Image restoration using the particle filter: handling non-causality

J.K. Prasanna and A.N. Rajagopalan

Abstract: A recursive-state estimation scheme for image restoration using the particle filter is described. Handling non-causal blurs within a recursive framework is a challenging problem. Most recursive image restoration schemes assume the blur to be causal or semi-causal in nature, but this is unrealistic. A novel choice for the state vector and a concurrent block estimation technique to incorporate full-plane regions of support for the image model as well as the blur are proposed. The particle filter-based framework enables general types of degradations to be tackled. The method assumes that the functional form of the distortion as well as the noise statistics are known but does not place any restrictions on them. Several experimental results are presented to validate the approach.

1 Introduction

The goal in image restoration is to recover an original image from its degraded version. This involves modelling the degradation and then applying an inverse transformation on the observed image. However, the presence of noise in the degraded image makes this problem ill-posed. Hence, direct functional inversion and pseudo-inversion methods are not very successful [1]. A classical approach to circumvent this problem is the Wiener filter, which accounts for the presence of noise [1]. It is basically a linear minimum mean square error (MMSE) estimator. Although the Wiener filter is quite popular, its scope is limited to space-invariant degradations, stationary images and noise. Iterative approaches such as projection onto convex sets [2], expectation maximisation [3] and so on, constitute yet another class of solutions. However, these techniques suffer from heavy computational load and storage problems. Recursive 2D estimation techniques also exist and have received a great deal of attention from researchers [1], as they are computationally less complex compared with iterative methods. They permit spatial adaptivity of image models and can be used to restore images degraded even by space-variant blurs. One of the drawbacks of recursive filtering is that incorporating non-causality is difficult. The popular 2D Kalman filter (KF) [4, 5] and most other recursive schemes assume the region of support (ROS) for the image to be either causal or semi-causal. Citrin and Azimi-Sadjadi [6] have proposed a scheme that uses full-plane support for the image formation model. Although a full-plane model provides better prediction than a causal one, the latter is still widely used because of its ease of implementation.

IEE Proceedings online no. 20045092

doi:10.1049/ip-vis:20045092

Paper first received 3rd July 2004 and in revised form 24th September 2005

The authors are with the Image Processing and Computer Vision Lab., Department of Electrical Engineering, Indian Institute of Technology Madras, Chennai 600036, India

E-mail: raju@ee.iitm.ernet.in

The use of causal blur kernels to model degradations in real images is even less justifiable. The assumption that a certain pixel is affected only by the pixels in its 'past' and not by the ones in its 'future' is more convenient than realistic. There have been some attempts in the past to come up with recursive schemes that can handle non-causal blurs. One simple solution [2] is to shift the support of the kernel until it is first-quadrant causal, that is, to introduce a fixed delay. However, this can result in instability of the filter if the shifted kernel is not a minimum phase bi-sequence. Tekalp *et al*. [7] have proposed decomposing the blur kernel into four quarter-plane convolutional factors, using spectral factorisation. Each of these convolutional factors is recursively realisable in their respective directions of recursion. However, this scheme is limited to space-invariant blurs. Moreover, it makes unrealistic assumptions about the nature of the blur kernel. Handling non-causal blur within a recursive framework continues to be a challenging problem. We pose the restoration problem as a state estimation problem and propose a new recursive approach for handling non-causality. Our work is motivated by the work of Citrin and Azimi-Sadjadi [6]. However, unlike in their work [6], we can handle a full-plane ROS for both the image and the blur models. Causality is maintained within the filtering process through a novel choice of the state vector and by concurrent block estimation.

The recursive and analytically tractable 2D KF [4, 5] gives optimum restoration results only when the deterministic degradation is linear and the observation noise is additive white-Gaussian. However, image distortion can, in general, be nonlinear, and the observation noise non-additive and/or non-Gaussian. In fact, most scanners and photographic films have nonlinear input–output characteristics [8]. To handle such general types of degradations, we incorporate state estimation within a recursive particle-filtering framework. The particle-filter (PF) can handle nonlinearity/non-Gaussianity and is a Monte Carlo implementation of a recursive-state estimator. The proposed method assumes that the functional form of the distortion and the noise statistics are known exactly but do not place any restrictions on them. We give several results to validate our approach to the image restoration problem.

2 Handling non-causality

A recursive framework requires a proper state vector to be conceptualised. It should take into account the ROS for the image formation model as well as the degradation model that is the size and shape of the blur kernel. It should not constrain the order of support of either of the two models to be space-invariant. The maximum among all orders of support should be taken into account while constructing the state vector. At any instant of time, all the pixels that are required to estimate the current pixel must be included in the state vector.

Consider an $M \times N$ image that is scanned from left to right and top to bottom in block rows of width N_1 and block columns of width N_2 (Fig. 1). Each block consists of $\mathcal{N} = N_1 N_2$ pixels. The state at any instant consists of the 13 blocks shown in Fig. 1. The size of each block (i.e., N_1 and N_2) is decided by the maximum among the sizes of the orders of support of both the models. The process of scanning is illustrated in Fig. 2. The state moves to the right by one block at every time step. The processing strip at any instant consists of three block rows. The goal is to estimate the pixels in block 6.

