

# Journal of Electronic Imaging

JElectronicImaging.org

## Advanced tools and framework for historical film restoration

Simone Croci  
Tunç Ozan Aydın  
Nikolce Stefanoski  
Markus Gross  
Aljosa Smolic

**SPIE**•



Simone Croci, Tunç Ozan Aydın, Nikolce Stefanoski, Markus Gross, Aljosa Smolic,  
“Advanced tools and framework for historical film restoration,” *J. Electron.  
Imaging* **26**(1), 011021 (2017), doi: 10.1117/1.JEI.26.1.011021.

# Advanced tools and framework for historical film restoration

Simone Croci,<sup>a,\*</sup> Tunç Ozan Aydın,<sup>a</sup> Nikolce Stefanoski,<sup>b</sup> Markus Gross,<sup>a,c</sup> and Aljosa Smolic<sup>a,d</sup>

<sup>a</sup>Disney Research Zurich, Stampfenbachstrasse 48, Zurich 8006, Switzerland

<sup>b</sup>unigFEED AG, Technoparkstrasse 1, Zurich 8005, Switzerland

<sup>c</sup>Swiss Federal Institute of Technology, Raemistrasse 101, Zurich 8092, Switzerland

<sup>d</sup>Trinity College, College Green, Dublin 2, Ireland

**Abstract.** Digital restoration of film content that has historical value is crucial for the preservation of cultural heritage. Also, digital restoration is not only a relevant application area of various video processing technologies that have been developed in computer graphics literature but also involves a multitude of unresolved research challenges. Currently, the digital restoration workflow is highly labor intensive and often heavily relies on expert knowledge. We revisit some key steps of this workflow and propose semiautomatic methods for performing them. To do that we build upon state-of-the-art video processing techniques by adding the components necessary for enabling (i) restoration of chemically degraded colors of the film stock, (ii) removal of excessive film grain through spatiotemporal filtering, and (iii) contrast recovery by transferring contrast from the negative film stock to the positive. We show that when applied individually our tools produce compelling results and when applied in concert significantly improve the degraded input content. Building on a conceptual framework of film restoration ensures the best possible combination of tools and use of available materials. ©2017 SPIE and IS&T [DOI: 10.1117/1.JEI.26.1.011021]

Keywords: digital film restoration; image processing; filtering; contrast enhancement.

Paper 16440SS received May 27, 2016; accepted for publication Dec. 12, 2016; published online Jan. 13, 2017.

## 1 Introduction

Until the relatively recent transition to digital media, film has been the primary medium for storing visual content for a long time. Some of this content is being preserved in archives due to their historical or artistic value, or their importance for other reasons. In fact, in some countries, the film content is perceived as a part of the cultural heritage, therefore, significant investments have been made for film archival. Another related effort is the digital restoration of the legacy content, such that it can be transferred and viewed in current hardware and, therefore, becomes available to a wider audience.

The digital restoration process involves scanning the film stock and accounting for various types of chemical and physical degradations without modifying the original artistic intent of the content. All these steps require specialized tools and knowledge, and as such the digital restoration process is often quite time consuming and labor intensive. Therefore, only a small percentage of the legacy content can be processed per year. While in some rare cases, a title's commercial value can justify its restoration costs (e.g., Disney's recently rereleased *Lion King*), other less popular but nevertheless culturally significant titles may become irreversibly destroyed before they are ever restored.

Interestingly, a number of research problems that are relevant to many of the challenges in digital film restoration have been investigated in computer graphics literature for decades. The restoration of chemically degraded film colors (due to humidity, heat, etc.) is closely related to the body of work in color transfer. The removal of excessive film grain is essentially a specialized instance of denoising and video

filtering. Similarly, restoration of the contrast of the positive film stock from the negative stock can be viewed as a special form of contrast enhancement through joint filtering. Despite the resemblances, however, we found that off-the-shelf computer graphic methods still require overcoming a multitude of research problems in order to be applicable to digital film restoration.

The goal of this interdisciplinary work is to present a set of semiautomatic key components for digital restoration with the hope of removing the barriers to a more practical restoration process. To that end, we start from the state-of-the-art in the computer graphics literature and make multiple contributions to enable their use for the purposes of digital film restoration. Specifically, we make the following contributions:

- We present practical digital film restoration tools for restoration of chemically degraded colors, excessive grain removal, and contrast transfer from negative stock, all by building upon the state-of-the-art in the computer graphics literature by adding components.
- We demonstrate the usefulness of our techniques on a multitude of case studies by restoring degraded film content.

We base our discussions on a conceptual framework of film restoration outlined in Sec. 3, which sets the foundation for optimum use of tools and available film materials. In the next section, we first start with a brief summary of the related work.

\*Address all correspondence to: Simone Croci, E-mail: [simone.croci@diastor@gmail.com](mailto:simone.croci@diastor@gmail.com)

## 2 Related Work

### 2.1 Film Restoration

For an overview on the fundamentals of film restoration, we refer the reader to Read and Meyer.<sup>1</sup> Digital restoration was further investigated by Kokaram.<sup>2</sup> Other related work discussed common degradations such as grain noise, dirt and sparkle, shake, and flicker, while also proposing a Bayesian approach for their removal. Rosenthaler and Gschwind<sup>3</sup> presented a model of film color degradation as well as a method for removal of scratches, dust, and other defects. Schallauer et al.<sup>4</sup> presented automatic algorithms for combating problems such as dust, dirt, image stabilization, flicker, and mold. Finally, Werlberger et al.<sup>5</sup> presented an optical flow guided framework for interpolation and noise removal for television content.

### 2.2 Color Manipulation

Two fundamental works in the area of film color restoration are by Gschwind and Frey.<sup>6,7</sup> These papers describe a simple but effective physical model of color degradation occurring in films. In our work, we build upon the work of Oskam et al.<sup>8</sup> that utilizes a radial basis function interpolation for balancing the colors between images, and was originally developed for virtual reality applications. We also studied other interpolation models that are described in Amidror et al.<sup>9</sup> Other relevant work in color transfer utilizes a moving least squares approach,<sup>10</sup> performs color transfer in two steps: nonlinear channel-wise mapping followed by cross-channel mapping,<sup>11</sup> relies on statistics in color images,<sup>12,13</sup> attempts to preserve the gradients of the source image,<sup>14</sup> tries to match the color distribution of the target image,<sup>15</sup> and performs histogram matching between the source and target images.<sup>16</sup>

### 2.3 Denoising

Denoising has been one of the most fundamental and practically relevant problems in visual computing and has been studied extensively. For a theoretical treatment for the filtering process involved in denoising, we refer the reader to Milanfar.<sup>17</sup> A popular denoising approach is the nonlocal-means (NLS) filtering,<sup>18,19</sup> which has been further discussed in comparison to other methods in Buades et al.<sup>20</sup> More recently, the performance of the NLS filter has been improved through adaptive manifolds.<sup>21</sup> Other work that is applicable to denoising includes the domain transform filter,<sup>22</sup> the permutohedral lattice,<sup>23</sup> the weighted-least-squares filter,<sup>24</sup> the bilateral filter,<sup>25</sup> and the BM3D filter.<sup>26</sup> Another relevant line of research to denoising is edge-aware spatiotemporal filtering.<sup>27,28</sup> Especially, Aydın et al.<sup>28</sup> demonstrated compelling denoising results for HDR video by relying on simple averaging over motion compensated spatiotemporal video volumes.

### 2.4 Contrast Transfer

A more general form of contrast transfer, namely, style transfer has been thoroughly investigated in the computer graphics community. Some recent examples include Shih et al.,<sup>29</sup> where they transfer the contrast and various other properties between headshot portraits. Wang et al.<sup>30</sup> transfer tone and color properties of photographs taken with high-end

DSLR cameras to photographs taken with cell phone cameras with lower image quality. Another interesting style transfer method based on learning the user's preferences has been presented by Akyuz et al.<sup>31</sup> specifically for HDR images. In our case, the style transfer occurs from the film negative to the film positive, where both the positive and negative can be chemically and physically degraded at various degrees.

