Film Grain Synthesis and its Application to Re-graining

Peter Schallauer*, Roland Mörzinger

JOANNEUM RESEARCH, Institute of Information Systems and Information Management,
Steyrergasse 17, 8010 Graz, Austria

ABSTRACT

Digital film restoration and special effects compositing require more and more automatic procedures for movie re-graining. Missing or inhomogeneous grain decreases perceived quality. For the purpose of grain synthesis an existing texture synthesis algorithm has been evaluated and optimized. We show that this algorithm can produce synthetic grain which is perceptually similar to a given grain template, which has high spatial and temporal variation and which can be applied to multi-spectral images. Furthermore a re-grain application framework is proposed, which synthesises based on an input grain template artificial grain and composites this together with the original image content. Due to its modular approach this framework supports manual as well as automatic re-graining applications. Two example applications are presented, one for re-graining an entire movie and one for fully automatic re-graining of image regions produced by restoration algorithms. Low computational cost of the proposed algorithms allows application in industrial grade software.

Keywords: Grain, noise, synthesis, template, restoration, film, video, cinema, re-grain, texture, quality

1. INTRODUCTION

Digital film restoration and video production systems are application areas demanding the artificial generation of film grain and electronic noise, two typical phenomena in film and video material. Film grain synthesis is a necessary function for semi-automated digital film restoration. Defect removal algorithms often impair the original grain pattern, as original film grain and noise are equalized due to their random occurrence. To reach optimal image restoration quality appropriate film grain needs to be reinserted in restored areas. Similarly, a common problem in film special effects is to replicate film grain for adjusting the grain pattern of computer generated image parts when composited with film footage or when different kinds of film footage shall be combined to produce a homogenously grained movie.

These tasks pose the following requirements on film grain synthesis:

- generation of artificial grain perceptually similar to a certain grain sample
- sufficient spatial and temporal variation
- expansion to capability to arbitrary image size
- support for multi-spectral images
- low computational cost

For artificial grain generation three main approaches have been proposed. The first relies on creating film grain by means of mathematical models ([7]), the second is a patch-based approach ([2]) and the third uses texture statistics derived from exemplary grain or noise templates ([3, 5]). The large diversity in the appearance of film grain, which varies with film stock and lighting conditions and the fact that film grain intensity is nonlinearly related to film exposure, makes it difficult to model film grain mathematically.
The method proposed in [7] assumes the availability of both, the corrupted and ideal image. When the latter is not known, which is often the case, statistical estimators for noise parameters give merely approximate solutions, involving models with a complicated form, polynomial equations and numerical integration at every pixel. The solution is computationally very expensive.

*peter.schallauer@joanneum.at; phone +43 316 876-1202; fax +43 316 876-1191; www.joanneum.at/iis/en

The patch-based approach presented in [2] proposes a non-parametric method for pixel-wise texture synthesis. To synthesize a single pixel, the algorithm first gathers regions in the sample image with small perceptual distance to the single pixel's neighbourhood. It selects one of the regions and uses its centre to be the new synthesized pixel in the context of a Markov Random Field. This works well on a large range of textures, but the success of the algorithm is dependent on the correct choice of neighbourhood size. The algorithm tends to replicate local sample image content, which is inappropriate for our intended grain synthesis application. Computational cost is high because an entire search of the sample image is necessary for each of the pixel to be synthesized.

The work in [3] synthesizes textures using features of a combination of Laplacian and steerable pyramids [6] applied to an input texture template. However, the histogram matching between synthesized noise image and template image, which is used in each of the corresponding pyramid subbands, does not add enough spatial correlation to capture the qualities of film grain. The issue of spatial grey value distributions is not sufficiently addressed by the grey level histogram matching approach. The lack of spatial correlation leads to results which contain additional high frequency image content when applied to film grain synthesis.

To overcome this limitation, Portilla and Simoncelli ([5]) improved the approach by introducing joint higher-order statistical texture measure for feature matching between synthesized grain image and template image.

Our approach is based on a modification of this texture synthesis algorithm (cf. Section 2). This paper investigates how it fulfils the stated grain synthesis requirements for restoration and special effects compositing applications.

In Section 2 we illustrate our instantiation of the texture synthesis algorithm; Section 3 presents an application framework for film grain processing, followed by results in Section 4 and discussion in Section 5.