We assume concurrent linear auto-regressive (AR) models for the image. Specifically, we assume the state space model equation to be

$$\boldsymbol{X}(k) = \boldsymbol{A}\boldsymbol{X}(k-1) + \boldsymbol{U}(k) \tag{1}$$

that is

$$\begin{pmatrix} X_1(k) \\ X_2(k) \\ \vdots \\ X_{13}(k) \end{pmatrix} = \boldsymbol{A} \begin{pmatrix} X_1(k-1) \\ X_2(k-1) \\ \vdots \\ X_{13}(k-1) \end{pmatrix} + \begin{pmatrix} U_1(k) \\ U_2(k) \\ \vdots \\ U_{13}(k) \end{pmatrix}$$

where $\boldsymbol{X}(k)$ is the current state vector consisting of 13 blocks, $\boldsymbol{X}(k-1)$ is the state vector in the previous time step and $\boldsymbol{U}(k)$ is the noise vector driving the process. The elements of matrix $\boldsymbol{A}$ depend on the AR coefficients and the ROS used to estimate each block. When the state propagates along the strip by one step, the blocks 2, 3, 4, 6, 7, 8, 9, 11, 12 and 13 are predicted using ten concurrent estimators. The supports for these blocks are given in Table 1.

The propagation of the state along the image boundary will be discussed subsequently. It should be noted that block 6 is predicted using a full-plane ROS. All other blocks are estimated only to provide support for estimating block 6. Blocks 1, 5 and 10 are propagated according to the equations $X_1(k) = X_2(k-1)$, $X_5(k) = X_6(k-1)$ and $X_{10}(k) = X_{11}(k-1)$. That is, these are obtained directly as estimates from the previous step. This is done because we will not be able to increase their respective regions of support even if we use estimators for these blocks. For example, if we decide to use an AR predictor for block 1

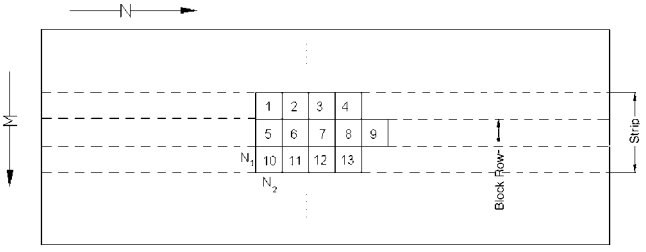

Fig. 1 *State that consists of 13 blocks at any time instant*

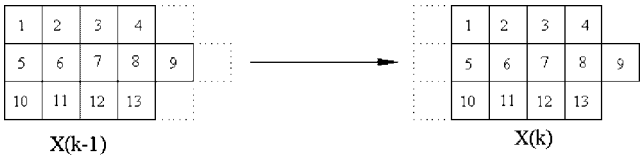

Fig. 2 *State movement to the right by one block at every time step*

as well, we would have to use the five blocks 1, 3, 5, 6 and 7 as the supporting blocks. But as block 2 was also estimated with a similar ROS, not much is gained by using AR predictors for block 1. Matrix $\boldsymbol{A}$ therefore becomes

$$\boldsymbol{A} = \begin{pmatrix} 0 & I & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & A_{2,2} & 0 & A_{2,4} & 0 & A_{2,6} & A_{2,7} & A_{2,8} \\ 0 & 0 & A_{3,3} & 0 & 0 & 0 & A_{3,7} & A_{3,8} \\ 0 & 0 & 0 & A_{4,4} & 0 & 0 & 0 & A_{4,8} \\ 0 & 0 & 0 & 0 & 0 & I & 0 & 0 \\ 0 & A_{6,2} & A_{6,3} & A_{6,4} & 0 & A_{6,6} & 0 & A_{6,8} \\ 0 & 0 & A_{7,3} & A_{7,4} & 0 & 0 & A_{7,7} & 0 \\ 0 & 0 & 0 & A_{8,4} & 0 & 0 & 0 & A_{8,8} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & A_{11,6} & A_{11,7} & A_{11,8} \\ 0 & 0 & 0 & 0 & 0 & 0 & A_{12,7} & A_{12,8} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & A_{13,8} \end{pmatrix}$$

$$\begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ A_{3,9} & 0 & 0 & 0 & 0 \\ A_{4,9} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ A_{6,9} & 0 & A_{6,11} & A_{6,12} & A_{6,13} \\ A_{7,9} & 0 & 0 & A_{7,12} & A_{7,13} \\ 0 & 0 & 0 & 0 & A_{8,13} \\ A_{9,9} & 0 & 0 & 0 & 0 \\ 0 & 0 & I & 0 & 0 \\ 0 & 0 & A_{11,11} & 0 & A_{11,13} \\ A_{12,9} & 0 & 0 & A_{12,12} & 0 \\ A_{13,9} & 0 & 0 & 0 & A_{13,13} \end{pmatrix} \tag{2}$$

Table 1: ROS for the ten blocks

Block	Support
$X_9(k)$	$X_9(k-1)$
$X_4(k)$	$X_i(k-1)$, $i = 4, 8, 9$
$X_8(k)$	$X_i(k-1)$, $i = 4, 8, 13$
$X_{13}(k)$	$X_i(k-1)$, $i = 8, 9, 13$
$X_3(k)$	$X_i(k-1)$, $i = 3, 7, 8, 9$
$X_7(k)$	$X_i(k-1)$, $i = 3, 4, 7, 9, 12, 13$
$X_{12}(k)$	$X_i(k-1)$, $i = 7, 8, 9, 12$
$X_2(k)$	$X_i(k-1)$, $i = 2, 4, 6, 7, 8$
$X_6(k)$	$X_i(k-1)$, $i = 2, 3, 4, 6, 8, 11, 12, 13$
$X_{11}(k)$	$X_i(k-1)$, $i = 6, 7, 8, 11, 13$

Each sub-matrix of $\boldsymbol{A}$ is of size $\mathcal{N} \times \mathcal{N}$ where $\mathcal{N}$ is the number of pixels in a block. The model parameters, that is, the $A_{i,j}$'s and the prediction variances can be calculated using the Yule–Walker equations [9].

