulating film
5 grain by (a) obtaining a pre-computed block of transformed coefficients from the memory; (b)
6 frequency filtering the transformed coefficients responsive to a frequency range that
7 characterize a desired pattern of film grain; and (c) performing an inverse transform on the
8 frequency-filtered transformed coefficients.

1 12. The apparatus according to claim 11 wherein the at least one of a programmed
2 processor, and dedicated logic circuit further scales the inversely transformed frequency-
3 filtered transformed coefficients.

1 13. The apparatus according to claim 11 further comprising a second memory for
2 storing the inversely transformed frequency-filtered transformed coefficients.

1 14. Apparatus comprising:
2 a first memory for storing at least one a pre-computed block of transformed
3 coefficients;
4 means for obtaining a pre-computed block of transformed coefficients from the
5 memory;

11-

6 means for frequency filtering the transformed coefficients responsive to a frequency
7 range that characterize a desired pattern of film grain; and
8 means for performing an inverse transform on the frequency-filtered transformed
9 coefficients.

1 15. The apparatus according to claim 14 further comprising means for scaling the
2 inversely transformed frequency-filtered transformed coefficients.

1 16. The apparatus according to claim 14 further comprising a second memory for
2 storing the inversely transformed frequency-filtered transformed coefficients.