

Modeling the Performance of Image Restoration from Motion Blur

Giacomo Boracchi and Alessandro Foi

Abstract—When dealing with motion blur there is an inevitable trade-off between the amount of blur and the amount of noise in the acquired images. The effectiveness of any restoration algorithm typically depends on these amounts, and it is difficult to find their best balance in order to ease the restoration task. To face this problem, we provide a methodology for deriving a statistical model of the restoration performance of a given deblurring algorithm in case of arbitrary motion. Each restoration-error model allows us to investigate how the restoration performance of the corresponding algorithm varies as the blur due to motion develops.

Our modeling treats the point-spread-function trajectories as random processes and, following a Monte-Carlo approach, expresses the restoration performance as the expectation of the restoration error conditioned on some motion-randomness descriptors and on the exposure time. This allows to coherently encompass various imaging scenarios, including camera shake and uniform (rectilinear) motion, and, for each of these, identify the specific exposure time that maximizes the image quality after deblurring.

Index Terms—Motion Blur, Camera Shake, Deconvolution, Image Deblurring, Imaging System Modeling.

I. INTRODUCTION

MOTION blur and noise are strictly related by the exposure time: photographers, before acquiring pictures of moving objects or dim scenes, always consider whether motion blur may occur (e.g., due to scene or camera motion), and carefully set the exposure time. The trade-off is between long exposures that reduce the noise at the cost of increasing the blur, and short exposures that reduce the blur at the cost of increasing the noise. Often there is no satisfactory compromise, and the captured image is inevitably too blurry or too noisy.

Many approaches to digital restoration of motion-blurred images [1]–[7] can be reduced to deblurring an image where the point-spread function (PSF) is known, by leveraging a standard deconvolution algorithm (e.g., [8]–[11]). While the actual PSF is often unpredictable, there are nevertheless several techniques to estimate the PSF by analyzing either the captured image alone [5], [6], [12], accelerometers and gyroscopes output [13], or additional extremely noisy images acquired with a short exposure [1], [2], [4]. Approaches relying on a single blurred image often enforce iterative procedures that alternate PSF estimation and image deblurring [14]–[16], whereas the restored image is always obtained by deconvolution of the observation and the final estimate of the PSF.

Giacomo Boracchi is with the Dipartimento di Elettronica e Informazione, Politecnico di Milano, Italy

Alessandro Foi is with the Department of Signal Processing, Tampere University of Technology, Finland

The restoration performance of any deblurring algorithm is determined by several concurrent factors: perhaps, the most relevant one is the exposure time, as this balances the amount of blur and noise in the observations. It is important to emphasize that the blurred image to be restored is never free of noise, and even a small amount of noise may compromise the effectiveness of the deblurring algorithm. Other elements, such as the scene or the specific trajectory generating the PSF, may significantly influence the restoration performance but these, differently from the exposure time, are difficult (if not impossible) to control in advance.

In our previous work [17] we studied how the performance of any image deconvolution algorithm varies w.r.t. the exposure time in the special case of rectilinear blur, for which an analytical formulation can be obtained. Here, we provide a more general result, addressing arbitrary motion blur. Our core contribution is a methodology for deriving a statistical model of the restoration error of a given deblurring algorithm in case of arbitrary motion, including random motion. More specifically, each restoration-error model describes how the expected restoration error of a particular image-deblurring algorithm varies as the blur due to camera motion develops over time along with the PSF trajectory, which we effectively handle by means of statistical descriptors. The peculiarity of the proposed methodology is that it simultaneously takes into account the exposure time, its interplay with the sensor noise, and the *motion randomness*.

Each specific restoration-error model allows us to identify the proper acquisition strategies that maximize the performance of the corresponding deblurring algorithm. In particular, in controlled imaging scenarios where the evolution of the PSF trajectory along with the exposure time can be statistically studied or analytically formulated, the restoration-error model can tell whether there exists an optimal exposure, i.e. an exposure time that minimizes the restoration error achievable by the corresponding deblurring algorithm; then, whenever the optimal exposure time exists, the restoration-error model provides its value. This issue, to the best of our knowledge, has so far been neglected, mainly because of the unpredictability of the PSF trajectory.

The proposed methodology is general, and customized restoration-error models can be derived for any type of deblurring algorithm although in what follows we focus on convolutional blur, thus mainly considering image deconvolution. Furthermore, it is convenient to decouple blur estimation from blur removal, as these two problems are typically faced by different algorithms that may behave differently w.r.t. the motion development: the proposed methodology is then mostly suited for non-blind deconvolution algorithms.

As an illustration of the proposed methodology, we model the restoration performance of three relevant deconvolution algorithms: the Anisotropic LPA-ICI Deblurring [10], the Deconvolution using Sparse Natural Image Priors [7], and the Richardson-Lucy Deconvolution [8]. Each derived restoration-error model expresses the expected restoration error of the corresponding algorithm as a function of the exposure time and the PSF standard deviations [18]: two simple, yet effective, descriptors of the blur PSF. We test these models in two situations: the uniform-motion blur and the camera-shake. In the former case, it is possible to describe the relation between the PSF standard deviations and the exposure time with an analytical expression, while in the latter case this relation follows from a statistical formulation. The resulting optimal exposure times agree with the acquisition strategies adopted in practice to cope with camera shake, and with our independent theoretical and experimental analysis of uniform-motion blur presented in [17].

The effectiveness of the proposed methodology is validated by comparing the outputs of these three restoration-error models against results of experiments on camera raw data, revealing that the actual performance of the corresponding algorithms follow, qualitatively as well as quantitatively, the trends outlined by their restoration-error models.

The remainder of the paper is organized as follows. We first present the related works and we outline some applications of our modeling approach. Then, in Section II we present the image formation model and we state the problem. The methodology for deriving restoration-error models is presented in Section III and, as meaningful examples, in Section IV we compute the restoration-error models for the Anisotropic LPA-ICI Deconvolution [10], the Deconvolution using Sparse Natural Image Priors [7], and the Richardson-Lucy Deconvolution [8]. Section V presents the results of deblurring a large dataset of camera raw images corrupted by uniform-motion and camera-shake blur. The proposed approach is validated in Section VI, by comparing the outputs of the three restoration-error models previously computed with the restoration performance measured on the considered dataset of raw images. Discussions concerning the different blur/noise trade-off for camera-shake and uniform-motion blur, as well as the model practical applicability are reported in Section VII, while concluding remarks are given in Section VIII.

A. Related Works

To put our contribution in perspective, let us briefly summarize some of the most important related works, where ad-hoc devices and controlled or customized acquisition strategies are devised to ease the restoration task. Differently from image stabilization techniques, which counteract/prevent the blur, most computational-photography techniques leverage particular acquisition strategies (or settings) that make the algorithmic inversion of the blurring operator easier. These algorithms can be divided into two classes: the first class consists of algorithms that couple the blurred image with some additional information [1], [3], [4], [19], [20], while the second class consists of algorithms that tweak the camera acquisition [21]–[25] to obtain PSFs that are easier to invert.

The first class of algorithms includes [3], [26], which exploit hybrid imaging systems (provided with two cameras having different resolutions) that are able to measure their own motion during the acquisition. The blur PSF is then computed from these motion information, and the blur is inverted using the traditional Richardson-Lucy Deconvolution [8]. Other techniques do not require ad-hoc hardware, and exploit images acquired with different exposures [1], [4]. These works focus on camera shake and pair a long-exposure image, which is dominated by blur, with a short-exposure one, which is corrupted by overwhelming noise: the short-exposure image is treated as blur-free, and used for computing the blur PSF. Differently, the algorithm proposed in [19] focuses on rectilinear PSF, and combines several blurred images acquired with different exposure times to compensate the frequencies suppressed by blur in each observation.

Algorithms of the second class aim at actively controlling the camera during the acquisition, thus piloting the resulting PSF, so that the blur inversion becomes a well-conditioned problem. In [21], [23] it is shown that the motion blur can be effectively handled by fluttering the camera shutter during the acquisition, following a coded-exposure. Such a coded exposure makes the resulting blur easier to invert. Other solutions consist of moving the camera (or the camera sensor) according to a parabolic motion during the exposure [24]: by combining the sensor and target motion one obtains a blur that can be inverted using a single PSF.

Our contribution can be an aid to techniques that rely on an estimate of the PSF for the deblurring, by supporting them with guidelines to design the acquisition settings for maximizing the restoration performance. These techniques include algorithms of the first class described above, as well as fully blind algorithms such as [5] and [6].

II. PRELIMINARIES

A. Observation Model

We model an image z_T acquired with an exposure time T as

$$z_T(x) = \kappa (u_T(x) + \eta(x)), \quad x \in X, \quad (1)$$

where $X \subseteq \mathbb{R}^2$ is the sampling grid and $\kappa > 0$ is a factor that can be used for scaling the signal into a usable (limited) dynamic range, thus mimicking the amplification gain in digital sensors (typically, $\kappa \propto T^{-1}$). The two terms $u_T(x)$ and $\eta(x)$ are independent random variables distributed as

$$\begin{aligned} u_T(x) &\sim \mathcal{P} \left(\lambda \int_0^T y(x - s(t)) dt \right), \\ \eta(x) &\sim \mathcal{N} (0, \sigma^2), \end{aligned} \quad (2)$$

where \mathcal{P} and \mathcal{N} denote respectively the Poisson and Gaussian distributions, and $\lambda > 0$ is a parameter characterizing the quantum efficiency of the sensor [27]. The function $y : \mathbb{R}^2 \rightarrow [0, M]$, $0 < M < \infty$, represents the original image, having range $[0, M]$ (the boundedness of y is always verified in practical applications). As this work focuses on motion blur, the function s is assumed to be a curve $s : [0, T] \rightarrow \mathbb{R}^2$,

which identifies the apparent motion taking place between the scene and the imaging sensor during the exposure time. In what follows we refer to s as the PSF trajectory, or simply the trajectory.

The motion blur in (2) is modeled by a linear and shift-invariant operator, as s is the same for the whole image. Thus, by using generalized functions, the argument of the Poisson distribution in (2) can be rewritten as

$$\begin{aligned} \lambda \int_0^T y(x - s(t)) dt &= \lambda \int_0^T (y \otimes \delta_{s(t)})(x) dt = \\ &= \lambda \left(y \otimes \int_0^T \delta_{s(t)} dt \right) (x) = \lambda (y \otimes h_T)(x), \end{aligned} \quad (3)$$

where $\delta_{s(t)}$ denotes the Dirac delta function at $s(t) \in \mathbb{R}^2$ and h_T is the motion blur PSF:

$$h_T(\cdot) = \int_0^T \delta_{s(t)}(\cdot) dt. \quad (4)$$

Since $h_T > 0$, we have

$$\int_{\mathbb{R}^2} h_T(x) dx = T. \quad (5)$$

Although the PSFs are often assumed to have unit mass, in our model their mass, given by the integral in (5), equals the exposure time T : in such a way we take into account how the signal expectation varies with the exposure time. The parameter κ in (1) will eventually take care of the normalization. Thus, (2) can be rewritten as

$$u_T(x) \sim \mathcal{P}(\lambda (y \otimes h_T)(x)), \quad x \in X. \quad (6)$$

Since we restrict to motion blur, other time-independent blurs such as out-of-focus or blur due to camera optics are neglected. Furthermore, we concentrate on PSFs resulting from trajectories characterized by random motions.