## 3 Digital Restoration Framework

Analog film has been used for over 125 years to tell stories and to communicate visual information. Although film has almost completely been replaced by digital video today, some film makers still choose to produce their movies on film due to its specific visual style. On the other hand, even for content shots on digital media, film is still a viable long-term alternative for archiving. Analog film technology from capture to display evolved tremendously over the decades, resulting in a vast amount of different types of film stocks and the processing techniques associated with them. A comprehensive review of historical film technologies with the main focus on color processes is presented by Flueckinger.<sup>32</sup>

The major steps of a digital restoration workflow are illustrated in Fig. 1. Typically analog film is captured on a negative  $N_O$  where  $O$  stands for original. In our simplified model (omitting, e.g., interpositives, etc.) the original positives  $P_O$  that are used for projection in cinemas are created from  $N_O$ . Generation of the positives usually involves color grading  $f_G$ , which is an artistic process for establishing the final look of the content on screen. As  $N_O$  and  $P_O$  are physical entities, they will undergo deformations over the years and decades resulting from mechanical/physical treatment and chemical processes inside the material. The negative is usually only used to create a few copies and otherwise is preserved carefully, whereas a positive may have run thousands of times through a projector and be exposed to various types of outside influences. A negative is typically in better shape but does not carry the grading information, which is a historical part of the final visual art. Furthermore,  $N_O$  and  $P_O$  are made of very different materials such that the chemical degradation processes may be completely different depending on the particular type of film stock. We formally express all these effects as two different degradation functions  $f_N$

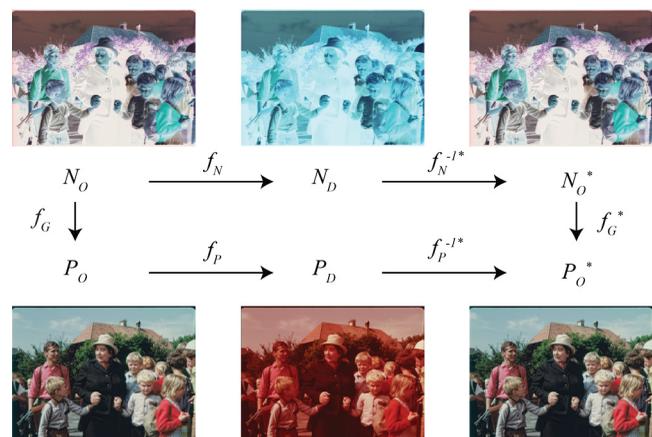


Fig. 1 A conceptual framework of digital film restoration.

and  $f_P$ , which transform the originals into their degraded versions  $N_D$  and  $P_D$ .

$N_D$  and  $P_D$  are degradations in the physical medium and should be removed. Importantly, the end result of this process should still remain faithful to the restored content's original state. The restoration may involve mechanical cleaning and repairing of pieces as sometimes only fragments are available. Additionally, often only either  $N_D$  or  $P_D$  can be obtained, and not both. After these initial steps, the film stock is digitized by a film scanner, which could introduce slight color shifts due to the scanner's inaccuracy. Conceptionally, we subsume these effects into the degrading functions  $f_N$  and  $f_P$ , and assume that  $N_D$  and  $P_D$  are final digitized content.

The goal of digital film restoration is the faithful reconstruction of the original visual information from the degraded content  $N_D$  and  $P_D$ . Conceptionally, this can be defined as estimation of inverse degrading functions, i.e., reconstruction functions  $f_N^{-1*}$  and  $f_P^{-1*}$  that would create estimates of the original inputs  $N_O^*$  and  $P_O^*$ . From this, we can estimate the original color grading function  $f_G^*$ .

The goal of our research was to improve film restoration for two practically relevant scenarios. The first scenario assumes that the positive stock is available, but the negative stock is either not available or is somewhat degraded because of not being well preserved. For this scenario, we present three methods for performing the improvements conceptualized by the reconstruction functions  $f_N^{-1*}$  and  $f_P^{-1*}$ . Specifically, we address the restoration of degraded colors, removal of excessive film grain, and (in case the negative stock is available) the contrast transfer from the negative to the positive.

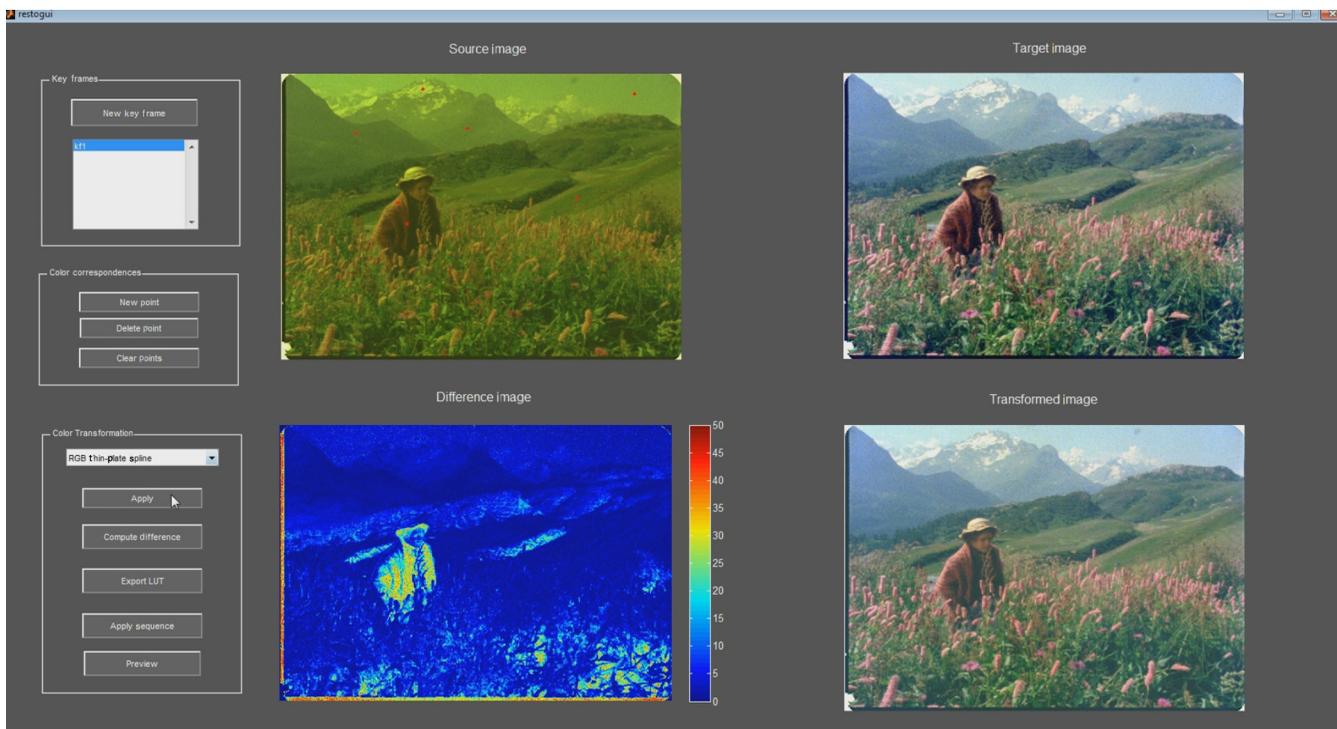
In the second scenario, we assume that both positive and negative stocks are available, and the negative stock is in

a good condition. In this case, we use the scans of the positive and negative films to estimate the color grading function  $f_G^*$  using the same color restoration method applied in the first scenario (Sec. 4.1). Then, we simply apply the color grading function to the negative ( $N_O^*$ ) in order to obtain a positive ( $P_O^*$ ) that has a similar level of quality to the negative, but is also faithful to the artistic vision estimated from the positive.

While the aforementioned scenarios are common in practice, every film restoration project is a unique exercise and can impose challenges. However, the methods presented in the remainder of this paper are useful additions to the currently available set of restoration tools, and can be used in tandem with others to facilitate the digital restoration process under various circumstances.

## 4 Advanced Tools

In our conceptual framework, any restoration method can be a part of the reconstruction functions  $f_N^{-1*}$  and  $f_P^{-1*}$ . In this section, we present three such tools which build upon the state-of-the-art in visual computing. In order to obtain the results presented in Sec. 5, we applied these methods in the presented order. However, they can be used in any order, standalone, or in combination with any other restoration tools and software (e.g., scratch removal, stabilization, etc.). In addition to the development of tools, another focus of our work was on the development of related user interfaces, which allow intuitive control over relevant parameters of the algorithms. An example of user interaction is shown in Fig. 2 (Video 1). These user interfaces enabled testing and optimization during algorithm development but also show how the tools can be integrated into professional products and further be used in real restoration projects.



**Fig. 2** User interaction with the color restoration tool via our interface. Video 1 (MP4, 10 MB) [URL: <http://dx.doi.org/10.1117/1.JEI.XX.XX.XXXXXX.1>]. (Please note that the large file size may result in a long download time.)

