



HAL
open science

Motion Driven Tonal Stabilization

Oriel Frigo, Neus Sabater, Julie Delon, Pierre Hellier

► **To cite this version:**

Oriel Frigo, Neus Sabater, Julie Delon, Pierre Hellier. Motion Driven Tonal Stabilization. 2015.
hal-01150558v1

HAL Id: hal-01150558

<https://hal.science/hal-01150558v1>

Preprint submitted on 22 May 2015 (v1), last revised 24 Jun 2016 (v5)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Motion Driven Tonal Stabilization

ORIEL FRIGO^{*,†} and NEUS SABATER^{*} and JULIE DELON[†] and PIERRE HELLIER^{*}

^{*}Technicolor R&I, Cesson-Sévigné, France

[†]Université Paris Descartes, MAP5, Paris, France

This paper addresses the problem of tonal fluctuation in videos. Due to the automatic settings of consumer cameras, the colors of objects in image sequences might change over time. We propose here a fast and computationally light method to stabilize this tonal appearance, while remaining robust to motion and occlusions. To do so, a minimally-viable color correction model is used, in conjunction with an effective estimation of dominant motion. The final solution is a temporally-weighted correction, explicitly driven by the motion magnitude, both visually efficient and very fast, with potential to real time processing. Experimental results obtained on a variety of sequences outperform the current state of the art in terms of tonal stability, at a much reduced computational complexity.

Categories and Subject Descriptors: I.4.8 [Image Processing and Computer Vision]: Scene Analysis—*Color*; I.4.3 [Image Processing and Computer Vision]: Enhancement—*Grayscale manipulation*

Additional Key Words and Phrases: Tonal stabilization; White balance; Exposure control; Video editing;

1. INTRODUCTION

The increasing number of amateur video footages facilitated by the proliferation of low-cost video cameras has made it visible a number of video artifacts, some of them being motion and tonal instabilities. While motion stabilization has been studied by several researchers, tonal instability has been far less discussed. Video tonal instability is a particular temporal artifact characterized by fluctuations in the colors of adjacent frames of a sequence. According to [Farbman and Lischinski 2011], in modern videos these instabilities are mainly caused by automatic settings of the camera, notably automatic white balance and automatic exposure. These common features of consumer digital cameras are intended respectively to provide color balanced and well exposed images, while facilitating the user experience. However, these features are mostly appropriated for still images and are not stable in time, resulting in unpleasant tonal instabilities that can be perceived in videos. A notable problem with automatic white balance algorithms is their dependency on illuminant estimation, which is considered an ill-posed problem. Classical assumptions of color constancy algorithms are easily violated in practice and in a context of temporal scene changes, it is likely to result in chromatic instability.

While automatic white balance can be simply turned off in some cases, low end cameras offer no control over setup parameters. In this case, the only alternative to avoid unpleasant tonal fluctuations is to further process the video. This preprocessing can also be crucial for computer vision applications relying on brightness or tonal constancy assumptions.

Generally speaking, tonal stabilization can be described as searching for the transformations that minimize undesired tonal variations in multiple images of a sequence. While surprisingly few works have specifically attempted to correct such color fluctuations in videos [Farbman and Lischinski 2011], [Wang et al. 2014], nu-

merous works are motivated by similar purposes in different research communities.

Let us start with a word on radiometric calibration approaches, which takes into account operations performed in the camera pipeline in order to retrieve a physically based color calibration model between images. Recent models [Chakrabarti et al. 2009], [Lin et al. 2011], [Kim et al. 2012] admit that the output image u of a camera is related to the irradiance vector e measured by the sensor by a combination of linear (white balance, color space conversion) and nonlinear (gamut mapping, camera response function) transformations. If the nonlinear part of the mapping could be estimated for a video sequence, it could be inverted and the resulted sequence could be corrected with linear transformations. Unfortunately, recovering the camera response function necessitates registered images under multiple exposures, as well training sets of RAW-sRGB image pairs [Grossberg and Nayar 2002], [Lin et al. 2011], both being generally not available for smartphone and point-and-shoot video cameras.

In a related direction, one could argue that color constancy algorithms would be the ideal solution to solve the tonal stabilization problem. Starting with Land and McCain work on the Retinex theory of color vision [Land et al. 1971][Land 1977][Brainard and Wandell 1986], color constancy has been extensively studied and remains an active research topic [Hordley 2006] [Gijssenij et al. 2011]. In digital cameras, the feature intended to approximate color constancy is known as automatic white balance (AWB). The most common approach to perform AWB is to first estimate the color temperature of the illuminant in the scene, and then correct the image by compensating the ratio of the estimated illuminant to the canonical (neutral) illuminant with some variant of a Von Kries diagonal transformation. However, illuminant estimation is a severely under-constrained problem, necessitating additional assumptions intended to cope with its ill-posed nature. Gijssenij *et al* [Gijssenij et al. 2011] in their extensive survey on computational color constancy, have categorized illuminant estimation approaches into *static* methods, based on fixed parameter settings and low-level statistics, in opposition to *gamut-based* and *learning-based* methods, which tune their parameters to a specific set of images. Automatic white balance in low cost video cameras is likely to be based on static methods, making strong assumptions about the scene content, so that illuminant color can be estimated in real time. Since this automatic feature is the cause of tonal instabilities, recovering the scene illuminant in each frame to correct the sequence is also prone to fail in practice. This is true even if the available computing power makes it possible to use more involved white balance methods than those used directly in the camera.

Another possibility is to draw on color transfer, which aims at modifying the colors of an input image according to the colors of an example image. Unlike color constancy, color transfer methods make very few assumptions about the physics of the scene or the camera (the scene illuminant or the camera response function are not explicitly estimated), being rather a general and pragmatic approach to estimate color transformations between images. Reinhard

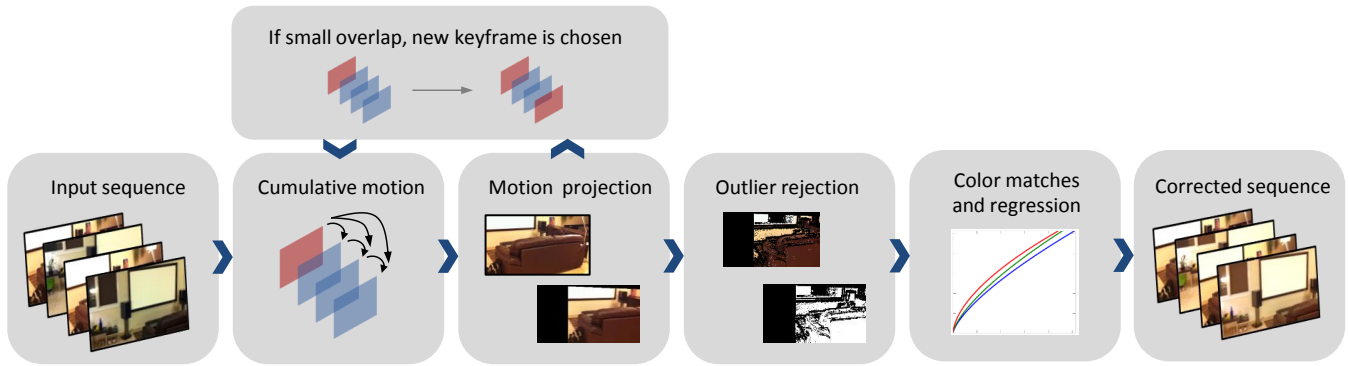


Fig. 1. Overview of proposed method for tonal stabilization of videos.

et al. [Reinhard et al. 2001] popularized the concept of color transfer 15 years ago with a very simple affine matching between 3D color distributions. Other works have since proposed more complex global color transfer in terms of nonlinear histogram matching [Neumann and Neumann 2005], [Pitié et al. 2007], [Papadakis et al. 2011], or optimal transportation between compact color signatures [Freedman and Kisilev 2010], [Ferradans et al. 2013], [Murray et al. 2012], [Wu et al. 2013], [Frigo et al. 2014]. Another class of methods assumes that there are spatial correspondences to be found between the input and example images, these correspondences being used to derive a color transformation [HaCohen et al. 2011]. If this assumption reduces the scope of the method for general images, it is particularly relevant for specific user cases such as optimizing color consistency photo collections from the same scene [HaCohen et al. 2013]. In the case of videos, we also expect to find numerous spatial correspondences between neighbor frames of the sequence. As we will see, the tonal stabilization algorithm presented in this paper draws on a first raw motion estimation between frames to compute a global color correction.