2. FILM GRAIN SYNTHESIS

The grain synthesis algorithm presented in this paper is based on a work of Portilla and Simoncelli ([5]). First, the grain template image is decomposed into a steerable pyramid, a linear, multi-scale and multi-orientation image transform. Each scale and orientation of the pyramid is analyzed with respect to several statistical texture features. The subsequent process of synthesis is initialized with an image of Gaussian white noise, onto which these features are iteratively imposed in the matching process. Their work introduces the following texture features: a) marginal statistics (image moments mean, variance, skewness, kurtosis, minimum and maximum grey value), b) correlation of subbands, c) correlation of magnitude responses and d) the relative local phase. The texture synthesis algorithm can be influenced by choosing a subset of these features in order to constrain the matching process.

For synthesizing film grain we have evaluated different subsets of constraints. Result images in Figure 1 illustrate that marginal statistics are not sufficient to capture the versatile structure of film grain. But they perform well in combination with the correlation measurements of subbands, as this feature models periodic structures of film grain. The figure also shows that the usage of all constraints does not contribute to a better visual appearance of the synthesized grain. On the contrary, adding the relative local phase may lead to undesired finer image structures.

Thus we selected the marginal statistics and the correlation of subbands constraints, but have omitted the remaining constraints. The correlation of magnitude responses is left out, as there is no necessity for representing structures like edges, bars and corners in the grain pattern. Similarly, we have not applied the relative local phase constraint, as it models features that are disadvantageous in our case, such as gradients due to shading and lightning effects and local features like edges and lines. Furthermore these constraints would significantly increase the computational complexity.

The requirements for film grain synthesis are highly met by this implementation: the fact that the synthesis starts with random noise assures high spatio-temporal variation. Multi spectral image support is implemented by applying the grain synthesis to the R, G, and B channel. In the RGB color model film grain is spatially correlated between the channels, which is reflected during the synthesis process by using the same random noise image as start image for the texture statistics matching procedure. The synthesized tiles are created in a way that they can be expanded seamlessly, e.g. a large grain pattern can be obtained by a block-wise copying of a smaller synthesized grain sample.

Grain template images

Results of [5], using marginal statistics (a) only

Results of [5], using preferred subset of constraints (a, b)

Results of [5], using all constraints (a, b, c, d)

Results of [5], using subset (c, d)

Results of [5], using subset (a, b, d)