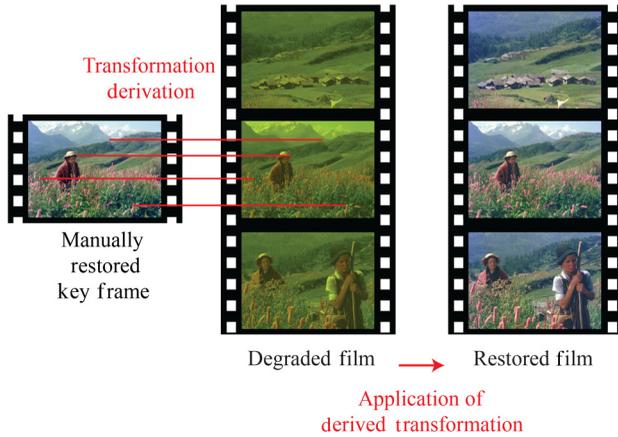


Fig. 3 Color restoration by color transformation estimation.

#### 4.1 Color Restoration

Color degradation is caused by chemical processes inside the film stock. (Fig. 3 shows an example.) These processes vary largely for different types of film stock and often even affect the color components differently. For instance, sometimes the blue component fades faster than others, causing a characteristic color shift. In some cases for certain types of film stock, it is possible to model the chemical degradation process and to derive corresponding algorithms for inversion as described in Refs. 6 and 7. However, such approaches require physical measurements of the film stock to derive appropriate model parameters and in the end still require interactive correction and finishing.

In our approach, we rather rely on color transfer algorithms that have proven to have high quality and flexibility in other application domains.<sup>8</sup> We make the assumption that the color degradation is a global transformation, which means that the degradation only depends on the colors of the original film and not on the position of the colors in the frame. This assumption is justified due to the fact that the evident color degradation is caused by a chemical process that affects the images globally.<sup>6</sup>

##### 4.1.1 Color restoration pipeline

The pipeline of our approach for color restoration is illustrated in Fig. 3. The key is a color transformation which is used to transform the degraded film. This transformation is generated in an interactive process controlled by the conservator, where the conservator provides several restored keyframes to our system. The number and spacing of keyframes depend on the length of a sequence and the amount of change. In our results, we mostly used just one keyframe for sequences of several seconds. Rarely a second keyframe was used. The conservator restores the keyframes using traditional grading tools. Then our system will estimate a color transformation of that particular grading by comparing the restored keyframe to the input as described below (indicated by red lines in Fig. 3). This color transformation is then applied to the whole sequence to create the final restored output.

##### 4.1.2 Estimation of color transformations

The estimation of the color transformation that was used to manually restore a keyframe follows the approach of Oskam

et al.<sup>8</sup> The task is to compute the transformation parameters according to some optimality criteria. Let us assume that the equation  $c' = f(c, p)$  defines the transformation  $f$  that takes as input a faded color  $c$  and computes as output the corresponding original color  $c'$ .  $f$  depends on  $p$  which represents the parameters that must be computed before using the transformation to correct other frames. For the computation of the parameter  $p$ , we use a least squares criterion defined as

$$\arg \min_p \sum_{i=1}^N \|c'_i - f(c_i, p)\|^2, \quad (1)$$

where  $(c_i, c'_i)$  ( $i = 1, \dots, N$ ) are color correspondences, i.e., pairs of restored and corresponding faded colors (Fig. 3). The least squares criterion simply minimizes the sum of the squared transformation errors for the given color correspondences and is computed by an appropriate solver. The color correspondences needed for transformation training can be obtained manually or automatically as outlined in the next section.

$f$  in Eq. (1) can be considered as an interpolation function that creates the restored output images from degraded input images, depending on a parameter set  $p$  as computed from the color correspondences of the keyframe. In principle, there are different options of interpolation functions  $f$ , such as polynomial, radial basis functions, thin-plate splines, moving least squares, or tetrahedral. Our user interface allows selection of all these options (see Fig. 2, Video 1). Figure 4 shows the results obtained using the previously mentioned interpolation functions based on the same color correspondences. In this example, the thin-plate spline function performs best, whereas the least sophisticated polynomial interpolation function achieves the worst result. Generally, we found for various examples that thin-plate spline interpolation provides very good results. It is formally defined by the following equation:

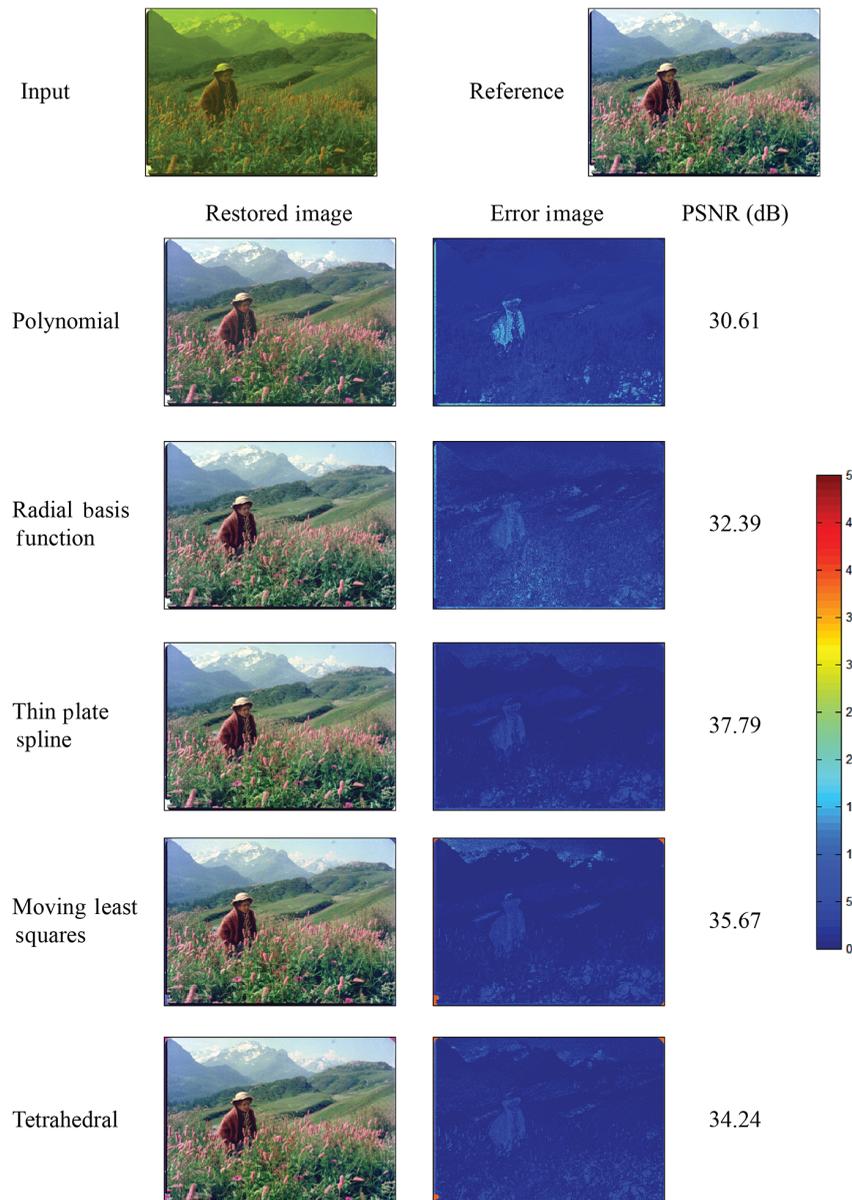
$$c' = Ac + b + \sum_{i=1}^N \phi(\|c - c_i\|)w_i, \quad (2)$$

where  $\phi(r) = r^2 \log(r)$ ,

and  $p := (A, b, w_1, \dots, w_N)$  are the parameters of the transformation.  $A$  is a  $3 \times 3$  matrix, whereas  $b$  and  $w_i$  are three-dimensional (3-D) vectors.

##### 4.1.3 Color correspondences extraction methods

For the selection of the correspondences needed to compute the parameters of the color reconstruction function, we developed an interactive and an automatic method. The interactive approach provides full control to the conservator. It follows Oskam et al.<sup>8</sup> and is shown in Fig. 5. First, the user selects some initial correspondences (usually 3 or 4) by clicking into the input image, typically inside characteristically colored regions. The system then solves the equations above and creates a first restored version, which is usually still very far off. Then the user continues adding correspondences in selected image regions where the error is still large, and the process is repeated until convergence and satisfaction of the user. This is typically the case after about 10 iterations depending on the complexity of the



**Fig. 4** Comparison of different interpolation functions. The error images illustrate the per pixel difference between the manually restored reference and the different restorations by color transformation from the same correspondences. The PSNR values in decibels provide a global measure of quality.

content. Figure 2 (Video 1) shows the user interaction with the color restoration tool. The user incrementally adds color correspondences until the color transformation is trained well enough. We developed an interface where the user can control properties and parameters of the algorithm as needed.