3 Propagation of the state density

For the linear-Gaussian model, an optimal estimate of the state can be obtained using the KF, which propagates and updates the mean and the covariance. However, when there is nonlinearity and/or non-Gaussianity involved either in the state model or in the observation model, the state density can become skewed or even multi-modal. In such cases, the probability density function (pdf) of the state will not be Gaussian and the KF is no longer optimal.

Our aim here is to construct the pdf of the state vector at each time step so that an optimal estimate of the state can be obtained from it. The state vector at time index k, $\boldsymbol{X}(k) \in \mathbb{R}^{13}$, is assumed to evolve according to the system model in (1). Note that the state dynamics forms a Markov chain in time. That is, if $\boldsymbol{\xi}(k) = \{\boldsymbol{X}(1), \boldsymbol{X}(2), \ldots, \boldsymbol{X}(k)\}$ is the set of all state vectors till time step k, then

$$p(\boldsymbol{X}(k+1)|\boldsymbol{\xi}(k)) = p(\boldsymbol{X}(k+1)|\boldsymbol{X}(k)) \tag{3}$$

At every time step, a noisy observation $\boldsymbol{Y}(k)$ becomes available from the degraded image. This observation is related to the state vector $\boldsymbol{X}(k)$ via the measurement model

$$\boldsymbol{Y}(k) = h_k(\boldsymbol{X}(k), \boldsymbol{V}(k)) \tag{4}$$

Here, $\boldsymbol{V}(k)$ is the observation noise, which is independent of $\boldsymbol{X}(r)$ for $r \leq k$ and $\boldsymbol{U}(r)$ for all r, whereas h_k is the degradation function.

We assume that the pdfs of the noise processes $\boldsymbol{U}(k)$ and $\boldsymbol{V}(k)$ are known, in addition to matrix $\boldsymbol{A}$ and the functional form of h_i for $i = 1, 2, \ldots, k$. Let $\boldsymbol{\Gamma}(k) = \{\boldsymbol{Y}(1), \boldsymbol{Y}(2), \ldots, \boldsymbol{Y}(k)\}$ be the set of all available observations at time step k. We also assume that the pdf of the initial state, that is, $p(\boldsymbol{X}(1)) = p(\boldsymbol{X}(1)|\boldsymbol{\Gamma}(0))$, is known.

The required pdf $p(\boldsymbol{X}(k)|\boldsymbol{\Gamma}(k))$ can be obtained recursively in two stages: prediction and update. Suppose that the pdf $p(\boldsymbol{X}(k-1)|\boldsymbol{\Gamma}(k-1))$ is available. Using the system model (1), we obtain the prior $p(\boldsymbol{X}(k)|\boldsymbol{\Gamma}(k-1))$ as

$$p(\boldsymbol{X}(k)|\boldsymbol{\Gamma}(k-1)) = \int p(\boldsymbol{X}(k)|\boldsymbol{X}(k-1)) \times p(\boldsymbol{X}(k-1)|\boldsymbol{\Gamma}(k-1))\, \mathrm{d}\boldsymbol{X}(k-1) \tag{5}$$

Here, $p(\boldsymbol{X}(k)|\boldsymbol{X}(k-1))$ is defined by the system model and the known pdf of the system model noise $\boldsymbol{U}(k-1)$. When the observation $\boldsymbol{Y}(k)$ becomes available, we can use it to update the prior density, using Bayes theorem to obtain the required posterior density

$$p(\boldsymbol{X}(k)|\boldsymbol{\Gamma}(k)) = \frac{p(\boldsymbol{Y}(k)|\boldsymbol{X}(k))p(\boldsymbol{X}(k)|\boldsymbol{\Gamma}(k-1))}{\int p(\boldsymbol{Y}(k)|\boldsymbol{X}(k))p(\boldsymbol{X}(k)|\boldsymbol{\Gamma}(k-1))\, \mathrm{d}\boldsymbol{X}(k)} \tag{6}$$

After the update, the density tends to peak in the vicinity of the observation. Note that the conditional pdf $p(\boldsymbol{Y}(k)|\boldsymbol{X}(k))$ is defined by the measurement model and the statistics of the observation noise $\boldsymbol{V}(k)$.

4 PF for restoration

Only for the simple linear-Gaussian estimation problem, does the density in (6) remain Gaussian at every time step. Our goal here is to solve the general image restoration problem. PFs have been successfully used to handle elements of nonlinearity/non-Gaussianity in 1D problems [10]. Studies have shown that filters that use analytical approximations such as the extended KF (EKF) [11] and the Gaussian sum filter [12] can severely distort the underlying structure of the state, which can in turn result in filter divergence [10]. The key idea in PF theory is to represent the required posterior state density as a set of random samples with associated weights. Estimates of moments can then be obtained directly from the samples. As the number of samples becomes very large, the PF estimate approaches the true value. We now propose an extension of the 1D PF to tackle images.