Note that the definition of z_T in (1) includes actually two noise terms: the time-dependent (and image-dependent) noise $E\{u_T\} - u_T$ inherent to the photon-acquisition process and modeled by the Poissonian distribution, and the time-independent (and image-independent) noise η , which accounts for electric and thermal noises and that is modeled by a Gaussian distribution. The signal-to-noise ratio at x , i.e. $\text{SNR}(z_T(x))$ becomes

$$\text{SNR}(z_T(x)) = \frac{E\{z_T(x)\}}{\text{std}\{z_T(x)\}} = \frac{\lambda (y \otimes h_T)(x)}{\sqrt{\lambda (y \otimes h_T)(x) + \sigma^2}}, \quad (7)$$

where the numerator corresponds to (3) and the denominator can be derived after trivial calculations. When both y and σ are zero, the $\text{SNR}(z_T(x))$ is formally defined as zero. Since $\text{SNR}(z_T(x))$ is monotonically increasing w.r.t. $\lambda (y \otimes h_T)(x)$, the average of $\text{SNR}(z_T(x))$ on the whole image is also monotonically increasing w.r.t. T . This implies that the corresponding observations become less noisy when the exposure time increases, which however does not necessarily imply that they improve, since, unfortunately, the amount of blur typically increases with T .

Probably, the best known example of observations modeled by (1) are the images acquired during the shake of a hand-held camera, which becomes particularly relevant at low-light conditions and in dim environments, where long exposure times are typically used to reduce the noise in the observations.

B. Problem Statement

The aim of this work is to model how, in presence of motion blur, the restoration error of a given deblurring algorithm varies as the motion develops with the exposure time. In practice, the performance of every algorithm depends on several elements other than the exposure time, and in particular it depends on the original image y , on the noise parameters, and on the trajectory s generating the PSF. While the noise parameters and the original image can be considered as deterministic elements and as such fixed (as for example when one takes a picture of a particular scene, with a particular camera, and with known acquisition settings), the PSF h_T cannot be typically predicted and thus it needs to be treated as a random process. Therefore, the analysis of the restoration error shall be carried out in a statistical manner, and for this purpose we pursue a Montecarlo approach where we assume that the PSFs h_T are drawn from a generating distribution \mathcal{H}_T (i.e. $h_T \sim \mathcal{H}_T$). It follows that our model does not provide the restoration error achieved/achievable for a specific trajectory s or PSF h_T , but as the expected restoration error over \mathcal{H}_T , which in practice can be computed as the average error over a large number of random motion PSFs drawn from \mathcal{H}_T . To make such expectation useful and meaningful, we shall condition it on few descriptors of these PSFs. The restoration-error model, which is the main objective of our analysis, is an estimate of the conditional expectation of the restoration error given some meaningful PSFs descriptors.

III. RESTORATION-ERROR MODEL: THE METHODOLOGY

Here are the key elements required to build a restoration-error model for a given deblurring algorithm \mathcal{D} :

- A collection H_T of m motion PSFs that are representative of the blur that may occur in the considered application scenario. These PSFs are considered as independently drawn from \mathcal{H}_T , in the sense that

$$H_T = \{h_T^j \text{ s.t. } h_T^j \sim \mathcal{H}_T, j = 1, \dots, m\}. \quad (8)$$

- The restoration error r , which is the metric for measuring the deblurring performance. This metric is a scalar function of the restored image and of the original image (when available) such as the RMSE, the PSNR, as well as perceptual metrics as the MSSIM [28]. In computing these metrics, we assume that the images are normalized to a standard intensity range (e.g., [0,255]), in order to get consistent error measurements across potentially very different ranges of intensities.
- An operator Θ that, for each PSF, provides a vector of n descriptors $[\theta_1, \dots, \theta_n]$, which can be effectively used to condition the expected restoration error. In order to obtain a meaningful restoration-error model, the values assumed by the PSF descriptors have to faithfully encompass




Fig. 1. Standard deviations associated to four motion-blur PSFs: a) $\varsigma_L = 5.71$ and $\varsigma_S = 0.75$, b) $\varsigma_L = 5.65$ and $\varsigma_S = 1.53$, c) $\varsigma_L = 4.91$ and $\varsigma_S = 3.08$, d) $\varsigma_L = 5.95$ and $\varsigma_S = 3.20$, e) $\varsigma_L = 5.44$ and $\varsigma_S = 5.26$. The red segments (principal axes) have length $2\varsigma_L$ and the green segments (minor axes) have length $2\varsigma_S$. Values are expressed in pixels.

the characteristics of the PSF that mostly influence the performance of \mathcal{D} .

- A statistical method \mathfrak{E} to estimate the conditional expectations from several realizations of the restoration errors on H_T . For instance, \mathfrak{E} can be an adaptive parametric or non-parametric smoother.

The restoration-error model consists in the expected restoration error r over $h_T \sim \mathcal{H}_T$, conditioned w.r.t. the PSF descriptors $\theta_1, \dots, \theta_n$, i.e.

$$R_{\lambda, \sigma, y}(\theta_1, \dots, \theta_n, T) = \underset{h_T \sim \mathcal{H}_T}{E} \{r(z_T) | \theta_1, \dots, \theta_n\}, \quad (9)$$

where the dependency of z_T on h_T (and on T , λ and σ) is expressed in (1) and (6). Note that λ , σ and y in (2) are considered fixed. In what follows, the conditional expectation is always intended given the PSFs descriptors.

A. Computation of the Restoration-Error Model

The restoration-error model is computed by estimating the expectation in (9) using \mathfrak{E} to smooth the restoration errors measured on the outputs of \mathcal{D} , when \mathcal{D} is applied on observations generated from PSFs in H_T (8). This means that (9) becomes

$$R_{\lambda, \sigma, y}(\theta_1, \dots, \theta_n, T) = \underset{h_T \in H_T}{\mathfrak{E}} \{r(z_T) | \theta_1, \dots, \theta_n\}, \quad (10)$$

where by $\underset{h_T \in H_T}{\mathfrak{E}}$ we indicate that the smoothing \mathfrak{E} takes place over the restoration errors obtained $\forall h_T \in H_T$. More specifically, \mathcal{D} is applied on observations that are synthetically generated according to (1) where h_T varies in H_T , y is a test image, λ and σ are fixed.

Note that the smoothed values of the restoration errors result in an approximation of the conditional restoration error defined for any arbitrary value of the descriptor $[\theta_1, \dots, \theta_n]$. The smoothed values corresponding to a fixed time instant T , i.e. $R_{\lambda, \sigma, y}(\dots, T)$, is referred to as the *restoration-error surface* at the exposure time T . Similarly, the restoration-error surfaces can be computed for any arbitrary value T' , defining $H_{T'}$ by scaling the norms of the PSFs in H_T . In practice, to define a restoration-error model, it suffices to consider a selected finite set of exposure times and resort to interpolation for the restoration-error surfaces corresponding to different values of T .

B. Evolution Through Time

While the restoration-error model $R_{\lambda, \sigma, y}$ considers the PSF descriptors $\theta_1, \dots, \theta_n$ and the exposure time T as independent, these in practice are always related because, when we consider motion blur, the PSF evolves over time. Thus, assume that the descriptors can be expressed as function of the exposure time $\theta_i = \theta_i(T)$, $i = 1, \dots, n$: these functions might be estimated empirically or, in some specific applications, these can be determined analytically. Then, upon substitution, the restoration-error model $R_{\lambda, \sigma, y}$ (10) becomes a univariate function of time:

$$R_{\lambda, \sigma, y}(T) = R_{\lambda, \sigma, y}(\theta_1(T), \dots, \theta_n(T), T). \quad (11)$$

As an important consequence, we can use $R_{\lambda, \sigma, y}(T)$ to analyze how the restoration performance of \mathcal{D} varies with respect to the exposure time. Thus, in practice, $R_{\lambda, \sigma, y}(T)$ provides a guideline to choose, before the acquisition, the exposure time yielding, after restoration, a higher image quality. Whenever such exposure time is unique and finite, we refer to it as the *optimal exposure time* T^* , formally defined as

$$T^* = \underset{T}{\operatorname{argmin}} [R_{\lambda, \sigma, y}(T)]. \quad (12)$$

C. Different Images

So far we have presented how to compute the restoration-error model for a given image y , which however is always unknown in deblurring applications. One way to circumvent this issue is to enforce an image prior, for example by computing the expected restoration error for multiple images drawn from the same prior distribution, and not for a specific image y . In practice, this means that \mathfrak{E} is applied to $r(z_T)$, computed from a bunch of images. Nevertheless such approach is rather cumbersome, as the amount of images to be drawn for accurately modeling the prior is potentially huge.

A more appealing situation would arise when the qualitative behavior of the surfaces does not significantly change with respect to the original image; the ideal situation being when different images yield the same surfaces modulo an additive constant or scaling factor because in this case the optimal exposure time, T^* of (12), is the same regardless of the specific image y . While this condition might seem improbable at a first glance, as a matter of fact it has been already experimentally verified in [17] for the case of uniform-motion blur. Whenever this condition holds, it is enough to smooth

using \mathfrak{E} the restoration errors computed from observations generated by different original images, which are conveniently chosen having the same range $[0, M]$, then (10) becomes

$$R_{\lambda,\sigma,M}(\theta_1, \dots, \theta_n, T) = \mathfrak{E}_{h_T \in H_T} \{r(z_T) | \theta_1, \dots, \theta_n\}, \quad (13)$$

being z_T generated from a set of original images having the same range $[0, M]$.

D. Model Portability

So far, all the assumptions made imply that the surfaces refer to the behavior of the restoration-error r for a specific deblurring algorithm \mathcal{D} , for specific noise parameters λ and σ , and for original images in a specific range $[0, M]$. Note that a linear scaling of the image range by multiplication of y against an arbitrary constant $c > 0$ is formally equivalent to using a sensor having noise parameters $c\lambda$ and σ , due to the linearity of the operations in (2). It follows that

$$R_{\lambda,\sigma,M}(\theta_1, \dots, \theta_n, T) = R_{\lambda M, \sigma, 1}(\theta_1, \dots, \theta_n, T) \quad (14)$$

thus, the restoration-error models (13) are more conveniently parametrized, upon normalization of the image range, by the product λM and σ only. Equation (14) shows that it is possible to avoid the re-computation of the surfaces when the observations are in different ranges, a situation that would arise, for instance, because of light changes in the scene.