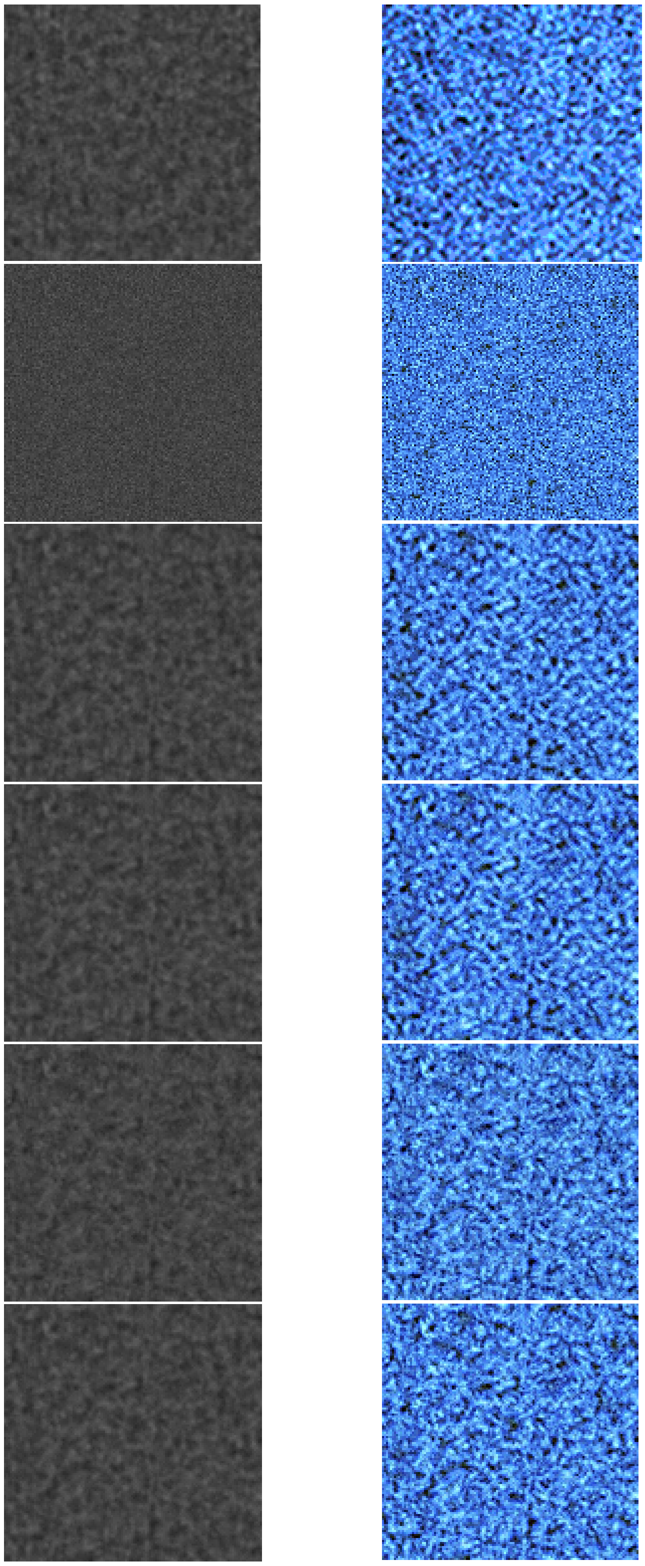

Figure 1: Grain synthesis results of [5] when using different constraints: (a) marginal statistics, (b) correlation of subbands, (c) correlation of magnitude responses and (d) relative local phase. For all results the same start noise image for synthesis is used.

3. FRAMEWORK FOR TEMPLATE-BASED RE-GRAINING

The proposed film grain synthesis procedure is part of a framework for template based re-graining. This framework enables a number of re-graining applications. Examples include:

- authentic missing grain reconstruction of image areas produced by restoration algorithms in film and video
- media graining either for computer generated and composited image content or for the reconstruction of movies out of different copies.

In the framework, shown in Figure 2, the principal information carriers for film grain are the grain template image and the parameters describing the level of grain and signal dependency. The extraction of a suitable grain template is possible by manual or automatic procedures ([4]), depending on the desired application. The grain synthesis, which is exclusively based on the grain template, generates synthetic grain, which is perceptually similar to the input grain in sufficient spatial and temporal variation and in arbitrary size. In the grainer the image signal is superimposed by the synthesized grain intensity pattern under consideration of a given signal dependency, expressed by:

$$f'(x,y) = f(x,y) + n(x,y)\,w(f_s(x,y)),$$

where *f'* refers to the resulting grained image, *f* to the input image signal, *n* denotes the synthesized grain noise pattern and the weighting function *w* models the signal dependency. In order to be robust against potential noise and grain in the input image signal, a smooth operator is applied before f_s is calculated.
The input image signal for the grainer is application specific. In the media grainer application, a clean, e.g. computer generated image is passed without modification, whereas in the missing grain restoration application, the input image is at first submitted to a defect removal process. These missing data reconstruction techniques, e.g. entire frame interpolation or spot removal, can cause unnatural looking areas, with grain smoothed or partly suppressed, which are then targeted by the missing grain restoration. The grain detector facilitates unsupervised missing grain restoration by automatically providing the grain template, the grain level and the signal dependency.