Assume that we have a set of samples $\{\boldsymbol{X}^{(i)}(k-1), i = 1, 2, \ldots, K_0\}$ that are distributed as $p(\boldsymbol{X}(k-1)|\boldsymbol{\Gamma}(k-1))$. To obtain a set of samples distributed as $p(\boldsymbol{X}(k)|\boldsymbol{\Gamma}(k))$, we do the following:

Step (i): Prediction. Each sample is passed through the system model (i.e. the image formation model) in (1) to obtain a new set of samples, which approximate the prior density at time step k. That is, for some $1 \leq i \leq K_0$, let

$$\tilde{\boldsymbol{X}}^{(i)}(k) = A\boldsymbol{X}^{(i)}(k-1) + \boldsymbol{U}^{(i)}(k-1) \tag{7}$$

If $\boldsymbol{X}^{(i)}(k-1)$ is a sample drawn from $p(\boldsymbol{X}(k-1)|\boldsymbol{\Gamma}(k-1))$ and $\boldsymbol{U}^{(i)}(k-1)$ is a sample drawn from the known noise pdf $p(\boldsymbol{U}(k-1))$, it is clear that $\tilde{\boldsymbol{X}}^{(i)}(k)$ in (7) is distributed as $p(\boldsymbol{X}(k)|\boldsymbol{\Gamma}(k-1))$. Repeating this procedure for all the samples, we obtain a new set of K_0 values that are independently distributed as $p(\boldsymbol{X}(k)|\boldsymbol{\Gamma}(k-1))$.

Step (ii): Update. On the receipt of the observation $\boldsymbol{Y}(k)$, the weight corresponding to each prior sample is evaluated and the prior values are re-sampled to obtain a new set. These samples are distributed as the required posterior density $p(\boldsymbol{X}(k)|\boldsymbol{\Gamma}(k))$.

Specifically, the normalised weight $q^{(i)}$ of the ith prior sample is computed as

$$q^{(i)} = \frac{p(\boldsymbol{Y}(k)|\tilde{\boldsymbol{X}}^{(i)}(k))}{\sum_{j=1}^{K_0} p(\boldsymbol{Y}(k)|\tilde{\boldsymbol{X}}^{(j)}(k))} \tag{8}$$

We then sample K_0 times from the set $\{\tilde{\boldsymbol{X}}^{(i)}(k), i = 1, 2, \ldots, K_0\}$, with probability mass $q^{(i)}$ associated with the value $\tilde{\boldsymbol{X}}^{(i)}(k)$, to generate a new set of samples $\{\boldsymbol{X}^{(i)}(k), i = 1, 2, \ldots, K_0\}$. An important theorem, proposed by Smith and Gelfand [13], states that samples drawn from the set $\{\tilde{\boldsymbol{X}}^{(i)}(k), i = 1, 2, \ldots, K_0\}$, with probability mass $q^{(i)}$ in (8), are asymptotically distributed as the required posterior density. That is, the set of samples $\{\boldsymbol{X}^{(i)}(k), i = 1, 2, \ldots, K_0\}$ represents $p(\boldsymbol{X}(k)|\boldsymbol{\Gamma}(k))$ as $K_0 \to \infty$.

Re-sampling can be accomplished as follows:

1. Define the cumulative weight $C^{(i)}$ of the ith particle as $C^{(i)} = C^{(i-1)} + q^{(i)}$, $i = 1, 2, \ldots, K_0$, $C^{(0)} = 0$.
2. Generate a uniform random number $\zeta \in [0, 1]$.
3. Find the smallest j for which $C^{(j)} \geq \zeta$.
4. Assign $\boldsymbol{X}^{(i)}(k) = \tilde{\boldsymbol{X}}^{(j)}(k)$.

The prediction and update steps form one iteration of the algorithm. The algorithm is initialised using samples drawn from the known pdf $p(\boldsymbol{X}(1)|\boldsymbol{\Gamma}(0))$. An optimal estimate of the state, such as the conditional mean, can be obtained from the updated samples. It is well known that the conditional mean is an MMSE estimate of the state and is

generally a nonlinear function of the observation. When the state is well within the image, only block 6 (which is predicted using a full-plane ROS) is saved as the final estimate of the original image at each time step. That is, if (m, n) is a pixel location well within the image field, then $\hat{x}(m, n) = (1/K_0)\sum_{j=1}^{K_0} X_6^{(j)}(k)$, where $\hat{x}(m, n)$ is the PF estimate of the pixel (m, n). At time step k, pixel 6 occupies the spatial position (m, n).

For $k = 0$, blocks 1, 5 and 10 lie outside the observed image. We initialise these three blocks by sampling from a Gaussian pdf, with mean and variance equal to the mean and variance of the degraded image itself. The other ten blocks lie inside the image and are initialised by sampling K_0 times from a Gaussian whose mean is equal to the observation at the corresponding location, and variance equal to the observation noise variance. The 13 K_0 samples obtained in this manner are fed into the image formation model (1). The state keeps moving to the right and only the full-plane estimate of block 6 is saved at each time step. When the state reaches the right side boundary of the image, non-full-plane estimates of blocks 7, 8 and 9 are saved at the corresponding spatial locations. For the first strip located at the top boundary of the image, block 2 (i.e. $X_2(k)$) is saved as the final estimate. Similarly, for the bottom boundary, $X_{11}(k)$ is saved as the final estimate.

4.1 Some examples

Without loss of generality and to fix ideas, let us assume from now on that each block consists of just one pixel, that is, $\mathcal{N} = 1$. This gives rise to a 13-dimensional state vector and a matrix $\boldsymbol{A}$ of size 13×13. This will allow us to handle space-variant blurs with blur kernels of 3×3 size. Similarly, $\mathcal{N} = 2 \times 2 = 4$ would allow us to handle a maximum blur kernel of 5×5 size.