Likewise, it can be shown that the exposure time T and the parameter λ can be interchanged. Indeed, a restoration-error surface $R_{\lambda,\sigma,M}(\dots, T')$ corresponding to a specific exposure time T' is formally equivalent to a restoration-error surface $R_{\lambda',\sigma,M}(\dots, T)$, where $\lambda' = \lambda T/T'$. This is because in (6) the factor λ scales h_T (having norm T from (5)), exactly as T'/T scales each PSF in the set H_T to yield the set $H_{T'}$. It follows that

$$R_{\lambda,\sigma,M}(\theta_1, \dots, \theta_n, T) = R_{\lambda M T, \sigma, 1}(\theta_1, \dots, \theta_n, 1). \quad (15)$$

Therefore, up to scaling the values of T , the same surfaces can be used to describe how the restoration error varies on sensors characterized by different values of λ . This is not surprising since the same observation z_T can be interpreted as acquired by a more quantum-efficient sensor having parameter 2λ and moving two times faster along the same fixed trajectory (thus with an exposure time $T/2$).

From (14) and (15) follows that the same restoration-error model allows us to describe the restoration performance on observations having different ranges, as well as on sensors having different values of λ , up to accordingly scaling the exposure time. More precisely, we have that, for any $\lambda' > 0$ and $M' > 0$,

$$R_{\lambda' M', \sigma, 1}(\theta_1, \dots, \theta_n, T) = R_{\lambda M, \sigma, 1}(\theta_1, \dots, \theta_n, \frac{\lambda' M'}{\lambda M} T). \quad (16)$$

Let us remark that the equations (14)-(16) are valid for a fixed value of σ . There is no direct scaling equation that relates restoration-error surfaces corresponding to different values of σ . Therefore, different restoration-error models have to be computed for different values of σ . Additional details

concerning the use of these models in practical applications are detailed in Section VII-C.

Unless otherwise noted, in what follows we assume a normalized image range and, for the sake of notation, we do not specify the unitary image range in the more compact notation $R_{\lambda M, \sigma}$.

E. Desiderata

In principle, it is possible to derive a restoration-error model (10) for any deblurring algorithm \mathcal{D} and any choice of H_T , r , Θ and \mathfrak{E} , by following the procedure described in Section III. Nevertheless, the resulting model does not necessarily provide meaningful results since, for instance, the PSF descriptors Θ may not properly condition the restoration error, the set of considered PSFs H_T may not fully represent the blur of the considered problem, or the smoother \mathfrak{E} may not provide reliable estimates of the expected restoration error. To obtain a meaningful restoration-error model it is then convenient to check whether the following conditions hold:

- C1 The PSF descriptor Θ has to correctly interpret the way the restoration error varies when varying the PSF. This means that PSFs having *similar descriptor values* yield *similar restoration errors* when the observations are generated from the same original image y . This allows the expected restoration error to be rightly conditioned on these descriptors.
- C2 The distributions of the restoration errors (for a fixed y) must be as localized as possible about their conditional expectation given $[\theta_1, \dots, \theta_n]$.
- C3 Finally, according to the considerations from Section III-C, it is desirable that the qualitative behavior of the conditioned restoration errors is not significantly altered when changing the original image y .

As an illustrative example, in what follows we derive the restoration-error models for three specific deconvolution algorithms. The proposed methodology is then extensively validated by comparing the trends of the predicted restoration errors with the RMSEs measured when restoring, with the corresponding deconvolution algorithms, a large dataset of camera raw images. The experiments show that outcomes of these models allow to reliably determine the best exposure to be used in practice.

IV. RESTORATION-ERROR MODEL: THREE EXAMPLES

In this section we compute the specific restoration-error models for three different deconvolution algorithms \mathcal{D} . These models are meant to illustrate how the proposed methodology can be exploited in practice, and they allow us to assess its effectiveness, as shown in Section VI. Before briefly recalling these algorithms, we introduce the restoration error r underlying the three models.

A. Restoration Error

We define the restoration error r associated to an observation z_T as the *ideal* root mean squared error (RMSE)

$$r(z_T, y) = \frac{255}{M} \frac{\|y - \hat{y}_T\|_2}{\sqrt{\#X}}, \quad (17)$$

where $\#X$ is the number of image pixels, and \hat{y}_T stands for the *ideal* restored image. The image \hat{y}_T is obtained by applying \mathcal{D} on z_T when the PSF is exactly known and the regularization parameters of \mathcal{D} have been chosen by an oracle, i.e. they are those that minimize (17). The same error metric has been also utilized in [17] for the case of rectilinear blur. Note also that, since \hat{y}_T represents the ideal restored image, (17) provides an upper-bound of the restoration performance achievable by \mathcal{D} on z_T , which implies that each restoration error model assumes a correct use of the deconvolution algorithm.

B. The Deconvolution Algorithms

The considered deconvolution algorithms are the Anisotropic LPA-ICI Deconvolution for signal-dependent noise [10], the Deconvolution using Sparse Natural Image Priors [7], and the Richardson-Lucy Deconvolution [8].

1) *LPA-ICI Deconvolution for Poissonian Images*: The Anisotropic LPA-ICI Deconvolution for signal-dependent noise¹ [10] relies on a nonparametric Poisson maximum-likelihood modeling and it couples a Tichonov regularized inverse in Fourier domain with an adaptive anisotropic filtering in space domain. Here, capital letters are used to indicate the Fourier transform of the corresponding quantities. The regularized inverse is thus expressed as

$$\hat{Y}_{T,\varepsilon}^{RI} = Z_T \frac{\kappa \overline{H_T}}{\kappa^2 |H_T|^2 + \text{PSD}_T \varepsilon^2}, \quad (18)$$

where $\varepsilon > 0$ is the regularization parameter, and PSD_T is the power spectral density of the noise characterizing observations acquired with exposure time T . It follows that

$$\text{PSD}_T = \sum_{x \in X} \text{var} \{z_T(x)\} = \kappa^2 \sum_{x \in X} (\lambda(y \otimes h_T)(x) + \sigma^2), \quad (19)$$

which is in practice approximated by

$$\text{PSD}_T \approx \kappa \sum_{x \in X} (z_T(x) + \kappa \sigma^2).$$

The filtering is realized in spatial domain, using directional polynomial-smoothing kernels $g_{\theta_i, h^+(\theta_i)}$ having pointwise-adaptive support-size $h^+(\theta_i)$ along the different directions $\{\theta_i\}$: the final restored image \hat{y}_T is computed as

$$\hat{y}_T(x) = \sum_{\theta_i} \beta(x, h^+(\theta_i), \theta_i) \int \hat{y}_{T,\varepsilon}^{RI}(\xi) g_{\theta_i, h^+(\theta_i)}(x - \xi) d\xi,$$

where $x \in X$, $\hat{y}_{T,\varepsilon}^{RI}$ is the inverse Fourier transform of $\hat{Y}_{T,\varepsilon}^{RI}$ (18), and the convex weights $\beta(x, h^+(\theta_i), \theta_i)$ are used to combine the directional estimates into an anisotropic one. For details we refer the reader to [11] and especially to [29].

The ideal RMSE is computed by selecting the parameter ε with an oracle, i.e. by minimizing (17).

2) *Deconvolution Using a Sparse Prior*: This algorithm² [7] formulates the deconvolution problem as determining the maximum a-posteriori estimate of the original image, given the observation z_T . Furthermore, the algorithm exploits a prior enforcing spatial-domain sparsity of the image derivatives. The resulting non-convex optimization problem is solved using an iterative re-weighted least square method. Although this algorithm has not been natively devised for Poissonian observations, it has been rather successfully applied to raw images, thanks to the oracle selection of the smoothness-weight parameter and by allowing a sufficient number of iterations (i.e. 200).

3) *Richardson-Lucy Deconvolution*: This classical deconvolution algorithm³ [30] assumes Poisson-distributed observations and the deconvolved image is obtained from z_T by an iterative expectation-maximization procedure. The ideal RMSE has been computed by selecting with the oracle the number of iterations.

C. PSF Generation

The PSFs constituting the collections H_T , which are used to compute the restoration-error models, are obtained by sampling continuous trajectories on a (regular) pixel grid. Each trajectory consists of the positions of a particle following a 2-D random motion in continuous domain. The particle has an initial velocity vector which, at each iteration, is affected by a Gaussian perturbation and by a deterministic inertial component, directed toward the current particle position. In addition, with a small probability, an impulsive perturbation aiming at inverting the particle velocity arises, mimicking a sudden movement that occurs when the user presses the camera button or tries to compensate the camera shake. At each step, the velocity is normalized to guarantee that trajectories corresponding to equal exposures have the same length. Each perturbation (Gaussian, inertial, and impulsive) is ruled by its own parameter and each set H_T contains PSFs sampled from trajectories generated by parameters spanning a meaningful range of values; rectilinear trajectories are generated when all the perturbation parameters are zero.

Each PSF $h_T \in H_T$ consists in discrete values that are computed by sampling a trajectory on a regular pixel grid, using sub-pixel linear interpolation. Collections corresponding to different exposure times are obtained by scaling the values of each PSF by a constant factor. In our simulations, each set H_T contains $m = 7471$ different PSFs.

Fig. 2 shows an example of a considered trajectory and the corresponding sampled PSF. The code generating the trajectories and the PSFs is publicly available for download⁴.

D. PSFs Standard Deviations

We adopt the PSF standard deviations proposed in [18] as a (bivariate) descriptor Θ ; these can be computed as follows. Each PSF h_T is normalized to unit mass by dividing it by T .

²Available at <http://groups.csail.mit.edu/graphics/CodedAperture>

³Implemented by the Matlab command `deconvlucy`.

⁴Available at <http://home.dei.polimi.it/boracchi/software/>.