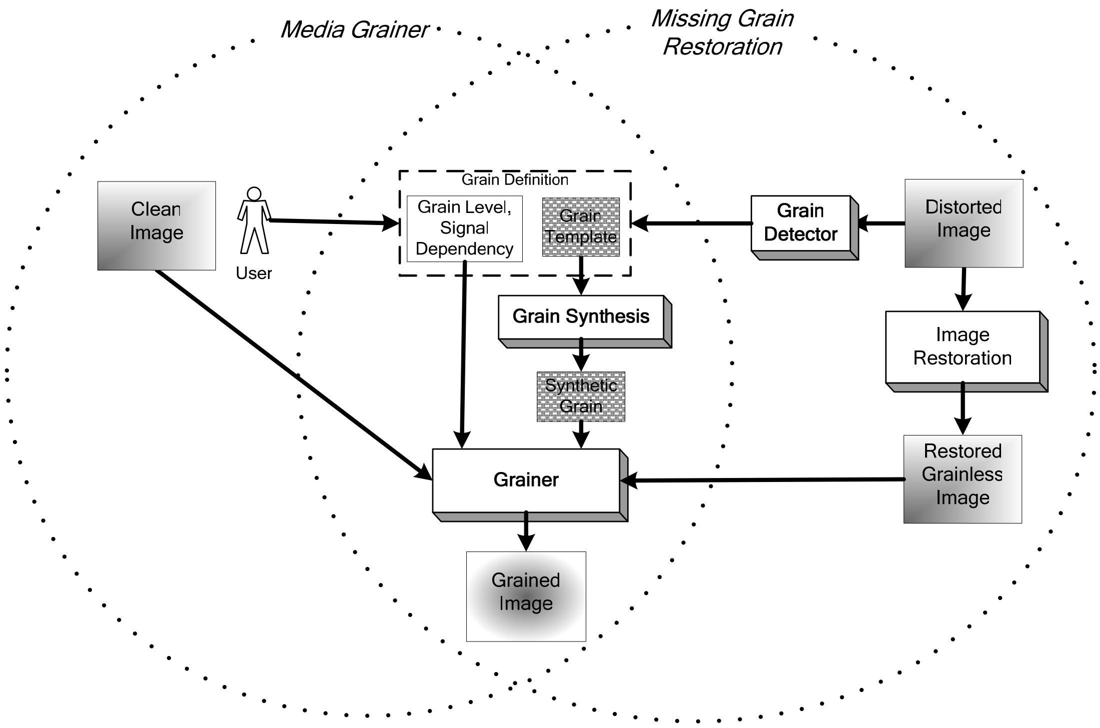

Figure 2: Application framework for template-based re-graining.

Information about the respective film grain or noise is exclusively defined in the grain template and the grain level and signal dependency parameters. This modular interface enables the setup of diverse grain applications, be it fully automated or manually driven.

4. RESULTS

Our current C++ implementation using Intel IPP and IPL requires roughly 600ms on a 3 GHz Pentium workstation for the results shown in Figure 3 and Figure 4. About 20% of run-time are due to the analysis step (building steerable pyramids with 4 scales, 4 orientations and a 9x9 neighbourhood) and 80% are spent in the synthesis step (synthesising a 256x256 grain pattern in 2 iterations). The synthesized grain image can be seamlessly repeated to obtain a grain pattern of arbitrary size.

Figure 3 shows typical grain template images and the corresponding synthesized grain. Subfigure 3d) displays electronic noise images. In this regard, it is worth mentioning that the proposed synthesis algorithm operates on film grain as well as on finer electronic noise. Whereas the results in this figure are homogenous and free of visible artifacts, the results demonstrated in Figure 5 obviously are not. There, the visible low frequent structures in the template are carried forward to the synthesized image, resulting in an inhomogeneous grain pattern.

Figure 5 exemplifies the missing grain restoration application. In this example dirt is removed by temporal interpolation only inside the image region indicated by the rectangle. When re-graining is not applied, due to the interpolation algorithm the output grain structure is suppressed. The result of fully automatic re-graining is shown in Subfigure 5c).

Figure 6 visualizes media graining of the clean Lena image. The re-grained image exhibits the grain pattern of the provided template. As to the signal dependent compositing of synthesized and original Lena image, more graininess is exhibited in the mid-grey areas and less in the very dark or bright tones.