Case A: Linear-Gaussian model: For the case in which the degradation is linear and the noise is additive white-Gaussian, the observation equation becomes

$$\boldsymbol{Y}(k) = H(k)\boldsymbol{X}(k) + \boldsymbol{V}(k) \tag{9}$$

In the case of a 3×3 non-causal blur, any given pixel is affected only by itself and the eight pixels surrounding it. Also, we note from the state vector diagram of Fig. 1 that all pixels surrounding blocks 6 and 7 are also a part of the state vector at any instant of time. Let $Y_6(k)$ and $Y_7(k)$ be the observed pixel values at the spatial locations corresponding to pixels 6 and 7, respectively, at some time index k. We can then write the observation equation as

$$\begin{pmatrix} Y_6(k) \\ Y_7(k) \end{pmatrix} = \boldsymbol{H}(k) \begin{pmatrix} X_1(k) \\ X_2(k) \\ \vdots \\ X_{13}(k) \end{pmatrix} + \begin{pmatrix} V_6(k) \\ V_7(k) \end{pmatrix} \tag{10}$$

where $\boldsymbol{H}(k)$ is a 2×13 blur matrix whose entries will change with k if the blur is space-variant.

Let $\{\tilde{\boldsymbol{X}}^{(i)}(k), i = 1, 2, \ldots, K_0\}$ be the set of vector samples that are distributed as the prior pdf. These samples are obtained by prediction using (1). We need to update these samples using appropriate weights. Upon receipt of the observation vector $\boldsymbol{Y}(k)$ at time step k, the normalised weight of each particle is calculated using (8). The expression for the normalised weight then becomes

$$q^{(i)} = \frac{\exp[-\|\boldsymbol{Y}(k) - \boldsymbol{H}(k)\tilde{\boldsymbol{X}}^{(i)}(k)\|^2/2\sigma_v^2]}{\sum_{j=1}^{K_0}\exp[-\|\boldsymbol{Y}(k) - \boldsymbol{H}(k)\tilde{\boldsymbol{X}}^{(j)}(k)\|^2/2\sigma_v^2]} \tag{11}$$

where σ_v^2 is the variance of the observation noise.

Case B: Nonlinear model: An example of a nonlinear image-recording medium is the photographic film wherein a point-wise nonlinear relationship exists between the scene and the silver density on the film and is of the form

$$r_{\mathrm{d}}(m, n) = \alpha \log_{10}\{h(s_{\mathrm{e}}(m, n))\} + \beta + v_{\mathrm{d}}(m, n) \tag{12}$$

The subscripts d and e denote the density and the exposure domain, respectively. The term $r_{\mathrm{d}}(m, n)$ is the corrupted image intensity in the density domain, $s_{\mathrm{e}}(m, n)$ is the original scene, $v_{\mathrm{d}}(m, n)$ is the additive white-Gaussian noise in the density domain with variance σ_{d}^2, h is the blur function and α, β are film-dependent parameters. For a linear space-invariant blur, the blur function is a convolution. To obtain the following equation, (12) can be rearranged

$$r_{\mathrm{e}}(m, n) = v_{\mathrm{e}}(m, n)\{h(s_{\mathrm{e}}(m, n))\} \tag{13}$$

where $v_{\mathrm{e}}(m, n) = 10^{v_{\mathrm{d}}(m,n)/\alpha}$ and $r_{\mathrm{e}}(m, n)$ is the observed image in the exposure domain. As can be seen from (13), because of sensor nonlinearity, the noise component in the density domain manifests itself as a multiplicative non-Gaussian noise in the exposure domain.

When the blur is linear and space-invariant, the observation equation in the density domain can be written as

$$\begin{pmatrix} Y_{\mathrm{d}6} \\ Y_{\mathrm{d}7} \end{pmatrix} = \alpha \log_{10}\{\boldsymbol{H}\boldsymbol{X}(k)\} + \begin{pmatrix} V_{\mathrm{d}6} \\ V_{\mathrm{d}7} \end{pmatrix} + \beta \tag{14}$$

Here, $V_{\mathrm{d}6}$ and $V_{\mathrm{d}7}$ are independent realisations of Gaussian noise in the density domain. We can obtain $Y_{\mathrm{d}6}$ and $Y_{\mathrm{d}7}$ from (12). The PF for handling the nonlinearity involved here differs from the one in Case A in the expression for the re-sampling weight. The weight of the ith particle $\tilde{\boldsymbol{X}}^{(i)}(k)$ is found as

$$q^{(i)} = \frac{\exp[-\|\boldsymbol{Y}_{\mathrm{d}}(k) - \alpha \log_{10}\{\boldsymbol{H}\tilde{\boldsymbol{X}}^{(i)}(k)\} - \beta\|^2/2\sigma_{\mathrm{d}}^2]}{\sum_{j=1}^{K_0}\exp[-\|\boldsymbol{Y}_{\mathrm{d}}(k) - \alpha \log_{10}\{\tilde{\boldsymbol{H}}\boldsymbol{X}^{(j)}(k)\} - \beta\|^2/2\sigma_{\mathrm{d}}^2]} \tag{15}$$

where $\boldsymbol{Y}_{\mathrm{d}}(k)$ is the observation vector in the density domain. The set of samples $\{\tilde{\boldsymbol{X}}^{(i)}(k), i = 1, 2, \ldots, K_0\}$ are updated using the respective weights $q^{(i)}$.