While very few works have attempted to correct color fluctuations in videos, several global [Decencière 1997], [Delon 2006] or local [Pitié et al. 2004], [Delon and Desolneux 2010] approaches have been proposed to remove high frequency brightness fluctuations, also known as flicker. In practice, the causes and effects of flicker in videos are quite different from those of tonal instabilities. In old archived films, these high frequency brightness fluctuations are mainly due to physical reasons, which may cause the flicker to vary along the image spatial coordinates. On the contrary, tonal instabilities observed in modern videos (fluctuations in camera exposure or white balance) tend to be global and are low frequency artifacts. As a consequence, parametric and non-parametric deflickering methods are not well suited to correct tonal instabilities. Moreover, the extension of these deflickering methods to color images is far from obvious.

To the best of our knowledge, only two previous works [Farbman and Lischinski 2011], [Wang et al. 2014] have proposed approaches specifically designed to correct color fluctuations in videos. The first one was proposed by Farbman *et al* in [Farbman and Lischinski 2011]. The method works by aligning tonally each frame of the video with one or several pre-selected reference (anchor) frames (for instance the first frame of the movie). Adjustment maps between successive frames are computed without explicit motion compensation, the authors claiming that many pixels in the same grid position in two successive frames are likely to correspond to the same surface in the scene. In practice, the method provides a

reasonable solution to stabilize tonal fluctuations in static videos, but at the cost of high space and time complexity. Besides, the lack of real motion estimation makes the computation of the adjustment map highly sensitive to noise or fast movements, and lack of accurate correspondences can result in flickering and error propagation.

The second one, due to Wang *et al.* [Wang et al. 2014], starts by estimating motion globally between successive frames, by relying on local features correspondences. A nine parameters affine color transformation is then used to model the exposure and white balance changes between two frames in the log domain. These affine color transformations are estimated by least squares for all neighboring frame pairs and accumulated to obtain a "color state" for each frame of the video, represented as a 4×4 matrix. To avoid data overfitting, a regularization term is added to force the affine transformations to be close to the identity matrix. In a second step, the frame color state function is smoothed by minimizing an energy which also tends to control the deviation from the original color state. The regularization step necessitates to use PCA on the set of color states in order to avoid smoothing the different parameters of the affine matrices independently. While the results provided by the authors are visually good and much more satisfying than those of [Farbman and Lischinski 2011], the method is surprisingly complex to implement and requires to tune several parameters.

We are aware of existing commercial tools (*Adobe After Effects*®, *Final Cut Pro*®) able to correct specific brightness fluctuations, such as high frequency flickering commonly seen in time-lapse photography. Nevertheless, these applications remain limited when it comes to correct tonal instabilities. The *Color Stabilizer Tool*¹ from *Adobe After Effects* is based on color sampling of selected points in a reference frame, adjusting the colors neighbor frames so that the color values of selected points remain constant throughout the duration of the layer. The *Flicker Free* plugin² for *Final Cut Pro* can be useful to remove high frequency flicker and equalize the exposure of footage, but is not targeted to stabilize white balance fluctuations.

The method presented in this paper draws on two elementary observations. First, we show that a simplified parametric color transfer model between the frames of a video is enough to correct convincingly color instabilities. Second, we observe that if motion estimation is necessary for tonal stabilization, dominant motion estimation is generally more than enough to infer large sets of correspondences between frames. These are the two key ideas permitting to

¹<https://helpx.adobe.com/after-effects/using/color-correction-effects.html>

²<https://digitalanarchy.com/Flicker/main.html>

develop a method which is both visually efficient and very fast, with potential to real time processing. Furthermore, we propose a temporal weighting scheme, where the intensity of tonal stabilization is directly guided by the motion speed.

The outline of the paper is the following. Our method is presented in Section 2, first in the simplified case of registered images, and then generalized to sequences with motion. Section 3 illustrates the efficiency of the proposed method on different sequences and shows that it compares favorably to [Farbman and Lischinski 2011], [Wang et al. 2014], and achieve these results at a much reduced computational cost. All results are available on the project website³.

2. PROPOSED METHOD

In this section, we present the main contributions of our method for video tonal stabilization. Our aim is to conceive a method that is:

- (1) Accurate enough to correct color instabilities observed between frames in a video;
- (2) Robust against motion, occlusion and noise;
- (3) Computationally simple enough to be implemented in a near real time application.

Figure 1 presents a general overview of the proposed method. In order to achieve robustness against motion and occlusion (an important limitation of [Farbman and Lischinski 2011]) while satisfying the simplicity requirement, we make use of dominant motion estimation (and compensation) to estimate the color correspondences between frames. By means of cumulative motion, we are able to quickly register images even if they differ by several frames in time. Color correspondences are then used to estimate a color transformation that is applied to correct the tonal instabilities. The requirement of physical accuracy for the tonal transformation model being in contradiction with the other requirements of robustness and simplicity, our color correction model tries to achieve a good tradeoff between these three properties.

The proposed algorithm makes two hypotheses on the sequences to be corrected. First, it is assumed that there are always spatial correspondences (or redundancy in content) between neighboring frames in the sequence (no scene cuts). This is generally true for every sequence composed of a single shot, as long as it does not pass through extreme variations of scene geometry (nearly total occlusion) or radiometry (huge changes in illumination or saturation). Second, we assume that there is a global transformation which can compensate the colorimetric aberrations between the frames. This implies that the observed color instability and consequently the camera response function are global (spatially invariant). In other words, the proposed method is not suitable for correction of local tonal instabilities. This assumptions also holds for every method of the state of the art [Farbman and Lischinski 2011], [Wang et al. 2014].

In the following subsections, we discuss in detail each step of the proposed method. For the sake of simplicity, we start the discussion with the tonal transformation model, first assuming the simplest case of color correction between registered images. We then present our model to deal with the general case of tonal stabilization of sequences containing motion.

2.1 Tonal Transformation Model

Let $\{u_t\}_{t=1,\dots,D}$ be a registered sequence of color frames $u_t : \Omega \rightarrow \mathbb{R}^3$ defined over the same spatial domain Ω and let u_k be a reference keyframe in the sequence. The color channels of u_t are written (u_t^R, u_t^G, u_t^B) . In order to tonally stabilize the sequence, for every following frame $u_t, t > k$ we look for a parametric color transformation $T_t : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ such that $T_t(u_t) \simeq u_k$. Now, as recalled in the introduction and according to [Kim et al. 2012], u_t is related to the irradiance vector $e = (e^R, e^G, e^B)$ measured by the sensor (RAW intensities) by the relation

$$u_t = f \circ h(\mathbf{T}_s \mathbf{T}_w e) \quad (1)$$

where \mathbf{T}_s is a 3×3 matrix transformation that accounts for color space conversion, $h : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is a nonlinear gamut mapping operator, $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is the nonlinear tone mapping of the camera and \mathbf{T}_w is a diagonal 3×3 matrix accounting for white balance and exposure (varying over time)

$$\mathbf{T}_w = \begin{bmatrix} \phi & 0 & 0 \\ 0 & \xi & 0 \\ 0 & 0 & \psi \end{bmatrix}. \quad (2)$$

If u_t and u_k are two perfectly registered images of the same scene, taken by the same camera, differing only with respect to white balance and exposure (so that these images have identical irradiance e), then u_t and u_k are related by

$$H^{-1}(u_t) = \begin{bmatrix} \frac{\phi_t}{\phi_k} & 0 & 0 \\ 0 & \frac{\xi_t}{\xi_k} & 0 \\ 0 & 0 & \frac{\psi_t}{\psi_k} \end{bmatrix} H^{-1}(u_k), \quad (3)$$

where $H = f \circ h \circ \mathbf{T}_s$. In theory, we can achieve tonal stabilization between images u_t and u_k with a simple diagonal transformation performed in the camera sensor space (given by the nonlinear transformation H^{-1}). This tonal stabilization model, inspired by radiometric calibration [Kim et al. 2012] [Xiong et al. 2012], can be seen as an accurate procedure to perform camera color transfer when irradiances are known in the form of RAW images, allowing an estimate of H . However, for the problem of tonal stabilization, we are faced with videos taken with low-cost cameras, from which we cannot make the usual assumptions that are necessary to compute radiometric calibration. The assumption of multiple exposures from the same scene, which is required to estimate the camera response function may not be valid for some sequences, and RAW-sRGB correspondences are also not available in practice.