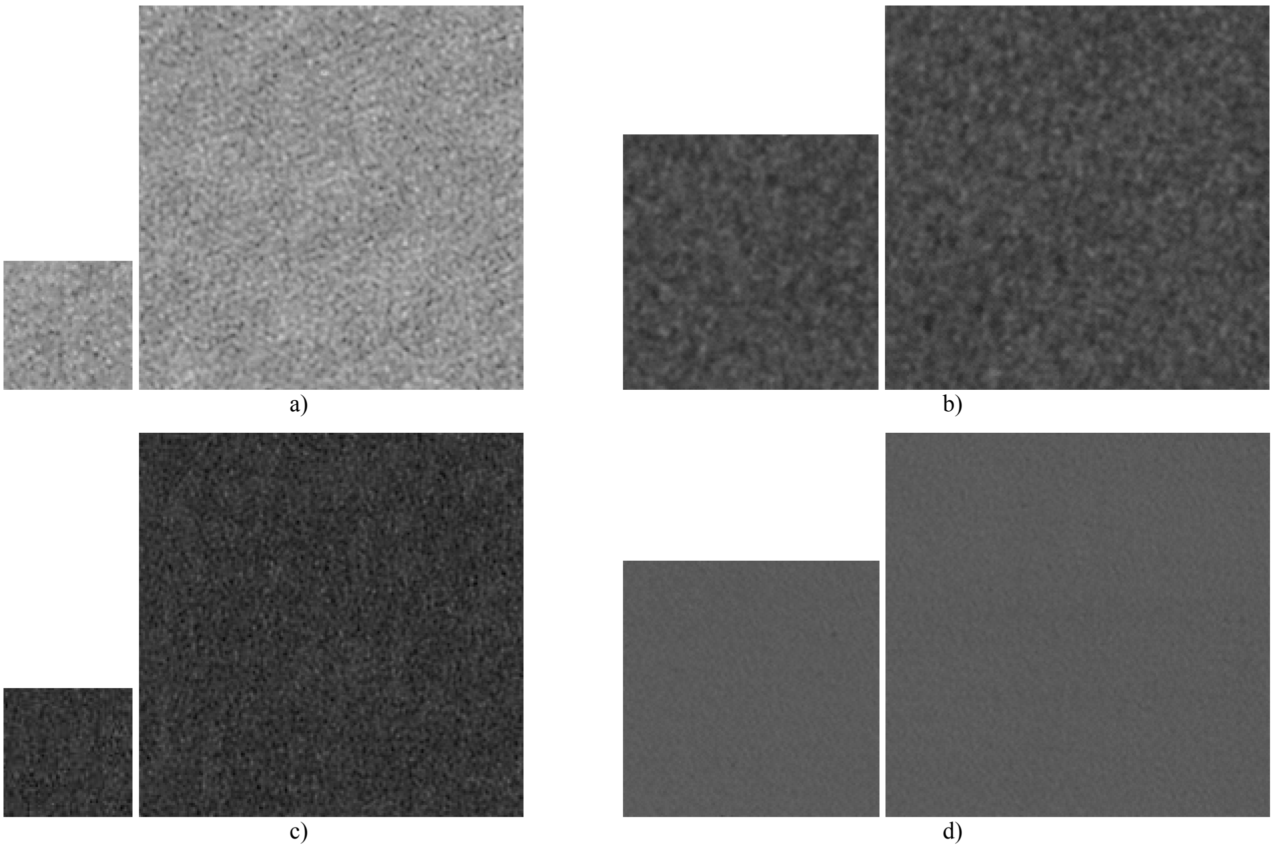

Figure 3: Results of grain synthesis with constraints proposed in Section 2. In each image pair the left image shows the grain template (64 or 128 pixel width) and the right image shows the generated synthetic grain (192 pixel width).

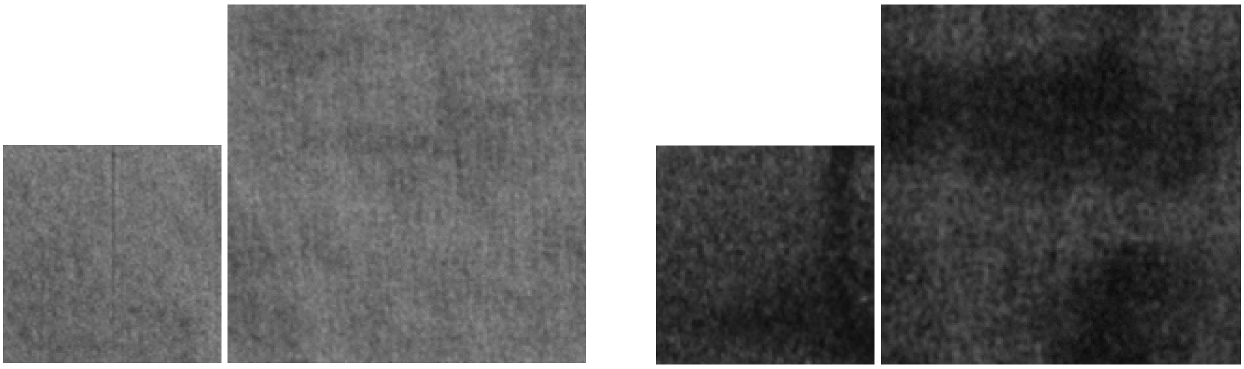

a) b)