We see that the proposed method is quite general and is amenable to handle any nonlinearity or non-Gaussianity. Only the expression for the re-sampling weight changes, depending on the degradation model and the noise statistics. The PF arrives at the conditional mean as the optimal estimate without involving any computationally complex equations, unlike the KF. In spite of its generality, it is a considerably simpler algorithm to implement than the KF.

5 Experimental results

The quality of restoration is typically quantified in terms of the improvement in signal-to-noise ratio (ISNR), which is defined as $\mathrm{ISNR} = 10 \log_{10}[\sum_{i,j}(x(i, j) - y(i, j))^2 / \sum_{i,j}(x(i, j) - \hat{x}(i, j))^2]$ dB. Here $x(\cdot,\cdot)$ is the original image, $y(\cdot,\cdot)$ is the degraded image and $\hat{x}(\cdot,\cdot)$, the restored image. Even though the ISNR is a widely used parameter, perceptual improvement in quality is also important to evaluate the performance of any restoration scheme.

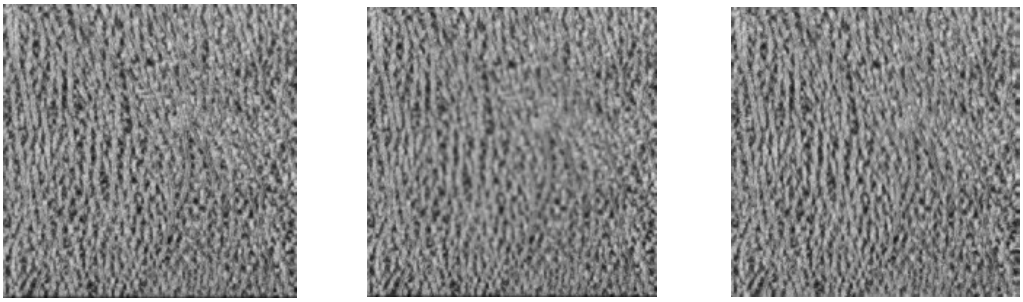

a *b* *c*

Fig. 3 *Restoration on the 'Calf' image*

a Original image
b Degraded image
c PF output, 8000 samples

Case A: Linear degradation: Here, we consider the case when the degradation is linear and the noise is additive Gaussian. We first present restoration results on a 200×200 'Calf' texture image (Fig. 3). The original image is modelled as an AR image and the parameters of the model are calculated using the Yule–Walker equations [9]. Space-variant, Gaussian blur with spatial distribution of the spread parameter σ, as shown in Fig. 4, is used to blur the original image, and Gaussian noise with variance 30 was then added to the blurred image. Fig. 3*a* shows the original 'Calf' image and Fig. 3*b* shows the degraded (blurred and noisy) image. The restoration obtained by using the PF with $K_0 = 8000$ samples is shown in Fig. 3*c*. As can be seen, the quality of restoration is very good if we use sufficiently large number of samples. The restored image is very close to the original image in appearance. The ISNR was 3.2 dB. Note that the proposed recursive technique has been able to remove the non-causal space-variant blur quite well. Equally good results were obtained for several other texture images taken from the Brodatz album.

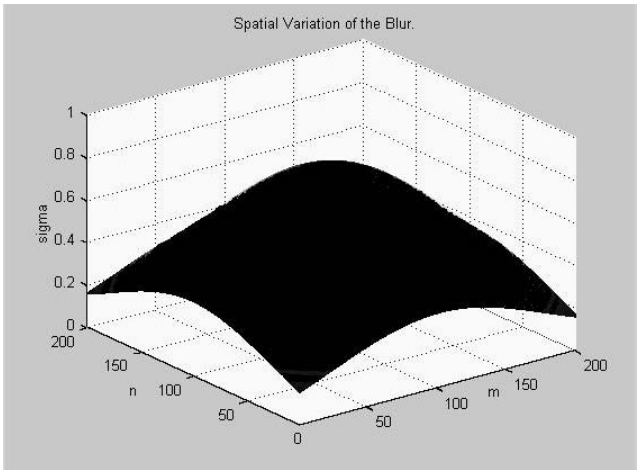

Fig. 4 *Spatial variation of the blur parameter σ for the 'Calf' image*

The proposed method was also tested on real images. Unlike texture images, the real images do not fit into an AR model very well and the restoration obtained when the algorithm was applied without any modification was not satisfactory. As is to be expected, the restored images were characterised by artefacts in regions with sharp edges. To overcome this problem, the following procedure was adopted. If the local variance of the image pixels corresponding to the current state vector exceeds the AR model variance by more than a certain pre-defined factor (chosen here as 5), then the local variance is used for the driving noise instead of the AR model variance in the prediction phase of the algorithm. This is done for all the ten concurrent block predictors. After this adjustment was made, good results were obtained even with real images. As an example, we present restoration result obtained on a 128×128 'Tank' image degraded by space-invariant Gaussian blur with blur parameter 0.7. The blurred image was corrupted with additive Gaussian noise of variance 10. Fig. 5*a* shows the original 'Tank' image and Fig. 5*b* shows its degraded version. As discussed previously, we increased the variance of the AR model noise as well as the number of samples. This non-stationary image model allows us to densely sample from a wide range of grey levels that in turn makes it possible to handle edges. The PF filter output obtained by using $K_0 = 9000$ samples is shown in