¹Available at <http://www.cs.tut.fi/~lasip/>




Fig. 2. An example of PSF trajectory generated from a random motion and the corresponding sampled PSF. This trajectory presents an impulsive variation of the velocity vector, thus mimicking the situation where the user presses the button or tries to compensate the camera shake. Another example of PSF clearly affected by a similar abrupt variation in the trajectory is shown in Fig. 1(e)

The normalized PSF is then treated as a bivariate probability distribution, for which we compute the covariance and hence the standard deviations along its principal axes, ς_L and ς_S . Specifically, for each PSF h_T , the values of ς_L and ς_S are computed as the square root of the eigenvalues of the covariance matrix of the normalized PSFs:

$$(\varsigma_L, \varsigma_S) = \sqrt{\text{Eig} \left(\text{Cov} \left(\frac{h_T}{T} \right) \right)}. \quad (20)$$

Fig. 1 illustrates four examples of PSFs with the corresponding principal axes; the corresponding values of ς_L and ς_S are reported in the figure caption. The main axis follows the direction along which the variance of the distribution is maximized: in each of the PSFs shown Fig. 1 it is represented by the red segment having length $2\varsigma_L$. The minor axis is represented by the green segment having length $2\varsigma_S$; of course, (20) guarantees that the two axis are orthogonal. A similar characterization of the PSFs has been used also in [31].

Fig. 3 illustrates some of the PSFs in H_T positioned in the Cartesian plane so that the center of each PSFs coincides with its values of ς_L and ς_S . As one can notice, there are almost no PSFs having $\varsigma_S < 1/\sqrt{6} \approx 0.4$. The value $\varsigma_S = 0$ would in fact correspond to rectilinear trajectories along horizontal or vertical directions that are perfectly centered with respect to the pixel grid. In any other case, the discretization of the trajectory s yields PSFs having larger values of ς_S . It is possible to model the effects of sampling the continuous trajectories on the regular pixel grid through a bivariate random distribution with independent marginals both having symmetric triangular density on $[-1, 1]$: this formulation justifies the choice of the value $1/\sqrt{6}$, which corresponds the standard deviation of such distribution. Further details are provided in Section VI-B.

E. Restoration-Error Surfaces

The cloud of points in Fig. 4(a) represent the restoration errors obtained by applying the Anisotropic LPA-ICI Decon-




Fig. 3. An example of synthetically generated random motion PSFs. The plot shows 278 PSFs that have been randomly selected within the dataset H_T of cardinality $m = 7471$. Each PSF is drawn in the ς_L, ς_S plane; the center of each PSF is located in the corresponding values of ς_L and ς_S (values expressed in pixels). As one can notice, there are almost no PSFs having $\varsigma_S < 1/\sqrt{6}$. The value $\varsigma_S = 0$ would correspond to rectilinear trajectories along horizontal or vertical directions that are perfectly centered with respect to the pixel grid; in all other cases, the discretization of the trajectory s yields PSFs having larger values of ς_S .

volution on observations synthetically generated from the PSFs in $H_{1/2}$ (corresponding to $T = 1/2$ sec.). Each point in this cloud is determined by a distinct PSF and has coordinates $(\varsigma_L, \varsigma_S, r(z_T, y))$, where ς_L and ς_S are the PSF standard deviations, and $r(z_T, y)$ is the *ideal* restoration error (17). Here $\lambda = 3000$, $\sigma = 0$, and the original image y corresponds to the standard test image *Lena*, normalized so that black equals 0 and white equals 1 (thus $M = 1$). Here and throughout the paper, the smoothing operator \mathfrak{E} is a third order, nonparametric polynomial smoother for data corrupted with Gaussian noise, having adaptive bandwidth defined by the Anisotropic LPA-ICI technique [29]⁵. As stated in Section III-A, the smoothed restoration errors constitute the restoration-error surfaces: the surface $R_{3000,0}(\cdot, \cdot, 1/2)$, computed from the above mentioned restoration errors, is shown in Fig. 4(a) and 4(c).

Fig. 5 illustrates how distinct clouds are separated: the three histograms show the distances between the restoration-error surface at $T = 2$ sec., and

- the cloud of restoration errors at $T = 2$ sec. (solid line),
- the cloud of restoration errors at $T = 1/2$ sec. (dashed line),
- the cloud of restoration errors at $T = 1/8$ sec. (dotted line).

⁵The choice of this smoother is irrespective of the deconvolution algorithm that yielded the cloud of RMSE values.

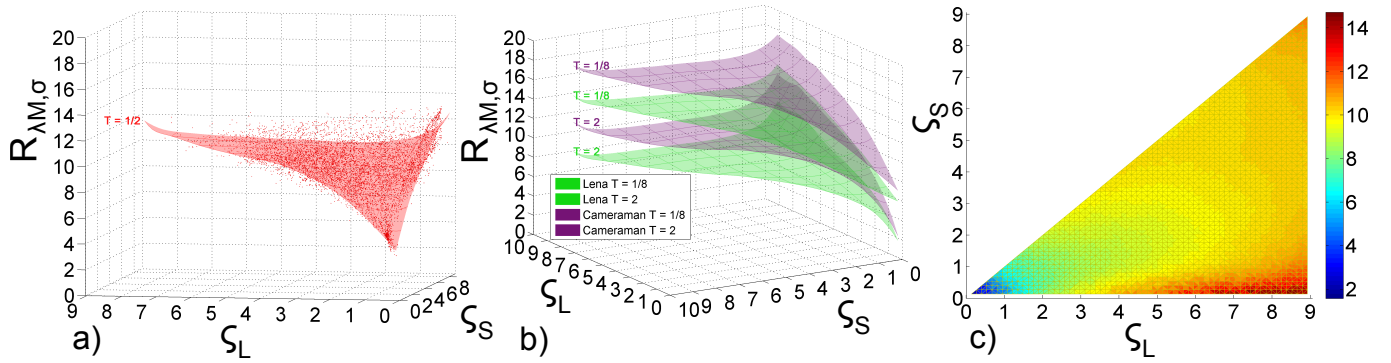


Fig. 4. (a) The restoration-error surface $R_{3000,0}(\cdot, \cdot, 1/2)$, and the cloud of restoration errors $r(z_T, y)$ obtained applying the Anisotropic LPA-ICI Deconvolution on observations generated from *Lena* image (having range $[0, 1]$) and PSFs in $H_{1/2}$, when $\lambda = 3000$, $\sigma = 0$, and $T = 1/2$ sec. (b) The comparison between surfaces corresponding to observations generated from different original images (*Lena* and *Cameraman*) at different exposure times: these surfaces qualitatively show the same behavior, and their differences can be roughly referred to an additive constant term. The label next to each surface indicates the corresponding exposure time. (c) The same surface shown in (a) as seen from above, i.e. displayed in the plane $(\varsigma_L, \varsigma_S)$, and using colors to represent the elevation of the surface. The other considered deconvolution algorithms yield qualitatively similar data and surface.




Fig. 5. Distribution of the residuals of the restoration-error surfaces corresponding to the Anisotropic LPA-ICI Deconvolution when $\lambda = 3000$, $\sigma = 0$, and y is *Lena* image ($M = 1$). The solid (blue) line represents the distribution of the residuals of the restoration-error surface at $T = 2$ sec. The dashed (red) line represents the distribution of the differences between the cloud of restoration errors measured at $T = 1/2$ sec., and the surface at $T = 2$ sec. Similarly, the dotted (black) line represents the distribution of differences between the cloud of errors measured at $T = 1/8$ sec. and the surface at $T = 2$ sec. Similar behaviors characterize the other deconvolution algorithms considered.

The shifts between these histograms show that the dispersion suppressed by the smoothing can be considered as nuisance, which can be neglected when comparing restoration-error surfaces having sufficiently different exposure times. Thus, the clouds are well clustered about their respective surfaces and the dispersion is relatively small: in other words, the surfaces correctly interpret the cloud behavior with respect to ς_L and ς_S . The PSFs standard deviations are indeed particularly effective PSF descriptors as PSFs having *similar* values of ς_L and ς_S yield *similar* restoration errors. Therefore, the restoration-error surfaces can be rightly used as a surrogate of the expected restoration error conditioned on ς_L and ς_S , and the conditions C1 and C2 of Section III-E can be assumed as satisfied. Note also that the distributions of the errors in Fig. 5 are not far from a Gaussian bell, thus confirming that the smoother \mathfrak{E} operated correctly.

Fig. 4(b) compares the surfaces obtained from the test image *Cameraman* with those from *Lena*: note that these surfaces do not differ qualitatively, although they have been computed using essentially different original images. Hence, also the condition C3 of Section III-E is satisfied and, as discussed in Section III-C, we assume that the overall behavior of each restoration-error surface is independent of the original image: the restoration-error surfaces are then computed using the smoother \mathfrak{E} on the restoration errors averaged over four standard test images (*Lena*, *Cameraman*, *Peppers*, *Boat*).

Fig. 4 shows the restoration error surfaces corresponding to the Anisotropic LPA-ICI Deconvolution while the same plots

for the Deconvolution using Sparse Natural Image Priors and the Richardson-Lucy Deconvolution have not been displayed due to space limitation. However, the corresponding surfaces are qualitatively very similar, as shown in Fig. 6.

F. The Restoration-Error Models

Since bivariate PSF descriptors are used, the restoration-error model $R_{\lambda M, \sigma}$ is a trivariate function that expresses how, in a camera having noise parameters λ , σ , the expected restoration error varies with respect to the PSFs standard deviations and the exposure time for images having range $[0, M]$. According to (10), the model is formalized as

$$R_{\lambda M, \sigma}(\varsigma_L, \varsigma_S, T) = \mathfrak{E}_{h_T \in H_T} \{r(z_T, y) | \varsigma_L, \varsigma_S\}. \quad (21)$$

For each deconvolution algorithm in Section IV-B we follow the approach described in Section III-A and we compute the restoration-error surfaces corresponding to a set of values of T that are powers of 2. This specific choice aims at mimicking the exposure stops typically used in photography and in particular, for the Anisotropic LPA-ICI and the Richardson-Lucy Deconvolution we take $T = 2^{-10}, \dots, 2^{10}$ sec., while we had to compute fewer surfaces for the Deconvolution using Sparse Natural Image Priors, due to its much longer computing time. Then, the model outcomes corresponding to different exposure times are obtained by resorting to interpolation.