Figure 4: Results of grain synthesis with constraints proposed in Section 2. Template images (left image of a) and b)) and resulting synthetic grain images (right image of a) and b)) contain inhomogeneous image content.

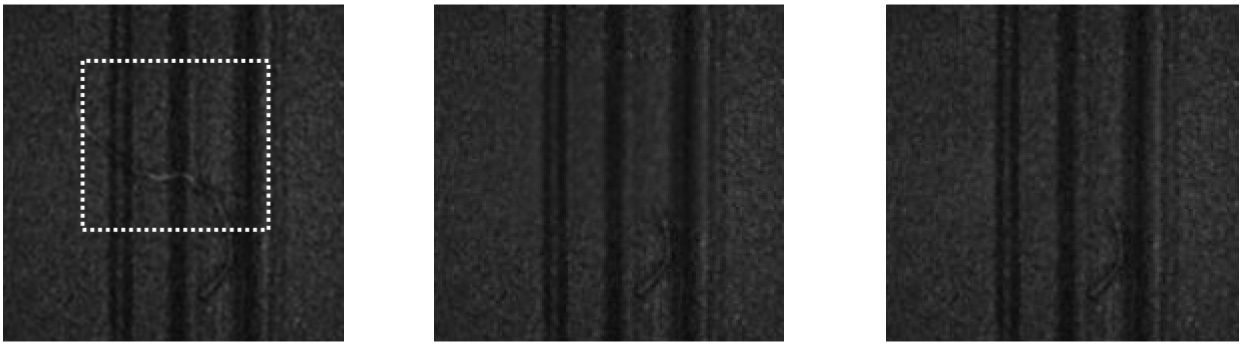

a) b) c)

Figure 5: Example of missing grain restoration application, a) original image containing dirt, the region indicated by the rectangle is restored in b) and c) images, b) interpolated dirt removed image content without re-graining, c) interpolated dirt removed image content with fully automatic re-graining.

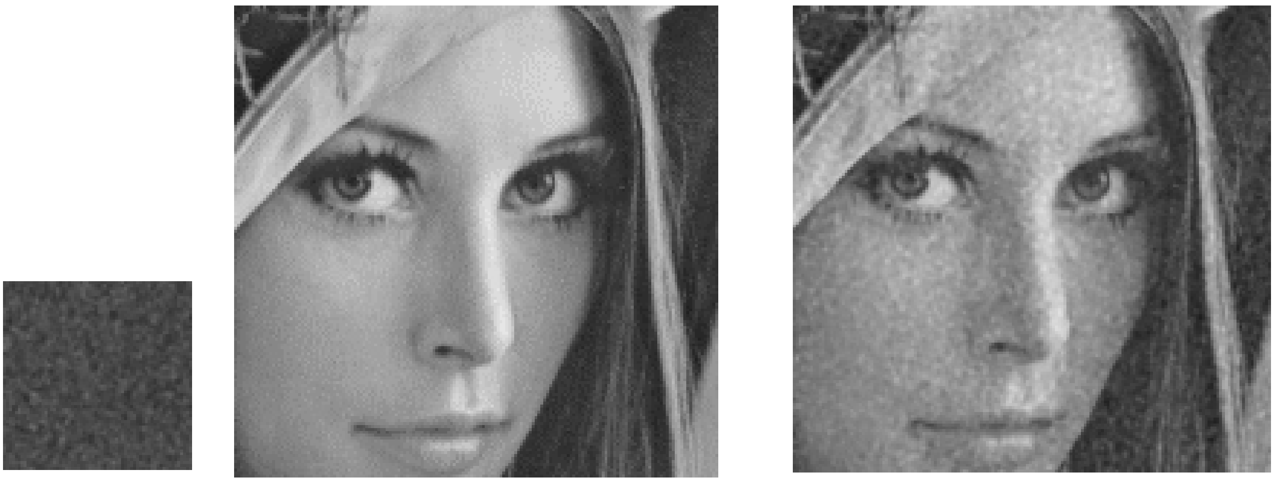

a) b) c)

Figure 6: Example of media grainer application; a) grain template image, b) Lena without re-grain, c) Lena re-grained with grain template a).