a *b* *c*

Fig. 5 *Restoration on the 'Tank' image*

a Original image
b Degraded image
c PF output, 9000 samples

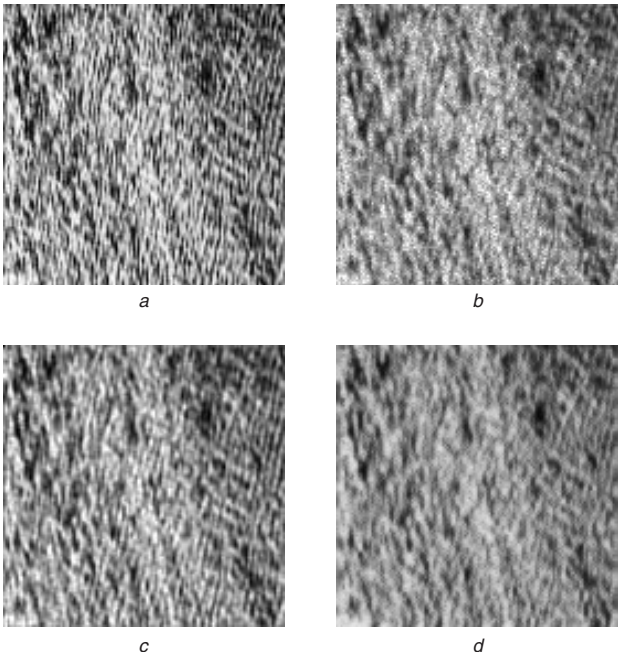

Fig. 6 *Comparison of MWF and PF for the 'Water' image*

a Original image
b Degraded image
c MWF output
d PF output, 10 000 samples

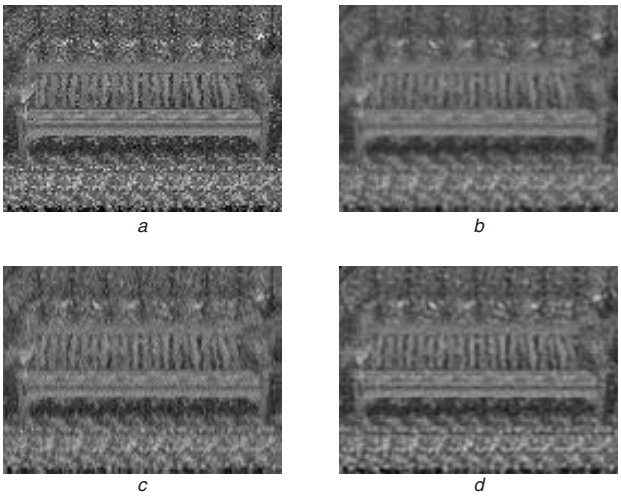

Fig. 7 *Comparison of MWF and PF for the 'Bench' image*

a Original image
b Degraded image
c MWF output
d PF output, 9000 samples

Fig. 5*c*. The recovered image is free of artefacts and preserves the details present in the original image. The ISNR value was found to be 0.8 dB.

The algorithm needs a large number of samples to arrive at an acceptable output because the dimensionality of the state space is 13, which is quite high. For most PF-based methods, it is well known [10] that the number of required samples increases rapidly with the dimension of the state vector.

Case B: Nonlinear/non-Gaussian degradation: In Section 4.1, we discussed the photographic film and showed that because of the logarithmic (nonlinear) film characteristic, the noise becomes multiplicative and non-Gaussian in the exposure domain. We now demonstrate the performance of the PF in handling film-grain noise. We also compare the PF output with that obtained by using the modified Wiener filter (MWF). The MWF was proposed by Pavlović and Tekalp [14] to restore images corrupted by film-grain noise. The MWF is a linear optimal filter that is designed to handle film-grain noise and blur. However, it cannot be applied to space-variant situations.

We first consider a 128 × 128 'Water' texture image shown in Fig. 6*a*. The image is degraded by a space-invariant Gaussian blur of spread 0.7 and additive noise with variance 0.06 in the density domain (Fig. 6*b*). We assume the film-dependent parameters α and β to be equal to 5 and -5, respectively. We use the MWF as well as the PF to restore the degraded image. Comparing Fig. 6*c* (the MWF output) and 6*d* (the PF output, 10 000 samples), we observe that the PF clearly outperforms the MWF. The MWF output is somewhat blurred. In contrast, the PF output is sharp and the details are better discernible in Fig. 6*d*. The ISNR values for the MWF and the PF were 0.197 and 0.99 dB, respectively.

Next, we consider the 95 × 128 'Bench' image shown in Fig. 7*a*. The image is degraded by a space-invariant Gaussian blur of spread 0.7. The variance of the Gaussian noise in the density domain was 0.01. The degraded image is shown in Fig. 7*b*. The MWF output is shown in Fig. 7*c* (ISNR for the MWF = 0.109 dB) and the PF output is shown in Fig. 7*d* (ISNR for the PF = 0.61 dB). The PF output is better than the MWF output in terms of perceptual quality also. For example, the vertical columns on the bench and the horizontal frame of the bench are less blurred and less noisy in the output image of the PF.