Our automatic approach is based on clustering of random correspondences and is shown in Fig. 6. The method randomly selects a predefined number of correspondences distributed over the image. We used 10,000 points for the results reported in this paper. This step reduces the number of samples to be considered, while still providing a representative set of the color distribution. Next, the set is further reduced to a number of typically 30 color clusters using a  $k$ -means algorithm in 3-D color space, to make the problem computationally feasible, while still providing a good representation. Finally, for each cluster, we select the correspondence which is closest to the centroid as representative, and these

representatives are used to solve for the color restoration transformation.

For most of our results reported in this paper and in Fig. 7 (Video 2), we used the automatic approach. For a few cases, we decided to fall back to the interactive approach when the material was too degraded. This limitation can always be handled in the interactive restoration workflow, where the automatic method serves as an additional tool which in most cases increases efficiency.

#### 4.2 Denoising/Degraining

Denoising is one of the most studied areas in image/video processing research. We extend a recently proposed edge-aware spatiotemporal filtering method based on permeability-guided filtering (PGF).<sup>28</sup> Here, we present a specific improvement for the application to film restoration.

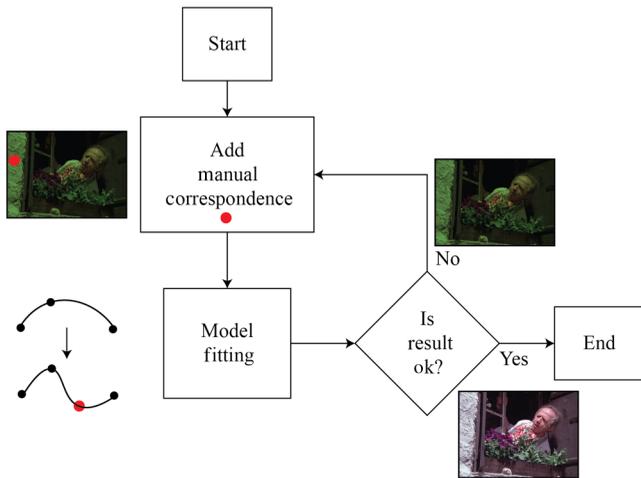


Fig. 5 Manual extraction of correspondences.

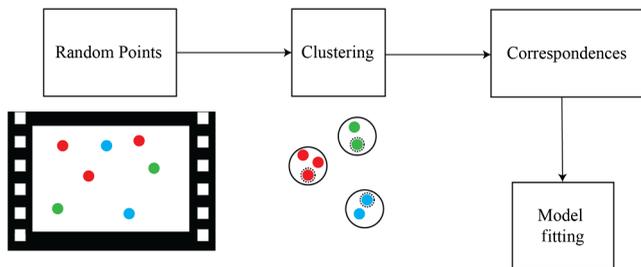


Fig. 6 Automatic extraction of correspondences.

#### 4.2.1 Grain noise characteristics

Exposed film always contains characteristic grain noise which creates the specific look that is considered as part of a piece of art. However, old degraded film usually exhibits a significant increase of film grain which is not desired, i.e., it is an artifact (Fig. 8). A major difficulty in film restoration is, therefore, to find the right amount of reduction to an approximately original level and nature of noise. The conservator

has to find this balance with the help of a powerful tool such as ours as described in the following, which also allows the retaining of texture and contrast of the content as unaffected as possible.

To characterize the grain properties of degraded film, we analyzed a patch of the sky region of the original image content shown in Fig. 8, which appears to show a uniformly colored region distorted with strong grain noise. The colored artifacts indicate that grain noise degraded differently in the color channels. A look at the energy spectral density of the luma channel of this patch (Fig. 9) shows that the spectral energy is isotropically distributed in the low and mid range of frequencies. The isotropic distribution indicates that grain noise has no direction of preference, whereas the lack of high-frequency energy reveals that grain noise is locally correlated, i.e., it is smooth to a certain degree. A closer look at the distribution of absolute  $x$ - and  $y$ -derivatives (Fig. 10) shows that they are concentrated in a small domain in comparison to the original image domain with values up to 255. Hence, we characterize nonstructured image regions with grain noise as having isotropically distributed noise and small  $x$ - and  $y$ -derivatives.

#### 4.2.2 Grain reduction

We reduce the grain noise by using PGF,<sup>28</sup> which we extend by employing permeability weights that take the characteristics of grain noise into account. The weights are computed adaptively from the available noisy image content to enable a strong spatial diffusion in areas with isotropic image content and small partial derivatives in  $x$ - and  $y$ -directions, while in areas deviating from this assumptions the level of diffusion is gradually reduced.

PGF is conducted by filtering with a one-dimensional (1-D) filter in a separable way along image rows, image columns, and in time along optical flow trajectories. First, all image rows are filtered which gives image  $J^{(1)}$  and then all image columns of image  $J^{(1)}$  are filtered and image  $J^{(2)}$  is obtained. We continue alternating between filtering image rows and image columns for a number of iterations. Finally, after spatial filtering, we apply the same filtering approach



Fig. 7 Different restoration results that can be seen in Video 2 (MOV, 190 MB) [URL: <http://dx.doi.org/10.1117/1.JEI.XX.XX.XXXXXX.2>]. (Please note that the large file size may result in a long download time.)

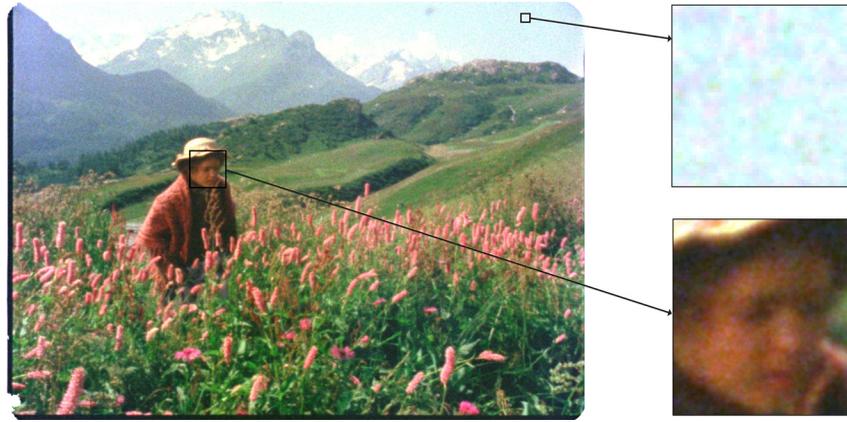


Fig. 8 Original film scan showing film grain in the sky and face region.

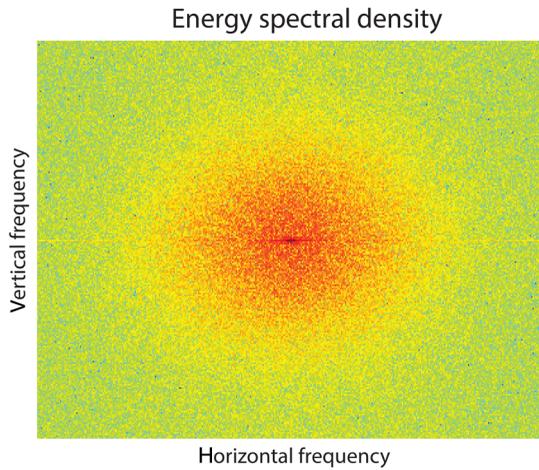


Fig. 9 Energy spectral density of a patch of the sky region of Fig. 8.