In this paper, we follow another direction and claim that a much simpler model can be used for achieving visually satisfying tonal stabilization when the transformation H is unknown. Several paths can be followed, from fully non parametric color transformations (such as weighted interpolation or optimal transport / histogram specification) to more or less complex parametric models (affine, spline, polynomial). Parametric models have the important advantage of being expressed by smooth and regular functions, well defined for the whole color range (extrapolation is not a problem even if all colors are not observed) and potentially described by few parameters, which reduces the risk of instabilities in time. A purely diagonal model applied to sRGB images is unfortunately not suitable to cope with non linearities inherent to camera response functions. In our quest for a tonal transformation model that is simple enough to be fastly computed, and accurate enough to produce a visually pleasant stabilized sequence, we converged to a separable power law transformation, which is shown to be a good compromise to fit the criteria described above. Perhaps not surprisingly, this model

³http://oriel.github.io/tonal_stabilization

is quite close to the ones used in the color post-production industry⁴.

Our deliberately simplified model for T_t assumes a separable transfer function on color channels:

$$T_t = (T_t^R, T_t^G, T_t^B), \text{ where } T_t^c(s) := \alpha_c s^{\gamma_c}, c \in \{R, G, B\}. \quad (4)$$

The accuracy (in term of goodness of fit) of this power law model will be demonstrated in the experimental section on several image pairs of the same scene, with varying white balance and exposure values. The ability of the model to achieve stabilization results more or equally satisfying than more complex models on image sequences will also be shown.

In order to estimate α_c, γ_c , we solve for every color channel of the frame u_t the linear least square fitting problem

$$\arg \min_{\alpha_c, \gamma_c} \sum_{\mathbf{x} \in \Omega} (u_k^c(\mathbf{x}) - \alpha_c u_t^c(\mathbf{x})^{\gamma_c})^2, \quad (5)$$

which can be rewritten as

$$\arg \min_{\alpha_c, \gamma_c} \sum_{\mathbf{x} \in \Omega} (\log u_k^c(\mathbf{x}) - \gamma_c \log u_t^c(\mathbf{x}) + \log \alpha_c)^2, \quad (6)$$

whose solution is given by

$$\gamma_c = \frac{\text{Cov}(\log u_t^c, \log u_k^c)}{\text{Var}(\log u_t^c)}, \alpha_c = \exp(\overline{\log u_k^c} - \gamma_c \overline{\log u_t^c}), \quad (7)$$

where $\bar{z} = \frac{1}{\#\Omega} \sum_{\mathbf{x} \in \Omega} z(\mathbf{x})$ is the average value of $z(\mathbf{x})$, $\mathbf{x} \in \Omega$. This is not equivalent to solving the least squares problem directly, since the loss function becomes logarithmic, which gives more weight to residuals computed from low values. However, this formulation permits to compute the coefficients α_c and γ_c exactly and very quickly (the computation is linear in the number of correspondent points). For higher regression accuracy in terms of linear mean squared error, the solution could be alternatively computed with a numerical method such as gradient descent [Levenberg 1944], [Marquardt 1963].

2.2 Image Registration

For estimating tonal transformations between different frames, it is desirable to have a dense set of correspondences, especially in homogeneous areas. In this paper, we choose to rely on global motion estimation. Estimating the dominant motion of the camera is computationally simpler (it can potentially be computed in real time) than dense optical flow, and in our experience, tonal instabilities seen in videos are usually highly correlated with the camera motion. In contrast to tasks that depend heavily on precise motion, we do not need a highly accurate motion description in order to estimate a color transformation that compensates tonal differences between frames. It is much more important that this dominant motion estimation is robust enough to be cumulated in time and permit color correction between non-neighbor frames.

In practice, for all pairs of consecutive frames u_l and u_{l-1} in the sequence, we make use of the robust algorithm of [Odobez and Boutheymy 1995] to estimate the affine motion transformation $A_{l,l-1}$ between the frames. This planar affine transformation, defined by 6 parameters, only accounts for the dominant motion of the camera without considering pixel-wise accuracy. Dominant motion has the advantage of being computationally simple to estimate, and is generally sufficient for the task of tonal stabilization.

Now, let u_k be a reference frame and u_t ($t > k$) a subsequent frame in the video. Before applying the transformation model in Eq. (10), u_t is warped towards u_k with the accumulated transformation

$$A_{t,k} = A_{t,t-1} \circ A_{t-1,t-2} \circ \dots \circ A_{k+1,k}. \quad (8)$$

We define the set of spatial correspondences $\Omega_{t,k}$ between u_k and u_t as

$$\Omega_{t,k} = \left\{ (\mathbf{x}, A_{t,k}(\mathbf{x})) \in \Omega \times \Omega \mid \frac{1}{3} \sum_c \left((u_k^c(\mathbf{x}) - \overline{u_k^c}) - (u_t^c(A_{t,k}(\mathbf{x}))) - \overline{u_t^c} \right)^2 < \sigma^2 \right\}, \quad (9)$$

where σ^2 is the empirical noise variance (for instance, estimated with the method in [Colom and Buades 2013]). Note that the constraint in Eq. (9) rules out possible motion outliers as well as occluded points (points visible in only one of the frames).

2.3 Motion driven tonal stabilization

Let us describe how keyframes are defined: the first frame of the sequence u_1 is defined as a keyframe. The following frames $u_t, t > 1$ are stabilized with respect to u_1 as long as there are enough correspondences between u_t and u_1 . More precisely, given a method parameter ω to be tuned, a new keyframe is defined as soon as the number of correspondences $\#\Omega_{t,1}$ is lower than $\omega \times \#\Omega$. In other words, a color correction is computed only if there are enough spatial correspondences to define an accurate and valid model. This process is repeated till the end of the sequence.

To ensure a fidelity to the colors of the input sequence $\{u_t\}_{t=1,\dots,D}$, we also introduce a regularization term for the transformation sequence $\{T_t(u_t)\}_{t=1,\dots,D}$ as following:

$$T_t' = \lambda T_t + (1 - \lambda) Id, \quad (10)$$

where $\lambda = \lambda_0 \exp(-\frac{\|V_{t,k}\|}{p})$ controls the amount of transformation of frame u_t , according to the motion amplitude between frame u_t and the keyframe, $\|V_{t,k}\|$ denotes the norm of the dominant motion vector $V_{t,k}$ and p is the maximum spatial displacement (number of rows + number of columns of the image), λ_0 is the initial weight (in practice we set $\lambda_0 := 0.9$). In other words, the smaller the motion magnitude, the closer λ is to 1 and the greater the transformation. In the following, this temporally-varying regularization is useful to correct the sequence while avoiding overexposure when huge changes in camera exposure occur in the original sequence.

Algorithm 1 sketches the proposed method. Note that the computation of $\Omega_{t,k}$ involves the computation of $A_{t,k}$ and the warping of u_t towards u_k based on $A_{t,k}$.