Fig. 6 shows the surfaces of the three restoration-error models with $\lambda = 3000$ and $\sigma = 0$, and using test images

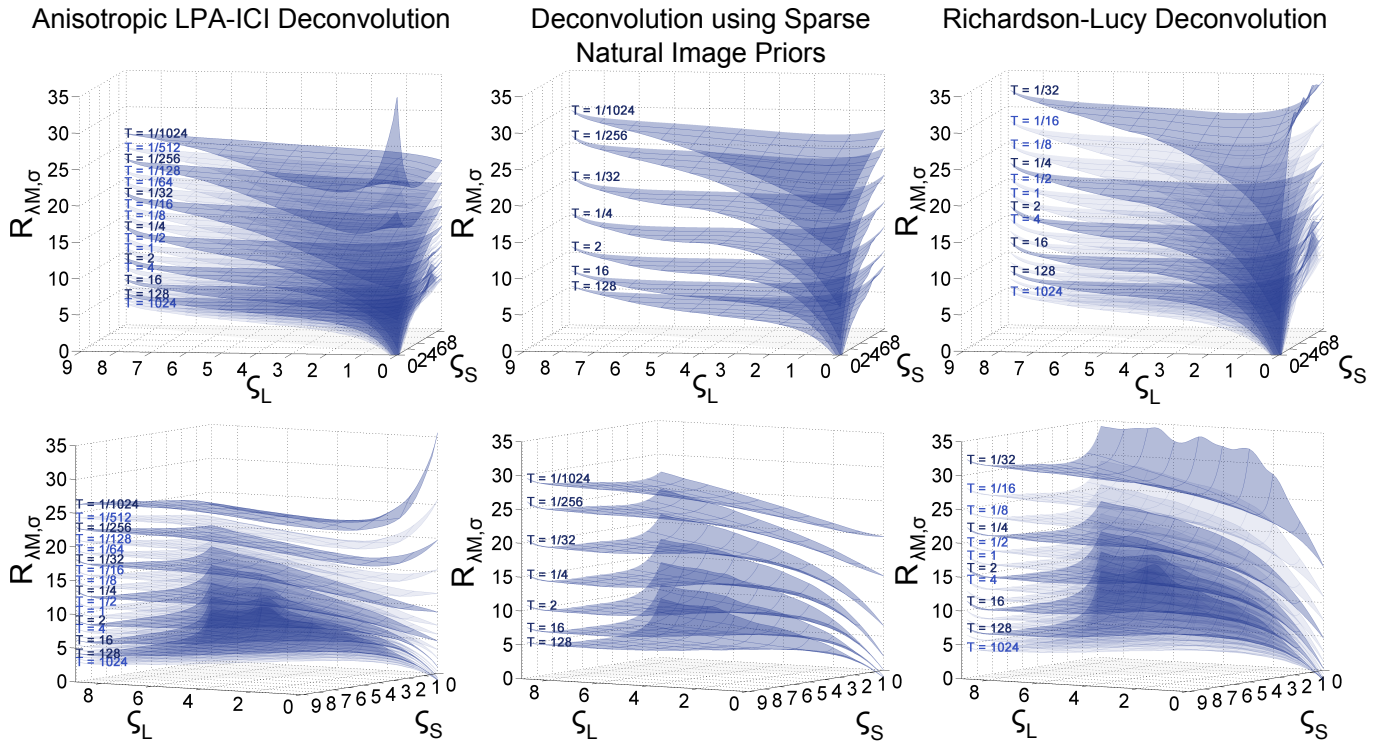


Fig. 6. Each row shows a different view of the restoration-error models for the three deconvolution algorithms. These models have been computed from the restoration errors corresponding to observations generated from test images having range $[0, 1]$ and noise parameters $\lambda = 3000$, $\sigma = 0$. The label next to each surface indicates the corresponding exposure time. To improve visualization, some surfaces are drawn in a darker shade.

having range $[0, 1]$. We wish to remark that, in this figure, the apparent negative RMSE values obtained for $\sigma_L \approx 0$ are due to the extrapolation of the fitted surfaces. In fact, H_T contains nearly pointwise PSFs having values of $\sigma_L \approx 1/\sqrt{6}$, for which exact deconvolution is possible and these yield $\text{RMSE} \approx 0$ when T (and hence the observation SNR) tends to infinity.

V. THE BLUR/NOISE TRADE-OFF AND THE DECONVOLUTION PERFORMANCE ON RAW IMAGES

In this section we consider two specific sorts of motion blur that are often encountered in practical applications: the camera-shake and the uniform-motion blur. To study how the blur/noise trade-off affects the deconvolution performance, we have acquired a large dataset of raw images in a controlled scenario and we have computed the ideal restoration error (17) for the three deconvolution algorithms presented in Section IV-B. In what follows we focus on the experimental settings and on the achieved restoration performance, then in Section VI we show that the outcomes of the restoration-error models derived in Section IV are coherent with the RMSEs computed from the dataset of raw images, thus proving that the proposed methodology is general and provides representative restoration-error models.

A. Raw Image Acquisition

Motion-blurred images have been acquired by fixing a Canon EOS 400D DSLR camera on a tripod, in front of a monitor running a short movie where a natural image is progressively translated along a motion trajectory. While the

movie was playing, we acquired several images that are rightly modeled by (1): the resulting blur is space-invariant since we have accurately positioned the camera to ensure parallelism between the monitor and the imaging sensor (in particular, we verified that in each picture the window displaying the movie was accurately rectangular).

Fig. 7 depicts the six trajectories used to render the movies: four of them represent camera-shake blur (from b to e), while the first one (a) corresponds to the uniform-motion blur⁶. The top row of Fig. 8 shows the four original images used; note that none of these is a standard test image or has been used to derive the restoration-error models in Section IV. For each trajectory/image pair we rendered a 3-second movie and we acquired several pictures according to the settings reported in Table I; a few examples are shown in the bottom row of Fig. 8. The respective ground-truth images, used to compute the restoration errors, have been acquired from the paused movies with 4-second exposure and ISO 100.

We independently process each image, cropping from each channel of the Bayer pattern an observation of 256×256 pixels to be deconvolved. The PSF was estimated via parametric fitting, through the minimization of the RMSE of the restored image, leveraging the knowledge of the continuous trajectory. The camera-raw dataset is finally composed of about 1700⁷ raw images (together with their relative PSF

⁶There is no need to consider uniform motion along different directions since any deblurring algorithm would provide essentially the same restoration quality when the blur direction varies, as discussed in [17].

⁷We excluded the few observations where the PSF estimation failed.

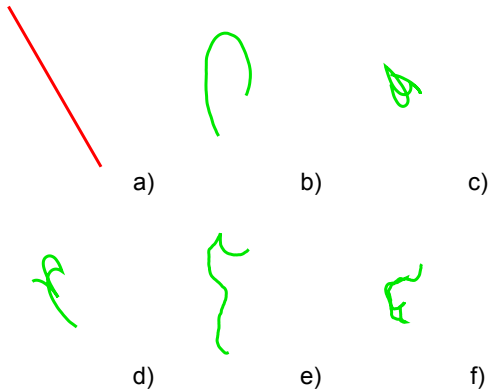


Fig. 7. Trajectories used to generate motions for the experiments on camera raw data. Trajectory (a) corresponds to uniform (rectilinear) motion, while the remaining (b - e) represent blur due to camera-shake.

TABLE I
ACQUISITION SETTINGS OF THE CAMERA RAW IMAGES.

setting	ISO	T	# of shots
1	1600	1/125 s	3
2	1600	1/60 s	3
3	1600	1/30 s	3
4	1600	1/15 s	3
5	1600	1/8 s	3
6	800	1/4 s	3
7	400	1/2 s	3
8	200	1 s	3
9	100	2 s	3
10	100	2.5 s	3

estimates), of which 270 corrupted by uniform-motion blur⁸. For each acquisition, we estimate the noise parameters λ and σ using the algorithm in [27]. Both the Anisotropic LPA-ICI [10] and the Richardson-Lucy [8] Deconvolution explicitly use these estimates, while the Deconvolution using Sparse Natural Image Priors [7] implicitly address the noise model by the selection of weights, which are chosen by an oracle that minimizes (17). Thus, all the restoration errors have been computed under comparable ideal conditions for the three algorithms. This experimental setup is analogous to that used in [17] to deconvolve raw images corrupted by uniform-motion blur.

Fig. 9 shows the ideal restoration errors (17) computed from the dataset of raw images, plotted as a function of the exposure time T . Each row reports the results from the red, green, and blue channels for the three considered deconvolution algorithms: it is important to present the RMSEs separately for each channel as these have very different ranges, thus very different SNR. Each marker refers to the restoration errors from either camera-shake (i.e. related to trajectories b - f) or uniform-motion blur (i.e. related to trajectory a). The two curves are obtained by a third-order nonparametric polynomial regression of the respective restoration errors, expressed w.r.t. the exposure time (to ease the visualization in Fig. 9, we slightly displaced the restoration errors about the true exposure times).

Fig. 10 - 12 illustrate the evolution of the observations and of the corresponding restored images as the blur develops in

⁸The raw images and the estimated PSFs are available at <http://home.dei.polimi.it/boracchi/software/>.

time along three distinct trajectories. These examples show that there is a clear blur/noise trade-off in the observations and that the way it affects the restoration performance significantly depends on the type of PSF trajectory: one can clearly see that in case of blur due to camera shake the considered deconvolution algorithms can achieve good deblurred estimates even at the longest exposures, in contrast to the case of uniform-motion blur, for which the deblurred estimates degrade past relatively short exposures. The same conclusions hold for the three deconvolution algorithms on the three channels of the Bayer pattern, as one can clearly see from the plots of Fig. 9 and the RMSEs values reported in the captions of Figs. 10 - 12.

VI. METHODOLOGY VALIDATION

To validate the proposed methodology for computing restoration-error models, we demonstrate that the outcomes of the restoration-error models obtained in Section IV are consistent with the trend of the RMSE curves computed from the dataset of raw images, shown in Fig. 9.

To this purpose, we need to take care of two aspects. First, the model parameters have to reflect the raw-data acquisition conditions; in particular, these are the noise parameters λ and σ , and the parameter M that specifies the image range. Second, we have to determine suitable expressions of the PSF standard deviations as functions of the exposure time for the considered types of motion blur. Denoting these functions as $\varsigma_L(T)$ and $\varsigma_S(T)$, the computed restoration-error models can be thus expressed according to (11) as univariate functions of the exposure time T :

$$R_{\lambda M, \sigma}(T) = R_{\lambda M, \sigma}(\varsigma_L(T), \varsigma_S(T), T). \quad (22)$$

Then, the validation of the proposed methodology can be carried out by comparing, on a suitable range of exposure times, the functions $R_{\lambda M, \sigma}(T)$ obtained from the restoration-error models of [7], [8], [10], with the RMSE curves in Fig. 9.

A. Selection of Model Parameters

The noise parameters λ and σ have been estimated directly from the raw data by using the procedure [27]. Based on this analysis, we selected $\lambda = 3000$ and, since for these images the signal-independent noise appeared much weaker than the signal-dependent one, we selected $\sigma = 0$. This approximation is supported by the experiments in [17], where it is shown that, at least in the case of uniform-motion blur, different amounts of noise η do not lead to qualitative differences in the trend of the deconvolution performance. The noise parameters are referred to a normalized camera dynamic range $[0, 1]$ over which the intensities of the raw data cover different ranges depending on the channel of the Bayer pattern. In particular, we found the following values for M : $M_R = 0.58$ (red), $M_G = 0.86$ (green), and $M_B = 0.40$ (blue).