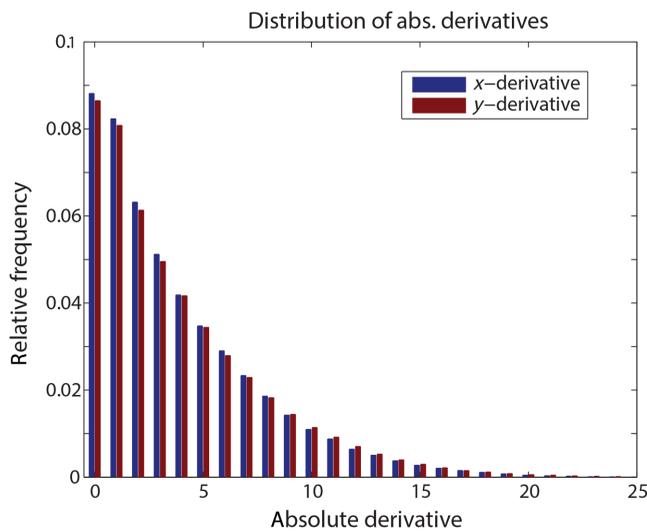


Fig. 10 Distribution of derivative magnitudes.

once in the time direction along optical flow trajectories. We experimentally found that the results converged after 10 spatial iterations. We used this setting as our default value, which, if desired, can be changed through our user interface.

We denote in the following a sequence of values of a one color channel along an image row, image column, or optical

flow trajectory as  $I_1, \dots, I_N$ . We filter this sequence according to

$$J_i^{(k+1)} = \frac{I_i + \sum_{j=1, j \neq i}^N \pi_{ij} J_j^{(k)}}{\sum_{j=1}^N \pi_{ij}}, \quad J_i^{(0)} := I_i, \quad (3)$$

where  $I_i$  is the original noisy image,  $J_i^{(k)}$  is the filtering result after  $k$  iterations, and  $\pi_{ij} \in [0, 1]$  is a permeability weight indicating the level of diffusion or permeability between pixels  $i$  and  $j$  ( $\pi_{ij} \approx 1$  is interpreted as high permeability). The permeability between two nonneighboring pixels  $i$  and  $j$  is computed as the product of permeabilities between neighboring pixels that lie on the interval between these pixels, i.e.,

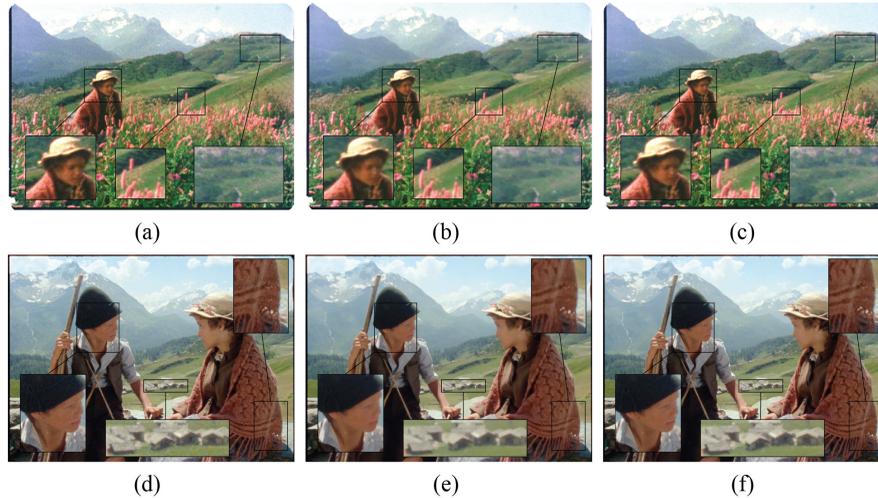
$$\pi_{ij} := \begin{cases} 1 & : i = j \\ \prod_{r=i}^{j-1} \pi_{r,r+1} & : i < j \\ \prod_{r=j}^{i-1} \pi_{r,r+1} & : i > j \end{cases} \quad (4)$$

Note that a small permeability between neighboring pixels along the way from pixel  $i$  to pixel  $j$  will lead to a low overall permeability  $\pi_{ij}$  and a low diffusion of pixel  $J_j^{(k)}$  into  $J_i^{(k+1)}$  [see Eqs. (3) and (4)]. This property is exploited for reducing or even stopping diffusion between certain pixels.

In our adaptive extension of the PGF filter (APGF), we compute the permeability between neighboring pixels  $\pi_{r,r+1}$  based on a variant of the Lorentzian stopping function which takes into account the grain noise characteristics. We essentially allow a stronger diffusion between neighboring pixels which create small magnitude differences and are located in an image area with isotropic structure. Diffusion is gradually reduced in areas which deviate from this assumption. Hence, the permeability  $\pi_{r,r+1}$  between neighboring pixels that have the desired properties is defined by

$$\pi_{r,r+1} := \left( 1 + \left| \frac{\bar{I}_r - \bar{I}_{r+1}}{\rho_r \gamma} \right|^\alpha \right)^{-1}, \quad (5)$$

where  $\bar{I}_r$  indicates the gray value of  $I$ ,  $\alpha$  is a shape parameter, and  $\gamma$  is a scaling parameter. Note that partial derivatives of magnitudes greater than  $\gamma$  are mapped to permeability weights  $\pi_{r,r+1}$  that are smaller than 0.5 (for  $\rho_r = 1$ ). In addition, the content-adaptive weight factor  $\rho_r \in [0, 1]$  reduces the permeability if the local image neighborhood



**Fig. 11** Degraining results. (a) and (d) two input frames from “Heidi,” (b) and (e) the results of applying PGF,<sup>28</sup> and (c) and (f) the results of applying APGF. APGF provides a better fidelity in comparison to PGF in anisotropic low gradient regions (edges, ridges, texture), while it shows similar fidelity in remaining regions.

is anisotropic. This is achieved by penalizing the deviation between the standard deviations in  $x$ - and  $y$ -directions, which are computed locally around pixel  $r$ , according to

$$\rho_r := e^{-\lambda|\sigma_x(r)-\sigma_y(r)|}, \quad (6)$$

where the parameter  $\lambda$  is used to adjust the sensitivity of the weight factor  $\rho_r$  to the level of deviation between the standard deviations  $\sigma_x$  and  $\sigma_y$ . The parameter  $\rho_r$  is computed according to Eq. (6) only if image rows or columns are filtered. In the case of filtering in the time direction, we use  $\rho_r = 1$ .

Hence, the permeability weights are computed in such a way that a strong spatial diffusion is only achieved if partial derivatives have a magnitude smaller than  $\gamma$  and the local image neighborhood is isotropic. If these conditions are not met then the permeability  $\pi_{r,r+1}$  is gradually reduced because  $\bar{I}_r - \bar{I}_{r+1}$  becomes large and  $\rho_r$  becomes small.

#### 4.2.3 Results

Our APGF method has three parameters:  $\alpha$ ,  $\gamma$ , and  $\lambda$ . In all filtering results shown in this paper, we use  $\alpha = 0.5$  and  $\lambda = 0.5$ . We select the scaling parameter  $\gamma$  by inspecting the distribution of partial derivatives in an image patch that contains a uniformly colored background (like the sky in Fig. 8), and visually test the impact of different  $\gamma$  parameters on the visual quality by selecting them from the range of the largest partial derivatives.

For temporal filtering, APGF requires optical flow as input, which we computed using the algorithm by Zimmer et al.<sup>33</sup> This method is quite robust against artifacts such as scratches and blotches. It further provides a smoothness parameter that can be increased in case such artifacts are evident.

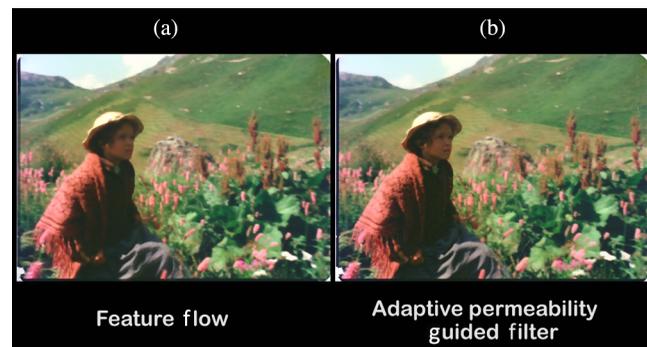
In Fig. 11, we show grain reduction results achieved with PGF and APGF, where both methods use the same  $\alpha$  and  $\gamma$  parameters. Note that PGF does not use a content-adaptive weight factor  $\rho_r$ , which is equivalent to using APGF with  $\lambda = 0$ . We observe that APGF provides a better image fidelity in anisotropic low gradient regions, such as the mountains

and flowers of Fig. 11(a) and the texture of the scarf and the settlement in the background of Fig. 11(d).

Figure 12 (Video 3) shows space-only and space-time filtering using APGF. It clearly shows the advantages of the spatiotemporal approach regarding temporal stability and smoothness, as well as overall quality. Figure 13 (Video 4) shows APGF with the previous state-of-the-art



**Fig. 12** Comparison of space-only and space-time filtering with APGF. Video 3 (MP4, 283 MB) [URL: <http://dx.doi.org/10.1117/1.JEI.XX.XX.XXXXXX.3>]. (Please note that the large file size may result in a long download time.)