5. CONCLUSION AND PERSPECTIVES

The grain applications dealt with in this paper are the *missing grain restoration* of defect removed image areas and *media graining* either of entire image sequences or of composited computer graphics content. Both tasks require the generation of artificial grain perceptually similar to a certain grain at hand in spatial and temporal variation and expansion capability to arbitrary size. In order to ensure widespread use, the algorithm must provide multi-spectral image support and must be of low computational cost.

This work presents an adaptation of the parametric texture model approach proposed in [5] specifically for film grain synthesis applications. Provided that the desired grain pattern is specified by means of a template image, the algorithm produces synthetic grain which matches the template very well. The random noise based approach inherently provides superb spatial and temporal variation. The complexity of the method is reduced by a targeted selection of the texture model parameters, so allowing deployment of the method in industrial grade software.

Unfavourable regular structures in the synthesized grain pattern can be introduced if the synthesized grain tile is small, that is, about 64x64 pixels or less or if low frequent structures are contained in the template pattern. While enlarging the synthesized grain tile size only results in higher computational effort, the low frequent information contained in the input template definitely requires future work, in order to provide stable fully automated re-graining.

Thanks to its modular design, the proposed framework for template based re-graining enables a number of re-graining applications, such as the user-operated *media grainer* and the unsupervised *missing grain restoration*. The key elements grain synthesis, grain template, grain level, signal dependency and the module for adding grain and signal form the basis for both applications, which are implemented and successfully applied in a digital film restoration software ([1]).

Further applications of the framework are considered, e.g. for boosting the efficiency of the digital cinema and broadcast delivery process. At the transmitter side, the content is grain suppressed and the grain properties (template image and signal dependency) are automatically detected and delivered jointly with the content. To obtain a result exhibiting the same graininess at the receiver side, the grain suppressed image is re-grained using the template. This method enables higher compression in the delivery process, requiring less transmission bandwidth and providing therefore cheaper delivery with same quality.

ACKNOWLEDGEMENTS

The authors would like to thank Georg Thallinger, Herwig Rehatschek and Werner Haas as well as several other colleagues at JOANNEUM RESEARCH, who contributed valuable feedback to the work. In this regard special thanks go to Michal Burian for his essential efforts in implementation and evaluation of the grain synthesis methods. This work has been funded partially under the 6th Framework Programme of the European Union within the IST project "PrestoSpace" (IST FP6-507336, http://www.prestospace.org).

REFERENCES

1. DIAMANT - Digital Film Manipulation System. http://www.joanneum.at/en/informatik/projekte_detail.php?p_iid=IIS&p_pid=102
2. A. Efros and T. Leung, “Texture Synthesis by Non-parametric Sampling”, *Proceedings of the Seventh International Conference on Computer Vision*, Corfu, Greece, 1999.
3. D. J. Heeger and J. R. Bergen, “Pyramid-Based Texture Analysis/Synthesis”, *Proceedings of SIGGRAPH ’95*, pp. 229-238, 1995.
4. R. Mörzinger, *Detection of Grain and Noise For Film and Video Re-graining*, Diploma Thesis, Upper Austrian University of Applied Sciences, Hagenberg, Austria, 2005.
5. J. Portilla and E. P. Simoncelli, “A Parametric Texture Model based on Joint Statistics of Complex Wavelet Coefficients”, *Journal of Computer Vision,* October 2000.
6. E. P. Simoncelli and W. T. Freeman, “The Steerable Pyramid: A Flexible Architecture for Multi-scale Derivative Computation”, *Proceedings of 2nd Annual IEEE International Conference on Image Processing*, Washington, DC, October 1995.
7. J. Yan and D. Hatzinakos, “Signal-dependent film grain noise removal and generation based on higher-order statistics”, *Proceedings of the IEEE Signal Processing Workshop on Higher-Order Statistics*, pp. 77-81, Banff, Canada, July 1997.