⁴http://en.wikipedia.org/wiki/ASC_CDL

Algorithm 1 Motion driven tonal stabilization

Input: Sequence of frames $\{u_t\}_{t=1,\dots,D}$
Output: Tonal stabilized sequence $\{T'_t(u_t)\}_{t=1,\dots,D}$

- 1: $k \leftarrow 1, t \leftarrow k + 1$
- 2: $T'_1(u_1) = u_1$
- 3: **while** $t \leq D$ **do**
- 4: Compute $\Omega_{t,k}$
- 5: **if** $\#\Omega_{t,k} \geq \omega \times \#\Omega$ **then**
- 6: **for** $c \in \{R, G, B\}$ **do**
- 7: $\alpha_c, \gamma_c \leftarrow \arg \min_{\alpha, \gamma} \sum_{(\mathbf{x}, \mathbf{y}) \in \Omega_{t,k}} (u_k^c(\mathbf{x}) - \alpha u_t^c(\mathbf{y})^\gamma)^2$
- 8: $T'_t(u_t^c) \leftarrow \lambda \alpha_c (u_t^c)^{\gamma_c} + (1 - \lambda) u_t^c$
- 9: **end for**
- 10: $t \leftarrow t + 1$
- 11: **else** # If there are not enough correspondences
- 12: $k \leftarrow t - 1$
- 13: $u_k \leftarrow T'_{t-1}(u_{t-1})$
- 14: **end if**
- 15: **end while**

For the sake of complexity, the original frames are rescaled (120 pixels wide) before estimating T' . Then, the estimated T' is coded into a Look-Up-Table (LUT) that is applied to the high resolved original frames. These implementation changes guarantee an algorithm with low complexity but do not produce noticeable loss in tonal stabilization accuracy.

3. EXPERIMENTAL RESULTS

In this section we first study the goodness of fit of our power law model. Then, we show the experimental results obtained with our tonal stabilization algorithm and we compare them with state-of-the-art results. On the one hand, qualitative evaluation based on visual inspection is performed and, on the other hand, quantitative results measuring the amount of tonal variation in the resulting sequence and the fidelity to the original sequence are provided. Both, quantitative and qualitative evaluation prove that the proposed algorithm is accurate and robust with all the tested sequences independently of the amount of tonal instabilities or motion.

3.1 Goodness of fit

In order to evaluate the accuracy of our power law model, we have estimated the mean R^2 (coefficient of determination) along color channels:

$$R^2 = \frac{1}{3} \sum_c \left(1 - \frac{\sum_{\mathbf{x} \in \Omega_p} (\log u_k^c(\mathbf{x}) - T_t^c(\log u_t^c(\mathbf{x})))^2}{\sum_{\mathbf{x} \in \Omega_p} (\log u_k^c(\mathbf{x}) - \log u_t^c(\mathbf{x}))^2} \right). \quad (11)$$

where Ω_p is the set of points selected from a color chart in the image. In particular, we consider images captured with a smartphone from the same scene containing a Macbeth color chart (see Fig. 2). Each picture is adjusted (using the camera settings) to have a different exposure or white balance (WB), so that we can analyze the tonal changes by studying the color transfer function between the reference picture (sunlight WB, medium exposure) and the other ones. More specifically, we use the median color value of each color in the color chart to estimate a power law transformation. In the right column of Fig. 2, we plot the functional relationship (in logarithmic domain) between the colors extracted from the color chart of the reference picture and the correspondent colors from the other

pictures. As an indication of goodness of fit, the computed R^2 value is shown for each plot (the closer is R^2 to 1, the better the observed tonal transformation fits the model). The coefficient is larger than 0.9 for all the computed regressions, which shows that the relationship between reference and test color intensities is approximately linear in a logarithmic scale and in general the model fits the data.

Note that the images with the color chart are useful to evaluate our model but they are not enough challenging to compare our results with other methods in the case of videos. In fact, all methods are comparable with this data since the presence of all the colors in the chart help the camera automatic white balance algorithm to work properly, without producing strong tonal instabilities.

3.2 Qualitative evaluation

In practice, our method has been tested on 18 different video sequences. While some sequences have been kindly provided by the authors of [Farbman and Lischinski 2011], we have completed our dataset with video sequences acquired with smart phones from different manufacturers. Complete video sequences (originals and results) are available at the project website⁵. We strongly recommend the reader to look at the electronic version of the paper and the videos in our website to appreciate the results.

First, we have considered video sequences in which camera motion is not complicated as in the sequence "sofa" (see Fig. 3). In the original sequence one same object appears with different colors (e.g., sofa) while in our resulting sequence all colors are stable.

In order to evaluate our results with respect to state-of-the-art methods, we have considered the methods of Farbman et al. [Farbman and Lischinski 2011] and Wang et al. [Wang et al. 2014]. In our comparison, the results from [Wang et al. 2014] have been provided by the authors while the results from [Farbman and Lischinski 2011] come from our implementation of their method, which is coherent with the results published in their paper and website. In particular, Fig. 4 compares our results with the sequence "building" that has been acquired with a Samsung Galaxy S smart phone. Note that the results from [Farbman and Lischinski 2011] are not perfectly stabilized due to the important camera motion of the sequence, and the results from [Wang et al. 2014] are stabilized but the sequence is much whiter and parts of it are completely saturated which is not visually pleasant. Fig. 5 compares with the sequence "graycard" the same algorithms which turn out to have the same behavior in terms of remnant tonal variation for [Farbman and Lischinski 2011] and wash-out look (white) for [Wang et al. 2014]. On the contrary, in the two sequences "building" and "graycard", our results are both stable and color coherent with a good dynamic range.

Note that all video sequences are processed with the same set of parameters, i.e., $\omega = 0.25$ and $\lambda_0 = 0.9$. Concerning the parameter robustness in our method we have observed that varying ω has very little impact on the results. On the contrary, the choice of λ (cf. Eq. 10) determines the amount of fidelity to the original sequence. Indeed, $\lambda = 1$ corresponds to a strict tonal stabilization and $\lambda = \lambda_0 \exp(-\frac{\|V_{t,k}\|}{p})$ weights temporally (in function of motion) the contribution to the original sequence. While temporal weighting reduces the strict tonal preservation, it may be argued that it produces a visually pleasant result by preserving some degree of the original colors. In particular, when the original sequence presents high tonal variations, temporal weighting avoids to create a resulting sequence too different from the original in which colors