**Fig. 13** Comparison between the (a) method of Lang et al.<sup>27</sup> (FeatureFlow) and (b) APGF. Video 4 (MP4, 30 MB) [URL: <http://dx.doi.org/10.1117/1.JEI.XX.XX.XXXXXX.4>]. (Please note that the large file size may result in a long download time.)

spatiotemporal approach presented by Lang et al.<sup>27</sup> Both methods perform similarly in terms of grain reduction and temporal stability, whereas APGF preserves fine image details, structures, edges, textures, etc., much better.

### 4.3 Contrast Enhancement

Contrast enhancement is another basic image/video processing operation and as such has also been widely explored. However, as there is typically no reference or original image information, tuning of these algorithms becomes an arbitrary operation. In film restoration, however, given our framework, a reference might be available. This is in cases where both a positive and a negative are available, and the conservator could decide to combine information from both sources. The positive may have lost contrast due to degradation, while still comprising the intended color grading information. The negative might be better preserved in terms of contrast, while lacking color grading.

In such a case, contrast transfer could be applied from the negative to the positive as shown in Fig. 14 and defined by the following equation:

$$L'_p = L_p + \beta D_{L_n}, \tag{7}$$

where  $L'_p$  and  $L_p$  are the output and input luminance channels of the positive,  $D_{L_n}$  is the detail luminance layer extracted from the negative, and  $\beta$  is a weighting parameter used to control the amount of contrast transfer. Figure 15 shows results of contrast transfer for varying  $\beta$ .

Two nontrivial tasks have to be solved to perform such contrast transfer, i.e., registration and detail extraction. Registration is challenging as negative and positive are by nature very different images (see Fig. 14). Furthermore, in a restoration scenario, both have been affected by different degradations that even after initial restoration lead to different image characteristics. In contrast, one of these differences is in fact the one that we want to explore here. Finally, both sequences were scanned from different film reels. Effects like shrinking may have affected both reels in different ways. Therefore, scanned historical film material is often temporally very unstable, i.e., jerky, warped, etc. Some of our examples exhibit such degradations. Stable alignment of two such different and very differently imperfect sequences is thus a challenging task. We experimented with various



Fig. 15 Results of contrast transfer for increasing  $\beta$ .

alignment approaches. SURF features<sup>34</sup> combined with a robust estimator<sup>35</sup> assuming an affine model turned out to provide very good results in most of our experiments. The implementations of SURF and the registration estimator used in this work are the ones provided by MATLAB with the functions detectSURFFeatures and estimateGeometricTransform. For the results, we mostly used the default parameters. For some images (1 out of 50 on average), interactive correction was necessary, as for contrast transfer almost perfect alignment is essential (see Fig. 14).

For extraction of the detail layer, we build on the very recent work by Aydin et al.<sup>28</sup> on temporally consistent video tone mapping. As part of their framework, they describe an algorithm for temporally coherent separation of a video sequence into a low-contrast base layer and a detail layer that captures all high-contrast information. The core of that approach is the same PGF algorithm that we also use for denoising (see Sec. 4.2).

The detail layer  $D_I$  of an image  $I$  is defined by the equation

$$D_I = I - B_I, \tag{8}$$

where the base layer  $B_I$  is computed as a low-pass version of the input image. It is worth noting that Eq. (8) differs from the common formulation where the detail layer is obtained by a subtraction in the log domain. Such a subtraction of a low-pass image may lead to halos and ringing artifacts as illustrated for a 1-D signal in Fig. 16, if a nonedge-aware filter is applied. Careful design of an edge-aware low-pass filter can avoid this as shown in Aydin et al.<sup>28</sup> Figure 17 shows the separation of an image into a base layer and a detail layer defined by Eq. (8).

Figure 18 shows the results of contrast enhancements using different low-pass filters for base layer generation. The nonedge-aware Gaussian filter produces halos on strong edges. The edge-aware WLS filter<sup>36</sup> and our PGF provide much better results. WLS is a state-of-the-art edge-aware filter, known to provide excellent results. The advantages of our approach over WLS are that WLS is computationally

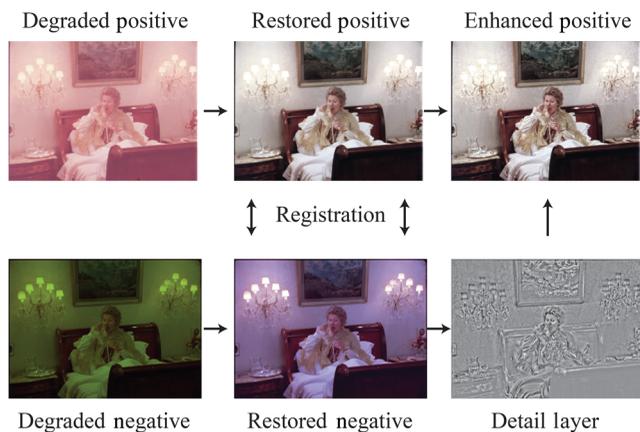
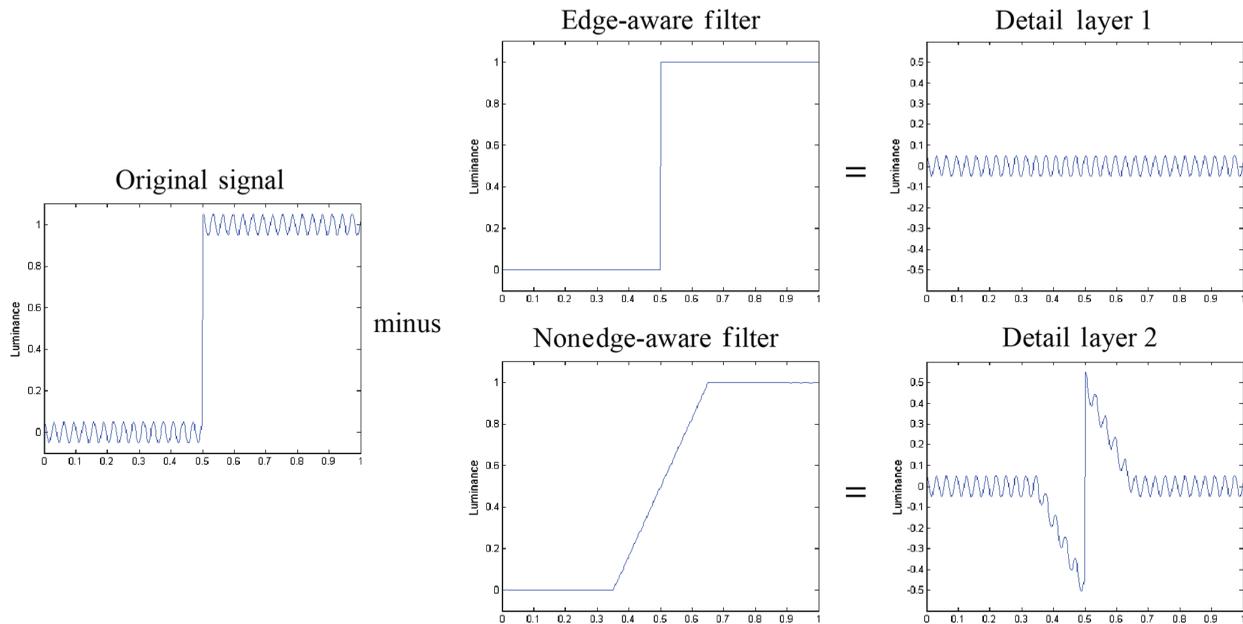
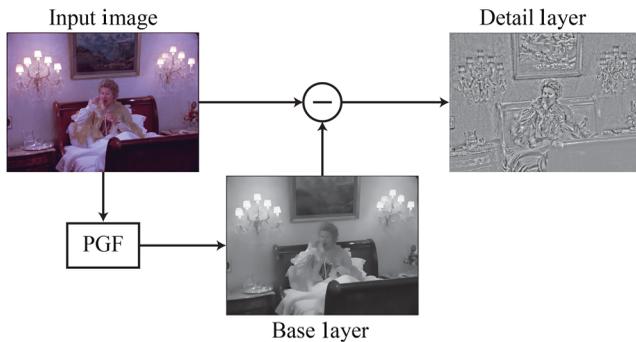


Fig. 14 Contrast transfer pipeline.



**Fig. 16** Detail layer computation using edge-aware versus nonedge-aware filtering. Halos and ringing artifacts may occur, if nonedge-aware filtering is applied for base layer computation.