Fig. 8. Top row: the natural images (*Balloons*, *Liza*, *Jeep* and *Salamander*) that were used to render the movies for the experiments. Bottom row: examples of the captured raw images demonstrating that there is a clear blur/noise trade-off; these four images have been acquired with settings 10, 8, 6, 4, 2 (see Table I), and come from the trajectory f), d), b), e) and a), respectively. Images were taken from the blue channel and rescaled for visualization purposes.

B. Sampling/using the Restoration Error Model

As stated in Section IV-C, the restoration-error models are constructed using PSFs sampled on a regular pixel grid and the values of the PSF standard-deviations ς_L and ς_S (20) have been computed from these sampled PSFs. Thus, any analytical formulation $\tilde{\varsigma}_L(T)$ and $\tilde{\varsigma}_S(T)$ that expresses on a continuous (i.e. non discrete) domain how the PSF standard deviations vary with respect to time, needs to adequately take into account the effects of sub-pixel interpolation, which become significant at low values of ς_L and ς_S . To this purpose, we model the effects of the sub-pixel linear interpolation, which is used for sampling a continuous trajectory on a regular grid, as an additive random vector following a bivariate distribution with independent marginals, each of which has a symmetric triangular density on the interval $[-1, 1]$. The 2×2 covariance matrix of this distribution is equal to $\mathbb{I}/6$. It follows that the formulas relating the standard deviations of the sampled PSFs to the exposure time (i.e. $\varsigma_L(T)$ and $\varsigma_S(T)$) can be approximated as the sum of the eigenvalues of the covariance matrix of the PSFs standard deviations in continuous domain (i.e. $\tilde{\varsigma}_L(T)$ and $\tilde{\varsigma}_S(T)$) with the eigenvalues of the covariance matrix of the distribution modeling the sub-pixel interpolation:

$$\begin{aligned}\varsigma_L(T) &= \sqrt{\tilde{\varsigma}_L(T)^2 + 1/6}, \\ \varsigma_S(T) &= \sqrt{\tilde{\varsigma}_S(T)^2 + 1/6}.\end{aligned}\quad (23)$$

Note that (23) implies that there exist a portion of the ς_L, ς_S plane, corresponding to the rectilinear PSFs (i.e. $\varsigma_S < 1/\sqrt{6}$) that should be practically empty, as shown in Fig. 3. We can in fact observe that even straight PSFs have values of ς_S which are strictly positive (and not zero, as one would expect from the formulation of the PSFs in continuous domain). We remark that when sampling the restoration-error model at a point $(\varsigma_L(T), \varsigma_S(T), T)$, like in (22), the values of ς_L and ς_S need to be those that correspond to the sampled PSFs, thus it is necessary to use (23) to compensate the functions $\tilde{\varsigma}_L(T)$ and $\tilde{\varsigma}_S(T)$, which have been instead derived from continuous

domain formulations.

C. PSF Standard Deviations Evolution in Time

In this section we present how to express, in the specific cases of uniform and camera shake blur, the PSFs standard deviations as functions of time, i.e. as the two functions $\varsigma_S(T)$ and $\varsigma_L(T)$.

1) *Camera Shake*: In [18], a statistical study of several blurred images acquired at different exposure times by different users (or groups of users), shows that the following power law reliably expresses the standard deviations as functions of the exposure time in case of camera shake:

$$\begin{aligned}\tilde{\varsigma}_L(T) &= \gamma \cdot a_L T^b, \\ \tilde{\varsigma}_S(T) &= \gamma \cdot a_S T^b,\end{aligned}\quad (24)$$

where $\gamma > 0$ is a conversion factor between degrees (which is the unit used in [18]) and pixels (which are used here). It is found that the value $b = 0.5632$ can be rightly used in every experimental condition considered in [18], while the values of a_L and a_S vary depending on the user photographic skills and on the camera mass. We choose $a_L = 0.092$ and $a_S = 0.0453$, which correspond to a DSLR camera [18]. Given (23) and (24), we can then compute a robust estimate of γ as

$$\hat{\gamma}^2 = \text{med} \left\{ \left[\frac{\left(\varsigma_L^{(j)} \right)^2 - \frac{1}{6}}{a_L^2} \left(\frac{1}{T^{(j)}} \right)^{2b} \right]_j \quad \text{II} \quad \left[\frac{\left(\varsigma_S^{(j)} \right)^2 - \frac{1}{6}}{a_S^2} \left(\frac{1}{T^{(j)}} \right)^{2b} \right]_j \right\}, \quad (25)$$

where the index j runs through the dataset of acquired raw images and $\varsigma_L^{(j)}, \varsigma_S^{(j)}, T^{(j)}$ are the measured values of the PSF standard-deviation and the exposure time for a the j -th raw

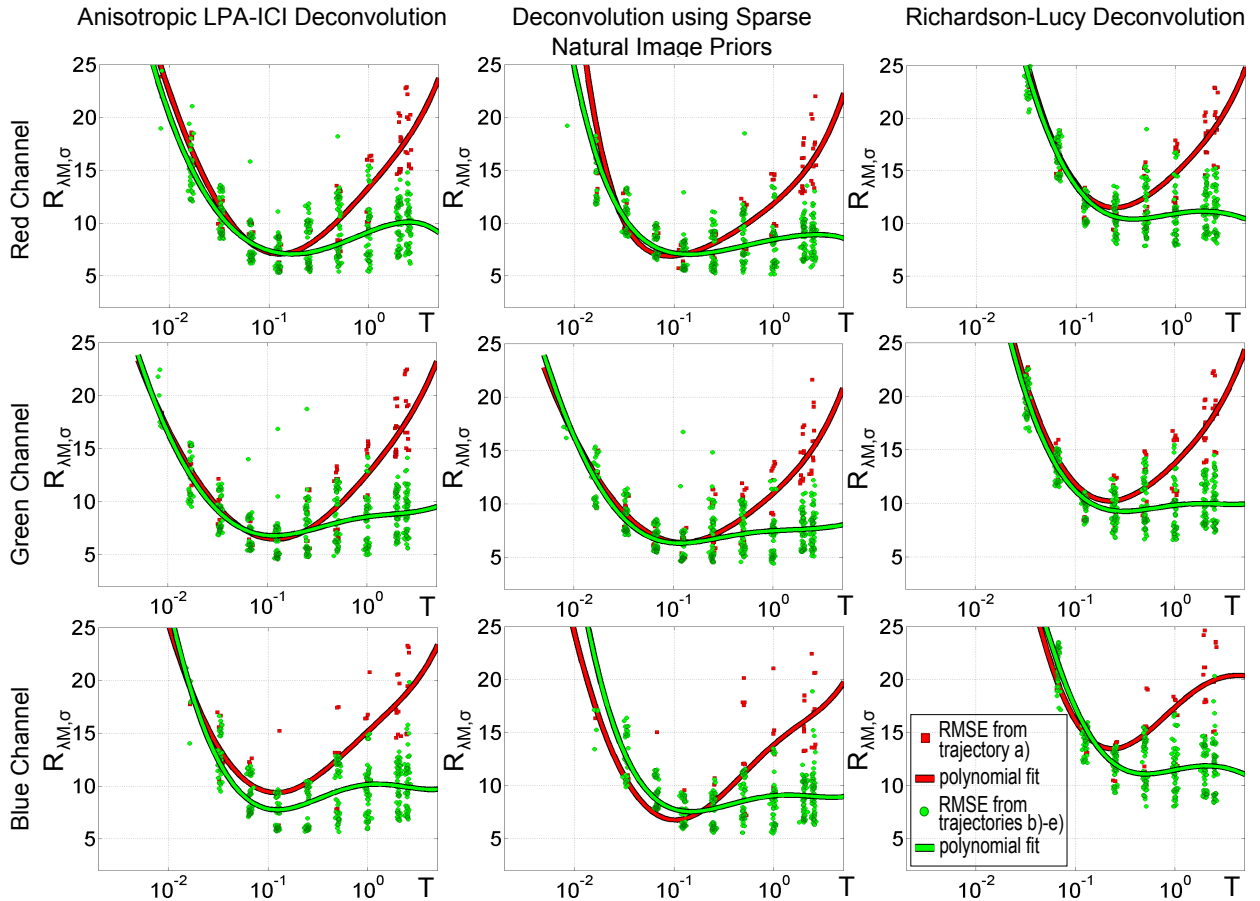


Fig. 9. Restoration performance of the three considered deconvolution algorithms on a dataset of camera raw data. The markers represent the measured values of the ideal RMSE: the red squares refer to the uniform-motion PSF, while the green circles refer to camera-shake PSFs. The curves are obtained by third-order nonparametric polynomial regression of their respective values. As one can clearly notice, the behavior in case of uniform-motion blur is coherent with the results in [17] and significantly different from the one of camera shake.

image of the dataset, respectively; the symbol Π stands for the concatenation of the two vectors. From the dataset of raw images used in Section V, we computed $\hat{\gamma} = 27.20$; such value is required to correctly substitute (24) in the restoration-error model.

2) *Uniform-Motion Blur*: The PSF standard deviations in case of uniform-motion blur evolve with T as

$$\begin{aligned}\tilde{\zeta}_L(T) &= \gamma \cdot a_L T, \\ \tilde{\zeta}_S(T) &= 0,\end{aligned}\quad (26)$$

where $\gamma > 0$ is the motion speed, which determines the length γT of the PSF at time T and $a_L = \frac{1}{\sqrt{12}}$ is the standard deviation of the uniform distribution in $[0, 1]$. These expressions can be easily derived from the uniform-motion equations in continuous domain. Analogous to (25), an estimate of γ can be computed from the acquired raw images as

$$\hat{\gamma} = \sqrt{\text{med} \left\{ \left[\left(12 \left(\zeta_L^{(j)} \right)^2 - 2 \right) \left(\frac{1}{T^{(j)}} \right)^2 \right]_j \right\}},$$

obtaining $\hat{\gamma} = 17.81$.

D. Validation

Fig. 13 illustrates the surfaces of the three considered restoration-error models corresponding to the green channel

(yielding $\lambda M_G = 930$) in the 3D space having coordinates $(\varsigma_L, \varsigma_S, R_{930,0}(\varsigma_L, \varsigma_S, T))$. The colored dots represent the points $(\varsigma_L, \varsigma_S, R_{930,0}(\varsigma_L(T), \varsigma_S(T), T))$, which correspond to the curves given by the right-hand side of (22). Note that the exposure times associated to the surfaces have changed w.r.t. to those in Fig. 6, since here $\lambda M = 930$.

The restoration errors predicted by the computed models are better shown in the plots of Fig. 14, where are reported, for each channel of the Bayer pattern and for each considered restoration-error model, the outcomes for both the camera-shake and uniform-motion blur.