⁵http://oriel.github.io/tonal_stabilization

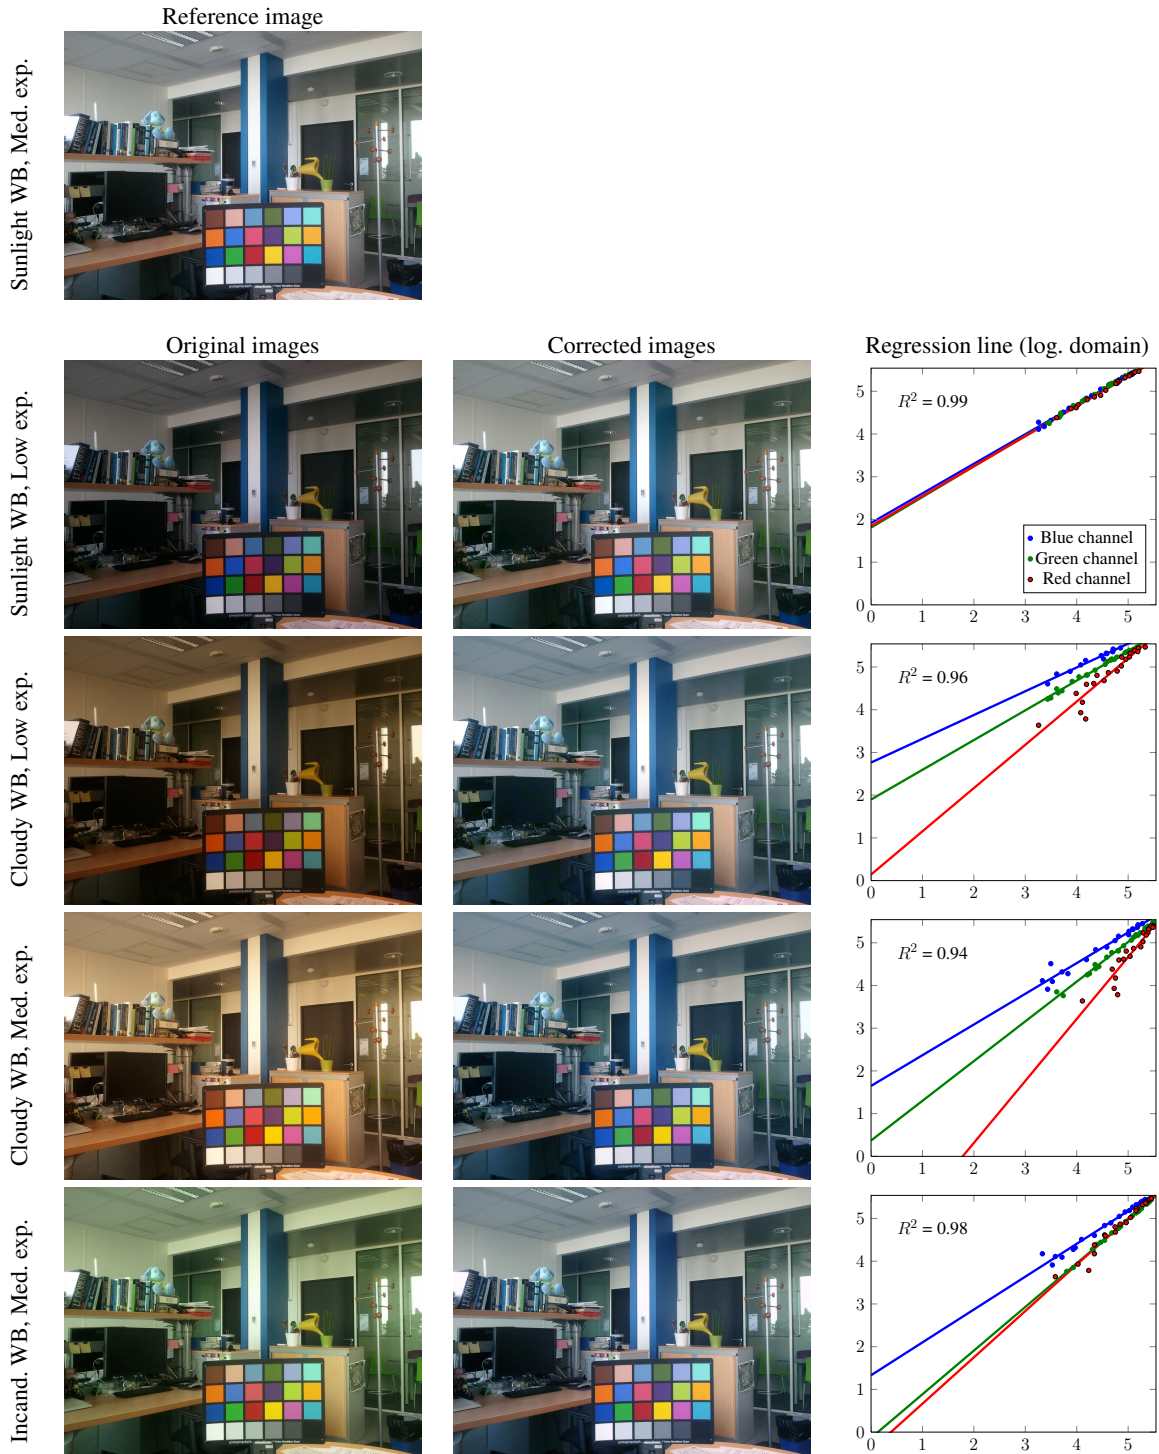


Fig. 2. Illustration of the goodness of fit of the proposed tonal transformation model. In this experience, a sequence of approximately registered images is taken from the same scene, and the color transformation is estimated from correspondent points extracted from the Macbeth color chart. On the first row, a reference image (keyframe u_k) is shown. For the following second to fifth rows, on the left column, tonally unstable images (u_t) taken with different exposures and white balance are shown. On the middle column, the color corrected images are shown, where it can be noted that tonal instability is largely reduced. Finally, on the right column, the extracted points from the color chart and the estimated regression lines are plotted in the logarithmic domain. The x -axis corresponds to $\log u_t^c(\Omega_p)$, while the y -axis corresponds to $\log u_k^c(\Omega_p)$ for the plotted points and $T_t^c(\log u_t^c(\Omega_p))$ for the plotted lines, where Ω_p is the set of color chart coordinates. The regression line has a close fit to the color points, reminding that R^2 values which are close to 1 are an indication of a good fit.

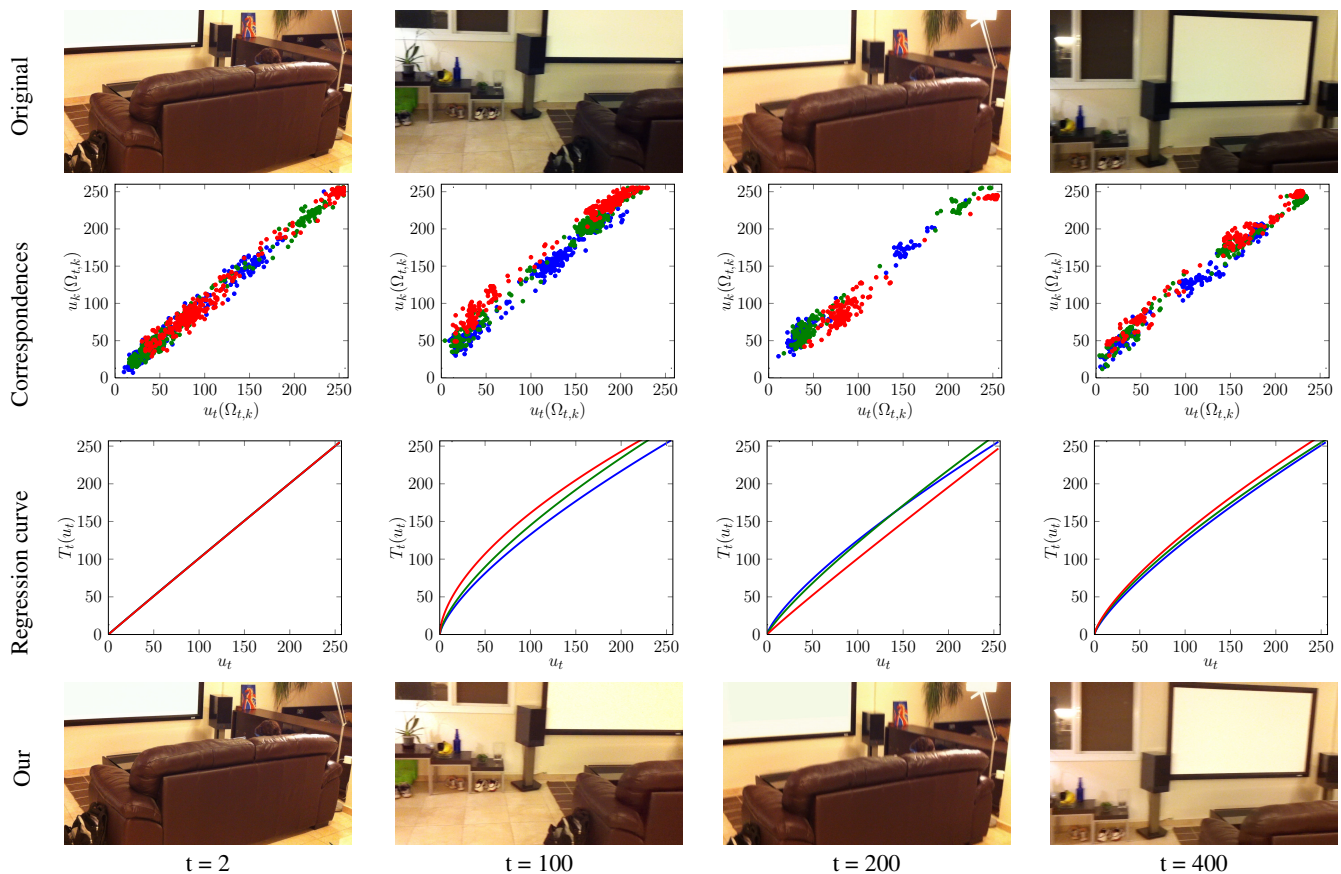


Fig. 3. Tonal stabilization of the sequence “sofa”. Top row: frames extracted from the original sequence, $t = (2, 100, 200, 400)$. Second row: plot of point correspondences between the original frame u_t and the keyframe u_k . Third row: estimated power law tonal transformation for each frame. Bottom row: same frames from top row, after tonal stabilization with our method. Note that objects appearing with different colors in the original sequence have the same color in our results. The color of the plotted points and curves correspond to the sRGB color channel (Red, Green, Blue) of the image.

may appear saturated or not natural. Fig. 6 shows an example of a sequence with important tonal instabilities in the original sequence and the results with the different choices of λ . The strict tonal stabilization result ($\lambda = 1$) saturates some parts of the resulting frames (e.g. $t = 170$) which motivates our choice of λ .