It must be noted that the MWF is optimal only among the class of linear filters. The estimates of linear predictors are optimal in the MMSE sense only when the observation and the original state are jointly Gaussian. As this condition is not satisfied for the photographic film, the MWF (being a linear estimator) cannot produce optimum results. The PF however yields optimal estimates (provided the number of

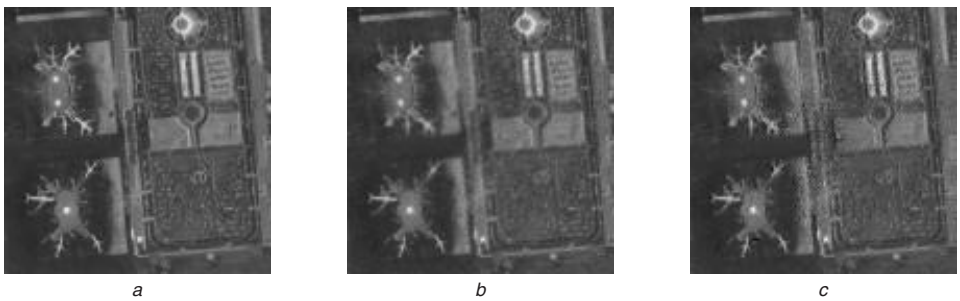

Fig. 8 *Restoration on the 'Airport' image*

a Original image
b Degraded image
c PF output, 10 000 samples

samples is sufficiently large) under all circumstances, as it does not constrain the blur or image models to be linear. This is the reason why the PF performs better than the MWF in restoring photographic images with film-grain noise.

Finally, we present a result obtained using the PF for the case when the blur is space-variant and the noise is multiplicative. The space-variant Gaussian blur has a maximum spread of 0.8 and has a spatial distribution as shown in Fig. 4. The noise had a variance of 0.01 in the density domain. The original 'Airport' image is shown in Fig. 8*a*. The degraded image and the PF output ($K_0 = 10\,000$) are shown in Figs. 8*b* and *c*, respectively. The improvement in perceptual quality in the restored image is quite evident in Fig. 8*c*, especially in the central region of the image. The ISNR value was found to be 1.87 dB.

The experiments described in this section clearly reveal the generality of the proposed method and its suitability for the image restoration problem.

6 Conclusions

A new method that incorporates full-plane ROSs for the image model as well as blur within a recursive framework was developed. The technique makes use of the PF to handle general types of degradations. The efficacy of the proposed method was tested on several images degraded by different types of degradations, and the results were found to be quite good. The algorithm is powerful and yet simple to implement. It is not constrained by considerations of analytical tractability, but is computationally intensive as it works by propagating a very large number of samples.

In future, one could explore inhomogeneous AR models to improve restoration results on real images. Also, the utility of linear AR models driven by non-Gaussian noise [15] for the state model can be investigated within the proposed framework.

7 References

1 Jain, A.K.: 'Fundamentals of digital image processing' (Prentice-Hall, Englewood Cliffs, NJ, 1989)

2 Patti, A.J., Ozkan, M.K., Tekalp, A.M., and Sezan, M.I.: 'New approaches for space-variant image restoration'. Proc. IEEE ICASSP, 1993, pp. 261–264

3 Lagendijk, R.L., Biemond, J., and Boekee, D.E.: 'Identification and restoration of noisy blurred images using the expectation-maximization algorithm', *IEEE Trans. Acoust. Speech Signal Process.*, 1990, **38**, pp. 1180–1189

4 Woods, J.W., and Radewan, C.H.: 'Kalman filtering in two dimensions', *IEEE Trans. Inf. Theory*, 1977, **23**, pp. 473–482

5 Woods, J.W., and Ingle, V.K.: 'Kalman filtering in two dimensions: further results', *IEEE Trans. Acoust. Speech Signal Process.*, 1981, **29**, pp. 188–197

6 Citrin, S., and Azimi-Sadjadi, R.: 'A full-plane block Kalman filter for image restoration', *IEEE Trans. Image Process.*, 1992, **1**, pp. 488–495

7 Tekalp, A.M., Kaufman, H., and Woods, J.W.: 'Identification of image and blur parameters for the restoration of noncausal blurs', *IEEE Trans. Acoust. Speech Signal Process.*, 1986, **34**, pp. 963–971

8 Sezan, M.I., and Tekalp, A.M.: 'Survey of recent developments in digital image restoration', *Opt. Eng.*, 1990, **29**, pp. 393–404

9 Lim, J.S.: 'Two-dimensional signal and image processing' (Prentice-Hall, Englewood Cliffs, NJ, 1990)

10 Gordon, N.J., Salmond, D.J., and Smith, F.M.A.: 'Novel approach to nonlinear/non-Gaussian Bayesian state estimation', *IEE Proc. F, Radar Signal Process.*, 1993, **140**, pp. 107–113

11 Arulampalam, S., Maskell, S., Gordon, N., and Clapp, T.: 'A tutorial on particle filters for online nonlinear/non-Gaussian tracking', *IEEE Trans. Signal Process.*, 2002, **50**, pp. 174–188

12 Anderson, B.D.O., and Moore, J.B.: 'Optimal filtering' (Prentice-Hall, Englewood Cliffs, NJ, 1979)

13 Smith, A.F.M., and Gelfand, A.E.: 'Bayesian statistics without tears: a sampling–resampling perspective', *J. Am. Stat. Assoc.*, 1992, **46**, pp. 84–88

14 Pavlović, G., and Tekalp, A.M.: 'Image restoration with multiplicative noise: incorporating sensor nonlinearity', *IEEE Trans. Signal Process.*, 1991, **39**, pp. 2132–2136

15 Kadaba, S.R., Gelfand, S.B., and Kashyap, R.L.: 'Recursive estimation of images using non-Gaussian autoregressive models', *IEEE Trans. Image Process.*, 1998, **7**, pp. 1439–1452