By comparing the trends of the corresponding plots in Fig. 9 and 14 it emerges a substantial similarity – in terms of qualitative behavior – between the outcomes of the restoration-error models and the RMSEs measured on the dataset of camera raw images after restoration with the corresponding algorithm. Furthermore, we can observe that the predicted RMSE values shown in Fig. 14 are indeed very close to those measured from the dataset of raw images.

Most importantly, the optimal exposure times identified by our model are consistent with those that emerge from the tests on camera raw images. Therefore, the hints provided by the restoration-error model are substantial, and in practice these allow to pilot the image acquisition in order to maximize the performance of the deblurring algorithm employed.




Fig. 10. An example of raw data acquired from *Balloons* (blue channel). The top row contains the estimated PSFs: the blur develops along trajectory f). The second row contains the camera raw data, acquired with settings (from left to right): 3, 4, 5, 6, 8, 10. The remaining rows show the restored images using different deconvolution algorithms. Third row: Anisotropic LPA-ICI Deconvolution, RMSE values (from left to right): 13.10, 7.23, 5.72, 5.76, 6.57, 6.89. Fourth row: Deconvolution using Natural Image Priors, RMSE values (from left to right): 12.33, 7.77, 6.30, 6.16, 6.18, 6.78. Bottom row: Richardson-Lucy Deconvolution, RMSE values (from left to right): 33.68, 22.41, 15.14, 11.04, 9.51, 9.10. Image intensities have been rescaled for visualization purposes.

In this sense, the derived models show satisfactory generalization properties on raw data, proving that the desiderata of Section III-E are satisfied by the choices that lead to computing these models, and in particular by the PSF descriptors and the smoothing operator used.

VII. DISCUSSION

A. Camera Shake vs. Rectilinear Blur

From the plots in Fig. 9 we can conclude that the camera-shake and the uniform-motion blurs require different acquisition strategies to maximize the quality of the restored image.

In case of camera shake the exposure time can be significantly increased, eventually improving the quality of the restored image, as the restoration error practically stabilizes after reaching its point of minimum. This also proves that the practice of minimizing the noise by using long exposure times (as in [2], [4], where a long-exposure image is deconvolved and the PSF is estimated exploiting an additional short-exposure image) is indeed an effective strategy to obtain good restoration quality out of images acquired with a hand-held camera in low-light conditions.

In case of uniform-motion blur, contrary to camera shake, the restoration errors show a clear minimum, which corre-

sponds to the optimal exposure time: beyond such optimal value, the error definitively increases. This result is coherent with the experimental validation in [17], and with analytical results concerning this qualitative behavior which indeed characterizes any deblurring algorithm \mathcal{D} when the observations are corrupted by uniform-motion blur. A very recent preprint by Tendero et al. [32] provides further mathematical proofs of this phenomenon.

B. Space-variant blur

The modeling used in our work relies on a space-invariant (convolutional) blur (2), (3) assumption, which implies that the PSF is the same at every location in the image. This simplification was employed rather successfully in many works concerning camera-shake removal (see e.g., [1], [4], [5]), and it is probably justified by the fact that, in most practical cases, the motion can be treated in first approximation as purely translational, therefore leading to space-invariant PSFs for sufficiently distant targets [31]. Even though methods and models for space-variant blur have lately gained popularity (see, e.g., [13], [26], [33]–[36]), we resorted to the space-invariant blur modeling for the sake of simplicity.



Fig. 11. An example of raw data acquired from *Liza* (green channel). The top row contains the estimated PSFs: the blur develops along trajectory *e*). The second row contains the camera raw data, acquired with settings (from left to right): 2, 3, 5, 6, 8, 10. The remaining rows show the restored images using different deconvolution algorithms. Third row: Anisotropic LPA-ICI Deconvolution, RMSE values (from left to right): 10.01, 7.56, 5.19, 5.61, 6.09, 8.99. Fourth row: Deconvolution using Natural Image Priors, RMSE values (from left to right): 9.79, 7.57, 5.18, 5.16, 5.54, 8.16. Bottom row: Richardson-Lucy Deconvolution, RMSE values (from left to right): 28.43, 18.25, 8.47, 7.33, 7.12, 9.43. Image intensities have been rescaled for visualization purposes.

On the one hand, based on the examples of space-variant PSFs from camera-shake blur shown in very recent works [13], [34], [35], [37], we argue that even though blur PSFs can be different at different locations in the image, typically they all share essentially the same values of the standard deviations ς_L , ς_R . Thus, our analysis may be applied also to such more general case, provided that the employed deconvolution algorithm \mathcal{D} for space-invariant PSF is replaced by one designed for the removal of space-variant blur.

On the other hand, scenes that include targets at close range on a distant background or targets with different relative motions, cannot be modeled by the proposed methodology, at least without introducing some form of image segmentation.

C. Practical Applicability

Let us discuss how to face few practical issues that need to be considered for using a suitably trained restoration-error model to determine the optimal exposure time before acquiring a picture.

Firstly, the restoration-error model assumes that the parameters λ , σ and M are known. The noise parameters λ and σ depend on the sensor hardware and (to a lesser extent) on the temperature. In practice, these parameters can be profiled

beforehand and stored in a look-up table, or they can be automatically estimated from a single auxiliary image (e.g., following the approach presented in [38]). The image range M can be easily estimated from an auxiliary exposure or by using a light meter.

Then, it is necessary characterize the specific motion blur affecting the acquisition, obtaining equations $\varsigma_L(T)$ and $\varsigma_S(T)$, like, e.g., (24) and (26). In the specific case of camera-shake blur, customized parameters a_L and a_S for (24) can be estimated for each camera user, studying the way the camera typically shakes, following the approach of [18] and possibly leveraging accelerometers. It is therefore worth stressing that, in the specific case of camera shake, the parameter γ depends on the camera optics and on the scene distance, while, in the case of uniform-motion blur, γ is determined by the apparent speed of the motion: in both cases it needs to be estimated before the acquisition. However, this is not a real issue for modern digital cameras that may estimate these using the target tracking or autofocus functionalities. Other exotic sorts of motion blur may be handled following the general approach presented in [18], or, in case the motion source is well known, through an analytical formulation.




Fig. 12. An example of raw data acquired from *Salamander* (red channel). The top row contains the estimated PSFs: the blur develops along trajectory a). The second row contains the camera raw data, acquired with settings (from left to right): 2, 3, 5, 6, 8, 10. The remaining rows show the restored images using different deconvolution algorithms. Third row: Anisotropic LPA-ICI Deconvolution, RMSE values (from left to right): 18.20, 13.44, 8.41, 9.86, 16.33, 22.89. Fourth row: Deconvolution using Natural Image Priors, RMSE values (from left to right): 18.10, 13.07, 7.90, 8.89, 14.67, 20.31. Bottom row: Richardson-Lucy Deconvolution, RMSE values (from left to right): 34.77, 25.53, 12.43, 12.11, 17.23, 22.93. Image intensities have been rescaled for visualization purposes.

VIII. CONCLUSIONS

We have detailed a methodology for deriving a statistical model of the performance of a given deblurring algorithm, when used to restore motion blurred images. Differently from our earlier work on rectilinear blur [17], we do not enforce any analytical formulation for the trajectories generating the motion-blur PSFs and we deal with random motion, which is effectively handled by means of statistical descriptors of the PSF.

Thanks to extensive experiments on camera raw images we investigated the blur/noise trade-off that rules the restoration performance in presence of motion blur, and we show that the computed restoration-error models provide estimates that are coherent with the results on real data.

In practice these models, combined with functions expressing how the PSF descriptors vary w.r.t. the exposure times, provide concrete guidelines for predicting the exposure time that maximizes the quality of the image restored by the corresponding algorithm. The outcomes of the restoration error models obtained from three different deconvolution algorithms (namely the Anisotropic LPA-ICI Deconvolution, the Deconvolution using Sparse Natural Image Priors, and the Richardson-Lucy Deconvolution), agree with the results shown

in [17], with the acquisition strategies followed in the practice to cope with camera shake, and with an extensive experimental evaluation performed on camera raw images.

ACKNOWLEDGMENT

The authors thank Tampere Center for Scientific Computing and CSC - Finnish IT Center for Science Ltd for providing access to the computing infrastructure that made this research possible.

This work was supported by the Academy of Finland (project no. 213462, Finnish Programme for Centres of Excellence in Research 2006-2011, project no. 129118, Postdoctoral Researcher's Project 2009-2011, and project no. 252547, Academy Research Fellow 2011-2016), by CIMO, the Finnish Centre for International Mobility (fellowship TM-07-4952).

REFERENCES

- [1] M. Tico and M. Vehvilainen, "Bayesian estimation of motion blur point spread function from differently exposed image frames," in *Proc. EUSIPCO 2006*, 2006.




Fig. 13. A different view of the restoration-error model for the three considered deconvolution algorithms. These surfaces have been rescaled according to (16) to model the green channel of the dataset of camera raw images, i.e. $\lambda M_G = 930$. The colored dots represent the points $(\varsigma_L, \varsigma_S, R_{930,0}(\varsigma_L(T), \varsigma_S(T), T))$, which correspond to the curves given by the right-hand side of (22) in case of uniform motion (red dots) or camera shake (green dots).