3.3 Quantitative evaluation

In an effort to quantitatively assess the performance of our algorithm we propose to study the tonal variation of a homogeneous patch with respect to the reference (first) frame through the sequence. This is, considering the resulting sequence, we compute the color differences (in CIELAB color space) between a homogeneous patch in the reference frame P_0 and its corresponding patches through the resulting sequence P_t , $t = 1, \dots, D$. Ideally, the patch tonal variation remains constant and equal to zero. However, this criterion is not sufficient to evaluate a tonal stabilization algorithm. For instance, a resulting sequence of completely homogeneous color frames would satisfy this criterion but would not be a good (pleasant) result. Because of this reason we also study the fidelity to the original sequence by computing the color difference between the same aforementioned patches P_t and the same patches on the original sequence P_t^o , for all $t = 1, \dots, D$. A resulting sequence with a large deviation from the original sequence would

produce undesired artifacts. With these two criteria being defined we consider a tonal stabilized sequence being a good result when the patch tonal variation is as constant and small as possible and *at the same time* the fidelity to the original sequence is as much preserved as possible.

Obviously, the two error curves (tonal variation and fidelity to original) can be computed, provided the video sequence has a homogeneous patch, as it is the case for the sequence “building” or the sequence “greycard”. Fig. 7 shows the error curves for these two sequences. We observe that the patch tonal variation is reduced with our method and the method of [Wang et al. 2014] when compared to the patch tonal variation of the original sequence but this is not the case for the results of [Farbman and Lischinski 2011]. Also, our method produces the closest results in terms of color fidelity to the original sequence. Notice that for the sequence “graycard” the fidelity to original is smaller for [Wang et al. 2014] between $t=75$ and $t=150$. We explain this behavior because we choose the first frame as reference. Indeed we obtain a smaller fidelity to original for the first frames (from $t=1$ to $t=75$), but then, when there is a big instability of the original sequence for these frames (see red curve of the tonal variation) our algorithm compensates this difference. As we have explained in Fig. 6 we believe that our weighting strategy provides more natural and artifact-free results.



Fig. 4. Comparison of tonal stabilization for the sequence “building” with the methods of [Farbman and Lischinski 2011] and [Wang et al. 2014]. This figure shows three frames of the sequence ($t=1,40,80$). Our algorithm is able to stabilize tonal variations without generating artifacts, while the results from [Farbman and Lischinski 2011] are not perfectly stabilized and the results of [Wang et al. 2014] tend to saturate parts of the building.

3.4 Implementation

Besides the qualitative and quantitative evaluation our method is also more efficient in comparison to the state-of-the-art. Our prototype implementation in Python processes a $1920 \cdot 1080$ resolution video in a rate of 11 frames per second considering image reading and writing and 20 frames per second without reading and writing⁶. On the contrary, the C++ implementation of [Wang et al. 2014] processes 1 frame per second (depending on video length) and a Python implementation of [Farbman and Lischinski 2011] processes 0.6 frames per second. We believe that an optimized implementation of our method could approach real time processing which is a major advantage and proves the feasibility of embedding robust tonal stabilization algorithms on smart phones.

4. CONCLUSION

In this work, we have proposed an efficient tonal stabilization method, aided by global motion estimation and a parametric tonal

transformation. We have shown that a simple six-parameters color transformation model is enough to provide tonal stabilization caused by automatic camera parameters, without the need to rely on any *a priori* knowledge about the camera model.

The proposed algorithm is robust for sequences containing motion, it reduces tonal error accumulation by means of long-term tonal propagation, and it does not require high space and time computational complexity to be executed.

One of the main advantages of the proposed method is that it could be applied in practice as an online algorithm, that has potential for real time video processing. The actual un-optimized software implementation is already near real-time processing. In addition, the proposed algorithm is suitable for an implementation on a chip, opening applications such as tonal compensation for video conferences or for live broadcast.

Experiments have demonstrated that the method performs favorably compared to state-of-the-art methods, both in terms of stabilization quality, fidelity to original colors, computational complexity and memory footprint.