**Fig. 17** Base and detail layer decomposition.

very expensive and not temporally consistent. As shown in Aydın et al.,<sup>28</sup> PGF closely approximates WLS, while providing temporal consistency and requiring significantly less computation. All results reported in this paper and video were generated by automatic detail layer extraction using this algorithm.

## 5 Results

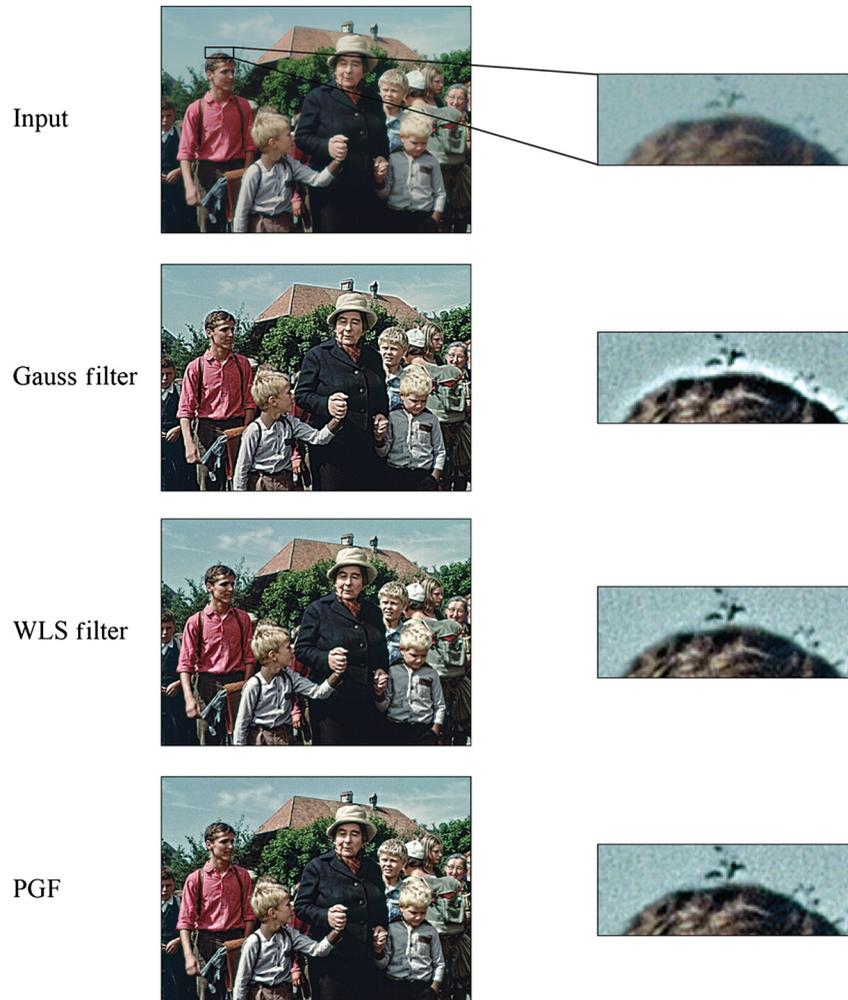
Our algorithms were developed and tested in the context of a real film restoration project called DIASTOR in cooperation with a film archive, a national broadcaster, and a post-production house. We selected sequences from two historical films Heidi and Kummerbuben. For all test sequences, we had mechanically/physically restored and scanned versions of negative and positive available as files (observations  $N_D$  and  $P_D$ ).

Heidi is a production from 1978 on Kodak Eastman film stock in 16 mm. As such it was produced for TV broadcast only and not for theatrical release. It features a diverse variety of shots, including landscapes, dialogs, closeups, etc., with

very cinematic camera work. The positive exhibits severe degradations in various classical ways (color fading, noise, loss of contrast, temporal instability, scratches, dirt, etc.). The negative is generally in better condition.

Kummerbuben is a production from 1968 also on Kodak Eastman, but in 35 mm for theatrical release. The available material is in better condition compared to Heidi, while still exhibiting all classical degradations. We selected a shot from Kummerbuben to illustrate the variety of different input material one is confronted with in film restoration, which contains significant camera and object motion.

Figure 19 shows selected results of our restoration pipeline for the positive ( $f_P^{-1*}$ ). The corresponding video results are available in Fig. 7 (Video 2). Severe color degradations were corrected using our efficient approach based on keyframes and color transformation estimation. Next, APGF was used for denoising, preserving image textures/details/structures and film look, while reducing extensive noise and grain. Finally, contrast was enhanced transferring details from the available negatives. We deliberately did not apply tools for scratch or dirt removal or other classical restoration tools, in order to show only the effect of our advanced tools and how the tools are robust against these artifacts. We must recognize that our tools can be affected by artifacts such as intensity flicker, blotches, and scratches. The result of our methods depends on the number and size of the blotches and scratches and also on the amplitude of the intensity flicker. If these artifacts are preponderant, we recommend removing them before applying our methods. We also decided not to crop the results to better illustrate the difficulties for alignment with such material, see especially the sequence with the two ladies walking uphill (WalkingLadies). The final results provide a decent quality, sufficient for reuse, while retaining the vintage character of the footage. Finally, it is the task of the conservator to create a faithful and appealing output, using the best available tools



**Fig. 18** Contrast enhancement results obtained with different filters for base layer extraction.

in an appropriate way. As our tools apply and extend the latest state-of-the-art in visual computing, they certainly expand possibilities and reduce manual efforts.

In a second set of experiments, we illustrate the framework idea (Fig. 1). If the negative  $N_D$  is available and in better condition, it makes more sense to apply the tools on that material ( $f_N^{-1*}$ ). The positive  $P_D$  may only be restored for a couple of keyframes, mainly focusing on color. The algorithm from Sec. 4.1 can then be used to get estimates for color grading  $f_G^*$  per keyframe for associated sequences. These estimates can then be used to regrade a new positive sequence  $P_O^*$  from the restored negative sequence  $N_O^*$ . The advantage of such an approach is that only selected keyframes from  $P_D$  are necessary only as a color reference. Heavily degraded images can be avoided. Figure 20 shows examples of this approach to results of the positive pipeline. Corresponding videos are available in Fig. 7 (Video 2). The positive of *WalkingLadies* exhibits severe degradations toward the end of the sequence, as shown in the last two examples in Fig. 20 (green line close to the left border in third image, destruction across the whole fourth image). It would be very cumbersome to repair these degradations. The negative on the other hand is not affected by these destructions and can, therefore, be much more easily restored. Furthermore, the

positive is temporally very unstable for this example, whereas the negative does not need much stabilization. These examples illustrate that film restoration is unique for every project. Especially for very old historical material, it is an interactive work similar to other types of art restoration. Awareness of the framework paired with powerful related tools will lead to the best possible results.

Regarding the computational times, the color restoration tool can potentially run in nearly real time, since each pixel is processed independently. The degrading filter does not run in real time because it temporally filters the video based on the optical flow. For the computation of the optical flow, we use the method by Zimmer et al.,<sup>33</sup> which computes high-quality flows and requires about 10 s for a full HD frame. Faster optical flow methods can be used at the cost of quality. Once the optical flow is computed, the filter requires about 3.5 s to filter a full HD frame using a spatial neighborhood of 21 pixels and a temporal neighborhood of 11 frames. Also, the contrast enhancement tool is not real time since it uses the PGF filter for the computation of the base layer that is similar to the degrading filter and it requires a comparable computational time. Once the base layer is computed, the contrast transfer tool requires three additional seconds for the registration, the computation of the detail layer, and the addition of it to the positive frame.



Fig. 19 Results after different steps of the pipeline for positive film restoration.



Fig. 20 Comparison of the results of the positive and the negative restoration pipeline.

## 6 Conclusions

In this paper, we addressed the problem of digital film restoration which we model as an estimation process given noisy and incomplete observations. We presented three different restoration tools that tackle specific imperfections of historical film material. First, we adapted a color transform approach for our specific needs and extended it by an automatic correspondence selection method, which in most cases significantly reduces user interaction. Next, we extended a very recent and very powerful noise reduction algorithm for our specific application by adding an adaptive component (APGF) tailored to grain noise in historic film. Finally, a contrast transfer method was introduced that extends a very recent spatiotemporal filtering algorithm to enhance the contrast of a positive using information from a corresponding negative. All three algorithms extend recent advances in visual computing and adapt them for specific tasks in digital film restoration. As such, they expand the toolbox from which a conservator can choose to handle a given restoration project at hand. Our results illustrate the power and quality of the tools on realistic examples from film restoration projects. We also illustrated the power of framework aspects in examples combining information from both positive and negative to create optimum final results.