- [2] —, “Image stabilization based on fusing the visual information in differently exposed images,” in *Proceedings of ICIP International Conference of Image Processing, San Antonio, USA, 2007*.
- [3] S. Nayar and M. Ben-Ezra, “Motion-based motion deblurring,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 26, no. 6, pp. 689–698, 2004.
- [4] L. Yuan, J. Sun, L. Quan, and H.-Y. Shum, “Image deblurring with blurred/noisy image pairs,” *ACM Trans. Graph.*, vol. 26, no. 3, p. 1, 2007.
- [5] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman, “Removing camera shake from a single photograph,” *ACM Trans. Graph.*, vol. 25, no. 3, pp. 787–794, 2006.
- [6] Y. Yitzhaky, I. Mor, A. Lantzman, and N. S. Kopeika, “A direct method for restoration of motion blurred images,” *J. Opt. Soc. Am. A*, p. 1998.
- [7] A. Levin, R. Fergus, F. Durand, and W. T. Freeman, “Image and depth from a conventional camera with a coded aperture,” *ACM Trans. Graph.*, vol. 26, July 2007. [Online]. Available: <http://doi.acm.org/10.1145/1276377.1276464>
- [8] W. H. Richardson, “Bayesian-based iterative method of image restoration,” *J. Opt. Soc. Am.*, vol. 62, no. 1, pp. 55–59, 1972. [Online]. Available: <http://www.opticsinfobase.org/abstract.cfm?URI=josa-62-1-55>
- [9] R. Neelamani, H. Choi, and R. Baraniuk, “Forward: Fourier-wavelet regularized deconvolution for ill-conditioned systems,” *Image Processing, IEEE Transactions on*, vol. 52, pp. 418–433, 2004.
- [10] A. Foi, S. Alenius, M. Trimeche, V. Katkovnik, and K. Egiazarian, “A spatially adaptive Poissonian image deblurring,” in *Proc. IEEE Int. Conf. Image Process.*, Genova, Italy, September 2005, pp. 925–928.
- [11] V. Katkovnik, K. Egiazarian, and J. Astola, “A spatially adaptive nonparametric regression image deblurring,” *Image Processing, IEEE Transactions on*, vol. 14, no. 10, pp. 1469–1478, October 2005.
- [12] S. Reeves and R. Mersereau, “Blur identification by the method of generalized cross-validation,” *Image Processing, IEEE Transactions on*, vol. 1, no. 3, pp. 301–311, Jul. 1992.
- [13] N. Joshi, S. B. Kang, C. L. Zitnick, and R. Szeliski, “Image deblurring using inertial measurement sensors,” *ACM Trans. Graph.*, vol. 29, pp. 30:1–30:9, July 2010. [Online]. Available: <http://doi.acm.org/10.1145/1778765.1778767>
- [14] T. Chan and C. Wong, “Total variation blind deconvolution,” *Image Processing, IEEE Transactions on*, vol. 7, no. 3, pp. 370–375, March 1998.
- [15] S. Cho and S. Lee, “Fast motion deblurring,” *ACM Transactions on Graphics (SIGGRAPH ASIA 2009)*, vol. 28, no. 5, p. article no. 145, 2009.
- [16] L. Xu and J. Jia, “Two-phase kernel estimation for robust motion deblurring,” in *Proceedings of the 11th ECCV: Part I*. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 157–170.
- [17] G. Boracchi and A. Foi, “Uniform motion blur in poissonian noise: Blur/noise tradeoff,” *Image Processing, IEEE Transactions on*, vol. 20, no. 2, pp. 592–598, 2011.
- [18] A. S. Feng Xiao and J. Farrell, “Camera-motion and effective spatial resolution,” in *International Congress. of Imaging Science, Rochester, NY*, May 2006.
- [19] A. Agrawal, Y. Xu, and R. Raskar, “Invertible motion blur in video,” *ACM Trans. Graph.*, vol. 28, no. 3, pp. 1–8, 2009.
- [20] D. G. Shaojie Zhuo and T. Sim, “Robust flash deblurring,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2010.
- [21] R. Raskar, A. Agrawal, and J. Tumblin, “Coded exposure photography: motion deblurring using fluttered shutter,” in *SIGGRAPH ’06: ACM SIGGRAPH 2006 Papers*. New York, NY, USA: ACM, 2006, pp. 795–804.
- [22] A. Agrawal and R. Raskar, “Resolving objects at higher resolution from a single motion-blurred image,” *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, vol. 0, pp. 1–8, 2007.
- [23] J. Telleen, A. Sullivan, J. Yee, O. Wang, P. Gunawardane, I. Collins, and J. Davis, “Synthetic shutter speed imaging,” *Comput. Graph. Forum*, vol. 26, no. 3, pp. 591–598, 2007.
- [24] A. Levin, P. Sand, T. S. Cho, F. Durand, and W. T. Freeman, “Motion-invariant photography,” in *SIGGRAPH ’08: ACM SIGGRAPH 2008 papers*. New York, NY, USA: ACM, 2008, pp. 1–9.
- [25] T. S. Cho, A. Levin, F. Durand, and W. T. Freeman, “Motion blur

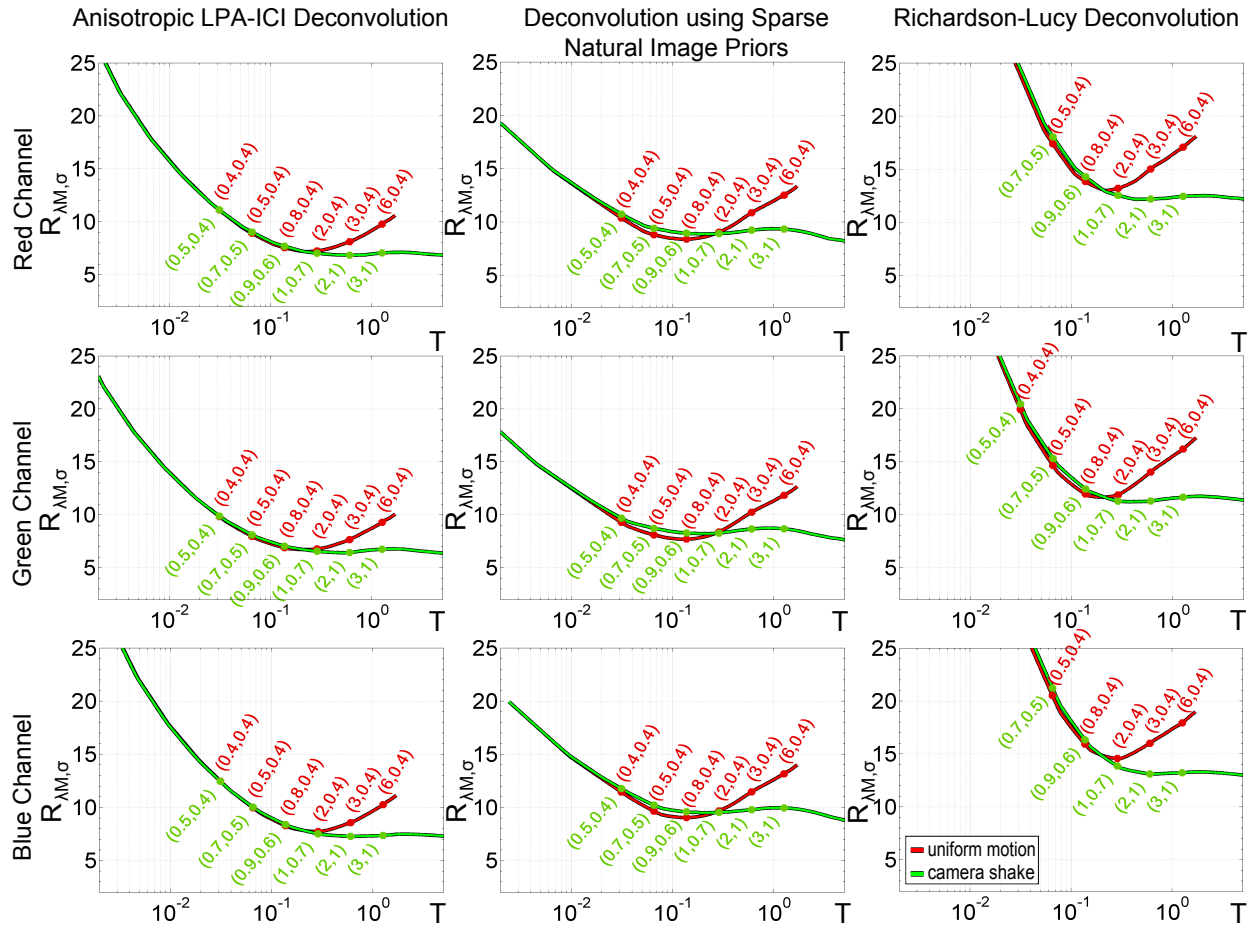


Fig. 14. The outcomes of the computed restoration-error models for the three different deconvolution algorithms in case of camera-shake and uniform-motion blur. The curves are presented as in Fig. 9, by separately considering each channel of the Bayer pattern since the original images have significantly smaller ranges. The points and the numbers shown in each plots refer to the values of the PSF standard deviation at a specific exposure time, which have been reported for illustrative purposes.

- removal with orthogonal parabolic exposures,” in *IEEE International Conference in Computational Photography (ICCP)*, 2010.
- [26] Y.-W. Tai, H. Du, M. S. Brown, and S. Lin, “Correction of spatially varying image and video motion blur using a hybrid camera,” *PAMI, IEEE Transactions on*, vol. 32, no. 6, pp. 1012–1028, 2010.
- [27] A. Foi, M. Trimeche, V. Katkovnik, and K. Egiazarian, “Practical Poissonian-Gaussian noise modeling and fitting for single-image raw-data,” *IEEE Trans. Image Process.*, vol. 17, no. 10, pp. 1737–1754, Oct. 2008.
- [28] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: From error visibility to structural similarity,” *Image Processing, IEEE Transactions on*, vol. 13, no. 4, pp. 600–612, 2004.
- [29] V. Katkovnik, K. Egiazarian, and J. Astola, *Local Approximation Techniques in Signal and Image Process.* SPIE Press, 2006, vol. PM157.
- [30] M. Bertero and P. Boccacci, *Introduction to Inverse Problems in Imaging.* Institute of Physics Publishing, 1998.
- [31] E. Or and D. Pundik, “Hand motion and image stabilization in hand-held devices,” *Consumer Electronics, IEEE Transactions on*, vol. 53, no. 4, pp. 1508–1512, Nov. 2007.
- [32] Y. Tendero, J. Morel, and B. Rougé, “The flutter shutter paradox,” http://dev.ipol.im/~morel/Dossier_MVA_2011_Cours_Transparents_Documents/2011_Cours6_Cours_Flutter_Shutter.pdf.
- [33] M. Šorel, F. Šroubek, and J. Flusser, “Recent advances in space-variant deblurring and image stabilization,” in *Unexploded Ordnance Detection and Mitigation*, ser. NATO Science for Peace and Security Series B: Physics and Biophysics, J. Byrnes, Ed. Springer Netherlands, 2009, pp. 259–272.
- [34] O. Whyte, J. Sivic, A. Zisserman, and J. Ponce, “Non-uniform deblurring for shaken images,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2010.
- [35] S. Harmeling, M. Hirsch, and B. Schölkopf, “Space-variant single-image blind deconvolution for removing camera shake,” in *NIPS 2010*, Max-Planck-Gesellschaft. Red Hook, NY, USA: Curran, 1 2011, pp. 1–9.
- [36] Y.-W. Tai, P. Tan, and M. Brown, “Richardson-lucy deblurring for scenes under a projective motion path,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 33, no. 8, pp. 1603–1618, aug. 2011.
- [37] A. Levin, Y. Weiss, F. Durand, and W. Freeman, “Understanding and evaluating blind deconvolution algorithms,” *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, vol. 0, pp. 1964–1971, 2009.
- [38] A. Foi, S. Alenius, V. Katkovnik, and K. Egiazarian, “Noise measure-

ment for raw-data of digital imaging sensors by automatic segmentation of nonuniform targets," *IEEE Sensors Journal*, vol. 7, no. 10, pp. 1456–1461, Oct. 2007.