⁶Processed by Intel(R) Core(TM) i5-3340M CPU @ 2.70GHz, 8GB RAM

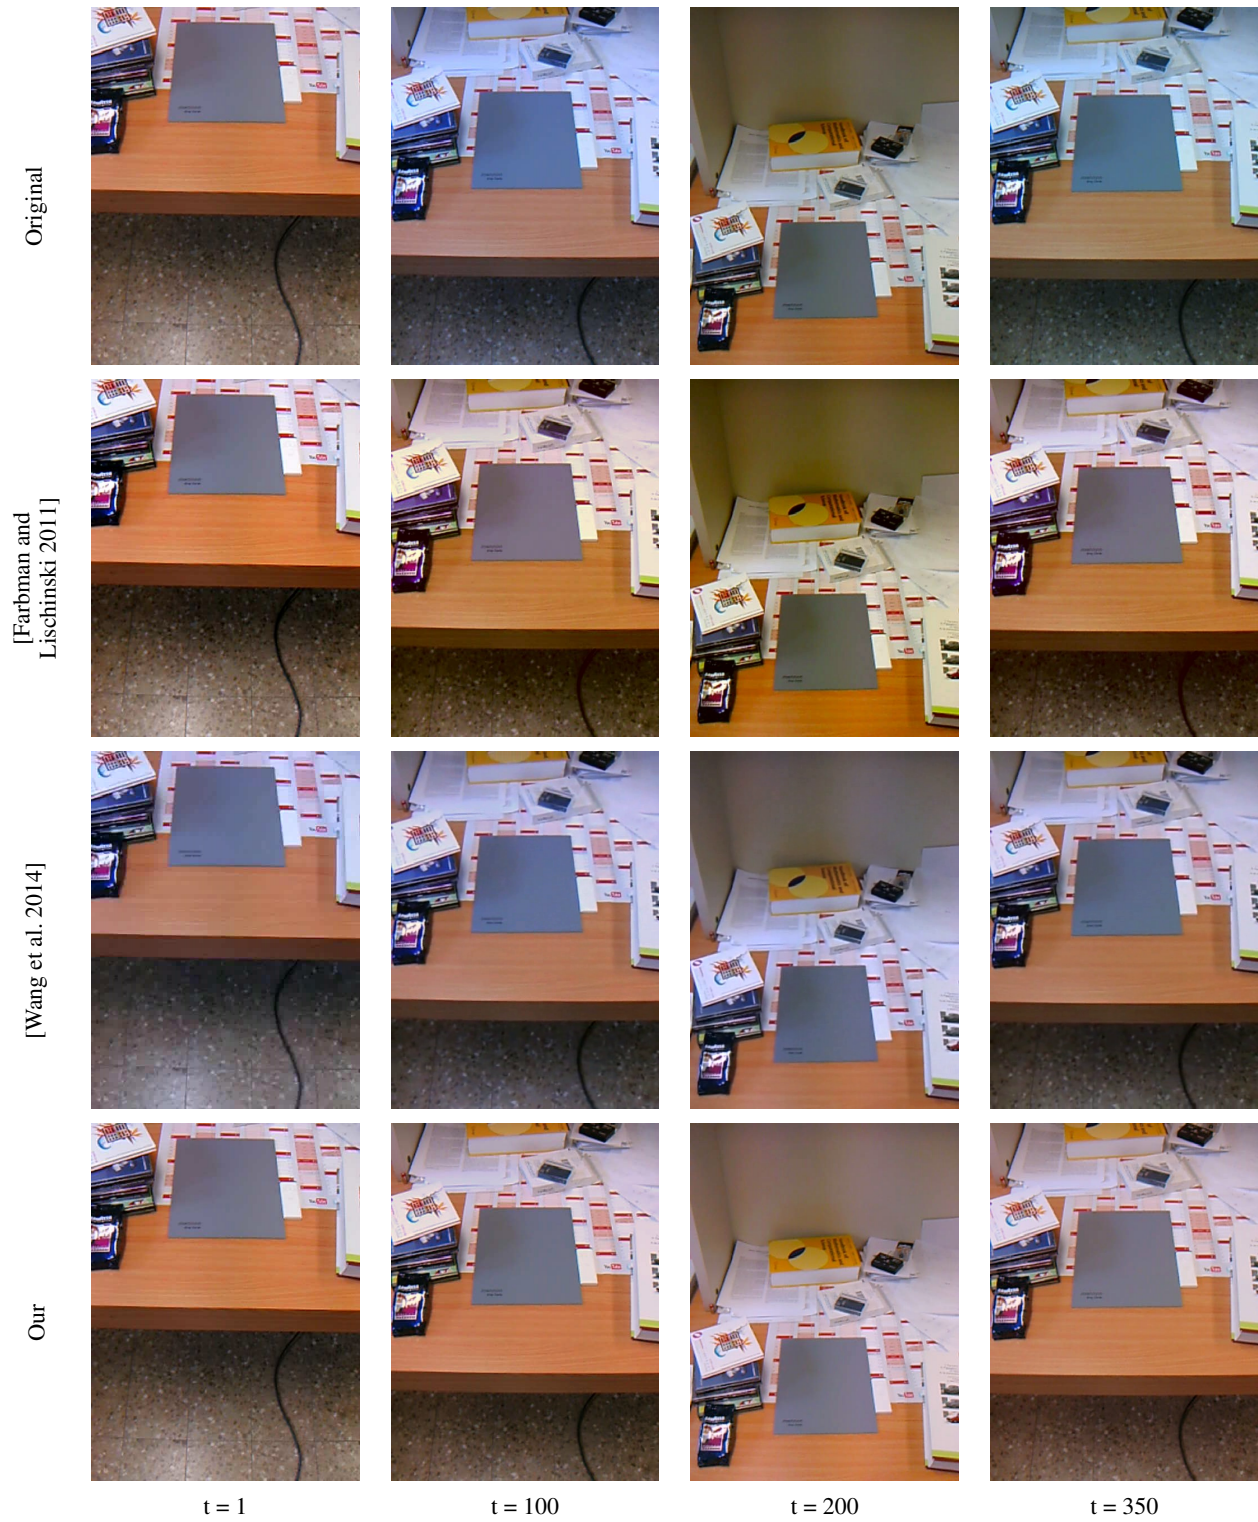


Fig. 5. Tonal stabilization of the sequence “graycard”. First row: frames extracted from the original sequence ($t=1, 100, 200, 350$). Second row: results from [Farbman and Lischinski 2011]. Notice the yellowish color for $t=200$. Third row: results from [Wang et al. 2014] with a wash-out appearance. Bottom row: our stabilized results without any visual artifact.

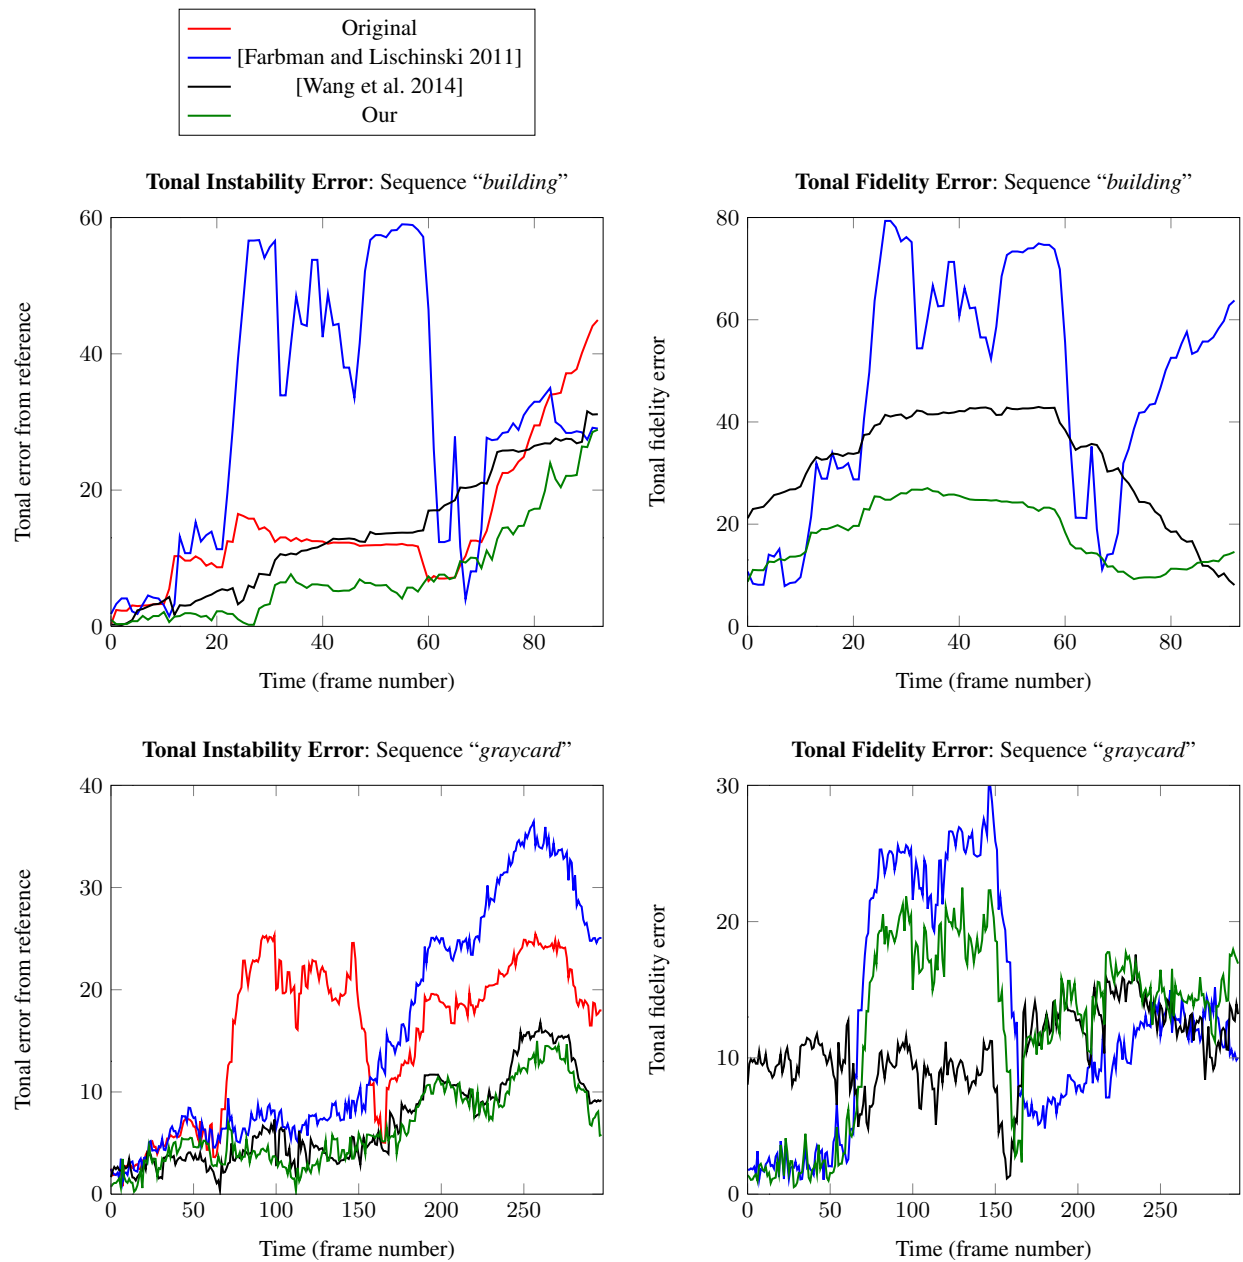


Fig. 7. Quantitative evaluation of the sequence “building” (top) and “graycard” (bottom). For each sequence, we show (i) the tonal instability error, computed as the color distance of a tracked patch to the reference (first) frame, and (ii) the tonal fidelity error, computed as the color distance at each instant, between the corrected frame and the original frame, which indicates the degree of fidelity between the original and the corrected sequence. The color distances are computed as the euclidian distance in perceptual color space CIELAB. Overall, our method compares favorably with the methods of [Farbman and Lischinski 2011] and [Wang et al. 2014], both in terms of *reduction of tonal instability* as in terms of *fidelity to original colors*.