In the end, film restoration will remain a very interactive art, where experience and intuition of the conservator play a decisive role, especially for very old and heavily degraded material. Improvements of tools may continue as long as progress in visual computing continues. The digital film restoration framework remains as a reference for involved processes and components.

### Acknowledgments

This work was supported by the Swiss Commission for Technology and Innovation (CTI/KTI) under Grant No. 15019.1 PFES-ES as part of the project DIASTOR.

### References

- P. Read and M.-P. Meyer, *Restoration of Motion Picture Film*, Butterworth-Heinemann, Oxford, United Kingdom (2000).
- A. Kokaram, "Ten years of digital visual restoration systems," in *IEEE Int. Conf. on Image Processing (ICIP '07)*, Vol. 4 (2007).
- L. Rosenthaler and R. Gschwind, "Restoration of movie films by digital image processing," in *IEE Seminar Digital Restoration of Film and Video Archives*, p. 6 (2002).
- P. Schallauer, A. Pinz, and W. Haas, "Automatic restoration algorithms for 35 mm film," *Videre: J. Comput. Vision Res.* **1**(3), 60–85 (1999).
- M. Werlberger et al., "Optical flow guided TV-L1 video interpolation and restoration," in *Energy Minimization Methods in Computer Vision and Pattern Recognition*, pp. 273–286 (2011).
- R. Gschwind and F. Frey, "Electronic imaging, a tool for the reconstruction of faded color photographs," *J. Imaging Sci. Technol.* **38**(6), 520–525 (1994).
- F. Frey and R. Gschwind, "Mathematical bleaching models for photographic three-color materials," *J. Imaging Sci. Technol.* **38**(6), 513–519 (1994).
- T. Oskam et al., "Fast and stable color balancing for images and augmented reality," in *Second Int. Conf. on 3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT '12)*, pp. 49–56 (2012).
- I. Amidor, "Scattered data interpolation methods for electronic imaging systems: a survey," *J. Electron. Imaging* **11**(2), 157–176 (2002).
- Y. Hwang et al., "Color transfer using probabilistic moving least squares," in *IEEE Conf. on Computer Vision and Pattern Recognition*, pp. 3342–3349 (2014).
- H. S. Faridul et al., "Approximate cross channel color mapping from sparse color correspondences," in *IEEE Int. Conf. on Computer Vision Workshops*, pp. 860–867 (2013).
- E. Reinhard et al., "Color transfer between images," *IEEE Comput. Graphics Appl.* **21**(5), 34–41 (2001).
- X. Xiao and L. Ma, "Color transfer in correlated color space," in *ACM Int. Conf. on Virtual Reality Continuum and its Applications*, pp. 305–309 (2006).
- X. Xiao and L. Ma, "Gradient-preserving color transfer," *Comput. Graphics Forum* **28**(7), 1879–1886 (2009).
- F. Pitie, A. Kokaram, and R. Dahyot, "N-dimensional probability density function transfer and its application to color transfer," in *Tenth IEEE Int. Conf. on Computer Vision (ICCV '05)*, Vol. 2, pp. 1434–1439 (2005).
- L. Neumann and A. Neumann, "Color style transfer techniques using hue, lightness and saturation histogram matching," in *Computational Aesthetics in Graphics, Visualization and Imaging*, pp. 111–122, Eurographics Association, Aire-la-Ville, Switzerland (2005).
- P. Milanfar, "A tour of modern image filtering," *IEEE Signal Process. Mag.* **30**(1), 106–128 (2013).
- A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *Proc. Int. Conf. Computer Vision and Pattern Recognition*, Vol. 2, pp. 60–65 (2005).
- A. Buades, B. Coll, and J.-M. Morel, "Nonlocal image and movie denoising," *Int. J. Comput. Vision* **76**, 123–139 (2008).
- A. Buades, B. Coll, and J.-M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Model. Simul.* **4**, 490–530 (2005).
- E. Gastal and M. Oliveira, "Adaptive manifolds for real-time high-dimensional filtering," in *Proc. of SIGGRAPH*, Vol. 31, p. 33 (2012).
- E. Gastal and M. Oliveira, "Domain transform for edge-aware image and video processing," in *Proc. of SIGGRAPH*, Vol. 30, p. 69 (2011).
- A. Adams, J. Baek, and A. Davis, "Fast high-dimensional filtering using the permutohedral lattice," *Comput. Graphics Forum* **29**, 753–762 (2010).
- Z. Farbman et al., "Edge-preserving decompositions for multi-scale tone and detail manipulation," *ACM Trans. Graph.* **27**(3), 67 (2008).
- C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Proc. 6th Int. Conf. Computer Vision*, pp. 839–846 (1998).
- K. Dabov et al., "Image denoising by sparse 3D transform-domain collaborative filtering," *IEEE Trans. Image Process.* **16**, 2080–2095 (2007).
- M. Lang et al., "Practical temporal consistency for image-based graphics applications," *ACM Trans. Graph.* **31**, 34 (2012).
- T. Aydin et al., "Temporally coherent local tone mapping of HDR video," *ACM Trans. Graph.* **33**, 196 (2014).
- Y. Shih et al., "Style transfer for headshot portraits," *ACM Trans. Graph.* **33**, 148 (2014).
- B. Wang, Y. Yu, and Y.-Q. Xu, "Example-based image color and tone style enhancement," *ACM Trans. Graph.* **30**(4), 64 (2011).
- A. O. Akyüz et al., "Style-based tone mapping for HDR images," in *SIGGRAPH Asia 2013 Technical Briefs*, SA '13, p. 23, ACM, New York, NY (2013).
- B. Flueckiger, "Timeline of historical film colors," 2012, <http://zauberklang.ch/filmcolors/> (1 May 2016).
- H. Zimmer, A. Bruhn, and J. Weickert, "Optical flow in harmony," *Int. J. Comput. Vision* **93**(3), 368–388 (2011).
- H. Bay et al., "Speeded up robust features (SURF)," *Comput. Vision Image Understanding* **110**(3), 346–359 (2008).
- P. Torr and A. Zisserman, "Robust computation and parameterization of multiple view relations," in *Sixth Int. Conf. on Computer Vision, 1998*, pp. 727–732 (1998).
- Z. Farbman et al., "Edge-preserving decompositions for multi-scale tone and detail manipulation," *ACM Trans. Graph.* **27**, 67 (2008).

**Simone Croci** is a research engineer with a master's degree in computer science from the Swiss Federal Institute of Technology Zurich (ETHZ). During his career, he has gathered experience principally in the field of visual computing working at ETHZ and at Disney Research Zurich. He contributed to papers in the areas of wave-based rendering, virtual view synthesis, HDR video, and camera calibration.

**Tunç Ozan Aydın** received his BS degree from the Civil Engineering Department of Istanbul Technical University in 2003, his MS degree from the College of Computing of Georgia Institute of Technology in 2005, and his PhD (summa cum laude) from the Computer Graphics Department of Max-Planck-Institut für Informatik in 2010. His dissertation received the Eurographics PhD Award in 2012. He is a research scientist at Disney Research Zurich. His research focuses on video postprocessing technologies.

**Nikolce Stefanoski** holds his diploma degree in mathematics and his Dr.-Ing. degree in electrical engineering and computer science from Leibniz University of Hannover, Germany. He is a head of video technology and development at uniqFEED. Before moving to uniqFEED, he worked as an associate research scientist at Disney Research

Zurich, where he was mainly involved in research projects related to novel production, postproduction, distribution, and display technologies for the television of the future.

**Markus Gross** received his Master of Science degree in electrical and computer engineering and his PhD in computer graphics and image analysis from Saarland University, Germany, in 1986 and 1989, respectively. He is a professor of computer science at Swiss Federal Institute of Technology Zurich (ETH), a head of the Computer Graphics Laboratory, and the director of Disney Research Zurich. He joined the ETH Computer Science Faculty in 1994. Before joining

Disney, he was a director of the Institute of Computational Sciences at ETH.

**Aljosa Smolic** is a senior research scientist and leader of the Advanced Video Technology Group at Disney Research Zurich. He has been responsible for numerous research projects in various fields of visual computing and video processing, and has published more than 140 scientific papers. He serves as an associate editor of the *IEEE Transactions on Image Processing*, and on various technical program committees of leading conferences. He is an associate lecturer at ETH Zurich, teaching multimedia communications.