ACKNOWLEDGMENTS

We thank the authors of [Wang et al. 2014] for providing their results in our sequences and [Farbman and Lischinski 2011] for helping with implementation details of their method.

REFERENCES

- BRAINARD, D. H. AND WANDELL, B. A. 1986. Analysis of the retinex theory of color vision. *JOSA A* 3, 10, 1651 – 1661.
- CHAKRABARTI, A., SCHARSTEIN, D., AND ZICKLER, T. 2009. An Empirical Camera Model for Internet Color Vision. *Proceedings of the British Machine Vision Conference 2009*, 51.1–51.11.

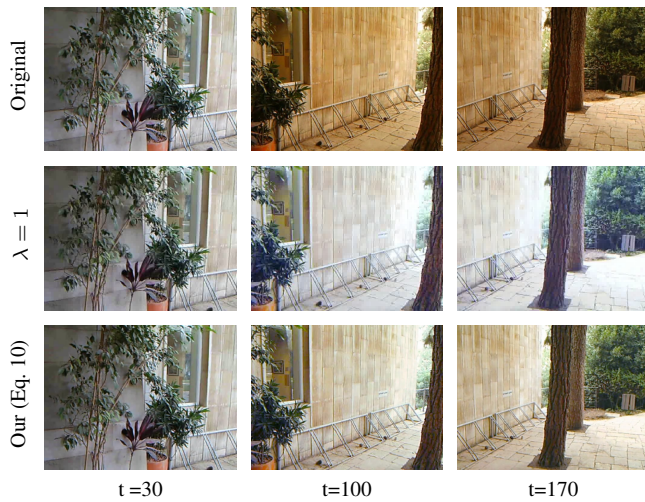


Fig. 6. Illustration of tonal stabilization in sequence “entrance”. First row: Frames extracted from the original sequence. Second row: stabilizing with $\lambda = 1$ (no temporal weighting) ensures tonal coherence between all the frames, but at the cost of clipping intensities. Bottom row: Tonal stabilization with temporal weighting (λ decreases exponentially in function of motion).

COLOM, M. AND BUADES, A. 2013. Analysis and Extension of the Percentile Method, Estimating a Noise Curve from a Single Image. *Image Processing On Line* 3, 332–359.

DECENCIÈRE, É. 1997. Theses. Ph.D. thesis, École Nationale Supérieure des Mines de Paris.

DELON, J. 2006. Movie and video scale-time equalization application to flicker reduction. *Image Processing, IEEE Transactions on* 15, 1 (Jan), 241–248.

DELON, J. AND DESOLNEUX, A. 2010. Stabilization of flicker-like effects in image sequences through local contrast correction. *SIAM Journal on Imaging Sciences* 3, 4, 703–734.

FARBMAN, Z. AND LISCHINSKI, D. 2011. Tonal stabilization of video. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2011)* 30, 4, 89:1 – 89:9.

FERRADANS, S., PAPADAKIS, N., RABIN, J., PEYR, G., AND AUJOL, J.-F. 2013. Regularized discrete optimal transport. In *Scale Space and Variational Methods in Computer Vision*, A. Kuijper, K. Bredies, T. Pock, and H. Bischof, Eds. Lecture Notes in Computer Science, vol. 7893. Springer Berlin Heidelberg, 428–439.

FREEDMAN, D. AND KISILEV, P. 2010. Object-to-object color transfer: Optimal flows and smp transformations. In *IEEE Computer Vision and Pattern Recognition*. 287–294.

FRIGO, O., SABATER, N., DEMOULIN, V., AND HELLIER, P. 2014. Optimal transportation for example-guided color transfer. In *Proceedings of Asian Conference on Computer Vision*.

GIJSENIJ, A., GEVERS, T., AND WEIJER, J. V. D. 2011. Computational color constancy: Survey and experiments. *Image Processing, IEEE Transactions on* 20, 9, 2475 – 2489.

GROSSBERG, M. AND NAYAR, S. 2002. What Can Be Known about the Radiometric Response from Images? In *Computer Vision - ECCV 2002*. 189–205.

HACOHEN, Y., SHECHTMAN, E., GOLDMAN, D. B., AND LISCHINSKI, D. 2011. Non-rigid dense correspondence with applications for image enhancement. *ACM Trans. Graph.* 30, 4 (July), 70:1–70:10.

HACOHEN, Y., SHECHTMAN, E., GOLDMAN, D. B., AND LISCHINSKI, D. 2013. Optimizing color consistency in photo collections. *ACM Trans. Graph.* 32, 4 (July), 38:1–38:10.

HORDLEY, S. D. 2006. Scene illuminant estimation: Past, present, and future. *Color Research & Application* 31, 4, 303 – 314.

KIM, S. J., LIN, H. T., LU, Z., SÜSSTRUNK, S., LIN, S., AND BROWN, M. S. 2012. A new in-camera imaging model for color computer vision and its application. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34, 2289–2302.

LAND, E. H. 1977. The Retinex Theory of Color Vision The Retinex Theory of Color Vision. 237, 6.

LAND, E. H., MCCANN, J. J., ET AL. 1971. Lightness and retinex theory. *Journal of the Optical society of America* 61, 1, 1 – 11.

LEVENBERG, K. 1944. A method for the solution of certain non-linear problems in least squares. *Quart. J. Appl. Maths.* II, 2, 164–168.

LIN, H., KIM, S. J., SUSSTRUNK, S., AND BROWN, M. S. 2011. Revisiting radiometric calibration for color computer vision. In *Proceedings of the IEEE International Conference on Computer Vision*. 129–136.

MARQUARDT, D. W. 1963. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *SIAM Journal on Applied Mathematics* 11, 2, 431–441.

MURRAY, N., SKAFF, S., MARCHESOTTI, L., AND PERRONNIN, F. 2012. Toward automatic and flexible concept transfer. *Computers & Graphics* 36, 6, 622–634.

NEUMANN, L. AND NEUMANN, A. 2005. Color style transfer techniques using hue, lightness and saturation histogram matching. In *Computational Aesthetics in Graphics, Visualization and Imaging 2005*, B. G. W. P. L. Neumann, M. Sbert, Ed. 111–122.

ODOBEZ, J. AND BOUTHEMY, P. 1995. Robust multiresolution estimation of parametric motion models. *Journal of visual communication and image representation* 6, 4, 348–365.

PAPADAKIS, N., PROVENZI, E., AND CASELLES, V. 2011. A variational model for histogram transfer of color images. *IEEE Transactions on Image Processing* 20, 6, 1682–1695.

PITIÉ, F., DAHYOT, R., KELLY, F., AND KOKARAM, A. 2004. A new robust technique for stabilizing brightness fluctuations in image sequences. In *Statistical Methods in Video Processing*. Springer, 153–164.

PITIÉ, F., KOKARAM, A. C., AND DAHYOT, R. 2007. Automated colour grading using colour distribution transfer. *Comput. Vis. Image Underst.* 107, 1-2, 123–137.

REINHARD, E., ADHIKMIN, M., GOOCH, B., AND SHIRLEY, P. 2001. Color transfer between images. *IEEE Computer Graphics and Applications* 21, 5 (sep/oct), 34–41.

WANG, Y., TAO, D., LI, X., SONG, M., BU, J., AND TAN, P. 2014. Video tonal stabilization via color states smoothing. *IEEE transactions on image processing* 23, 11, 4838–4849.

WU, F., DONG, W., KONG, Y., MEI, X., PAUL, J.-C., AND ZHANG, X. 2013. Content-based colour transfer. *Computer Graphics Forum*, no–no.

XIONG, Y., SAENKO, K., DARRELL, T., AND ZICKLER, T. 2012. From pixels to physics: Probabilistic color de-rendering. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. 358–365.