BLIND IMAGE RESTORATION USING A BLOCK-STATIONARY SIGNAL MODEL

Tom E. Bishop and James R. Hopgood

IDCOM, School of Engineering and Electronics, University of Edinburgh,
King's Buildings, Mayfield Road, Edinburgh EH9 3JL, UK
{t.e.bishop,james.hopgood}@ed.ac.uk

ABSTRACT

We present a novel method for Blind image restoration which is a multidimensional extension of an approach used successfully for audio restoration. A nonstationary image model is used to increase reliability of blur estimates. This source model consists of a separate autoregressive model in each region of the image. A hierarchical Bayesian model for the observations is used, and a maximum marginalised *a posteriori* (MMAP) blur estimate is obtained by optimising the resulting probability density function.

1. INTRODUCTION

With the recent rapid growth of digital photography, using small low-cost optics in conjunction with signal processing to enhance or restore images is an ever more attractive option, due to increasingly abundant processing power. In practice, many images suffer from blur due to optical imprecision, mis-focussing or motion blur; exact specification of these blurs is usually unknown *a priori*.

Blind image restoration (BIR) attempts to tackle the problem by estimating the degradation from the image itself, using available prior knowledge about the general nature of the image and blur. Existing methods have used either simplistic models for the image and blur or imposed deterministic constraints on the image. A summary of earlier methods is presented in [1]. Methods may be classed as *a priori* (estimating the blur then restoring the image), *joint* (simultaneous estimation and restoration), or *direct* (bypassing blur identification, as in some recent multichannel methods).

We propose a new *a priori* blur identification method, using a maximum marginalised *a posteriori* (MMAP) formulation. In essence, this involves evaluating the goodness of fit of the restored image to a nonstationary image model, thereby finding a blur estimate.

The majority of existing work has been on joint methods, where it is common to estimate the image and blur alternately. Many use a form of alternating minimisation (AM) algorithm, as an extension to regularisation theory. Total variation (TV) smoothness criteria have featured in recent work [2], attempting to minimise a cost function based on data fidelity and piecewise smoothness of the image and blur. Extensions to this idea have shown promise [3]. Other earlier methods use deterministic constraints [1] imposed at each iteration, also switching between the image and blur; these tended to be either unstable or not general-purpose. Convergence of many of these algorithms is either ill-defined or is to local maxima.

The use of parametric models for the observed image has been considered by several authors, typically using autoregressive moving average (ARMA) models; for example, maximum likelihood (ML) estimation has been considered in [4]. Though the ML formulation itself is essentially an *a priori* parameter identification, use of the expectation maximisation (EM) algorithm provides an external problem which incorporates image restoration as well as blur identification resulting in a joint procedure. Again this method may converge to a local maximum.

This work is supported by a UK-EPSRC studentship.

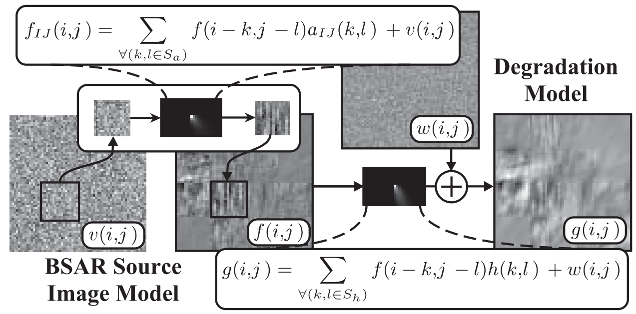

Fig. 1. Generative source-image & blur model

The earliest work on *a priori* blur identification is based on methods that looked for patterns of zero crossings or spikes in the spectral or cepstral domains [5]. Similar to our proposed technique, these methods in fact use a block-stationary approach. Presented through empirical arguments, results show some limited success. By averaging together the blocks, the stationary term in the blur becomes more identifiable over the mean nonstationary image. We aim to combine these ideas with a more formal image model.

In addition — although not stated as such — many other techniques using cost functionals or deterministic constraints impose implicit stationary image models that are overly simplistic, assuming homogeneous images. In this paper, a more robust optimal solution is developed using stochastic models and Bayesian parameter estimation techniques. Nonstationary models aid in blur identifiability compared to existing stationary image models in [4]. Interestingly, a block-stationary AR (BSAR) model is also used in [3] for blur support estimation; in this work we apply these models to blur identification.

2. BIR PROBLEM FORMULATION

The $m \times n$ observed image, $g(i,j)$, is modelled as the discrete convolution of an unobserved ideal (non-degraded) image, $f(i,j)$, and a spatially invariant point-spread function (PSF), $h(k,l)$, with additive white Gaussian noise (WGN), $w(i,j)$. This general linear model (see Fig. 1) can be expressed in matrix-vector form:

$$\boldsymbol{g} = \boldsymbol{H}\boldsymbol{f} + \boldsymbol{w}, \tag{1}$$

with lexicographically ordered images $\boldsymbol{f}$ and $\boldsymbol{g}$ and block Toeplitz with Toeplitz blocks (BTTB) degradation matrix $\boldsymbol{H}$ [6]. The task of

BIR is to estimate $\boldsymbol{f}$ from $\boldsymbol{g}$ where the PSF is unknown. This discrete ill-posed problem requires prior knowledge to constrain the solution. We use a hierarchical Bayesian image model to achieve this.

2.1. MMAP Methodology

A stochastic BSAR image model and a spatially-invariant blur model are used. The likelihood is formed for the observed image, given the true image model and its parameters. The Bayesian paradigm allows the *posterior* probability density function (PDF) to be expressed, using priors for the model parameters. The actual values of the image model parameters are not needed for deblurring, only those of the blur. They may be *marginalised* to yield the PDF of the blur parameters alone given the observed image. This function may then be optimised numerically, giving the blur estimate which restores the image optimally under the constraints of the model. In theory, while optimisation may be performed on the full joint PDF, marginalisation reduces the search space dimension.

2.2. Nonstationary Image Model

A linear model is used for the idealised source image (see Fig. 1). The BSAR model splits f into $M \times N$ blocks. Each block is modelled by a 2D autoregressive (AR) process of order $(m_a - 1) \times (n_a - 1)$. Present work considers a causal model, with $P = MN(m_a n_a - 1)$ AR coefficients, $a_{IJ}(k, l)$. The stacked vector of these lexicographically ordered coefficients is $\boldsymbol{a}$. The image model is written:

$$\boldsymbol{f} = \boldsymbol{A}\boldsymbol{f} + \boldsymbol{v}, \tag{2a}$$

or equivalently,

$$\boldsymbol{f} = \boldsymbol{F}\boldsymbol{a} + \boldsymbol{v}. \tag{2b}$$

In (2a), the matrix $\boldsymbol{A}$ is of a similar form to $\boldsymbol{H}$, with $m \times m$ square $n \times n$ blocks; however it is only BTTB if $M = N = 1$. In (2b), the *data matrix*, $\boldsymbol{F}$, is again BTTB only in the case of a single block, and is not square. Moreover, it is of size $mn \times P$; the $((i-1)n + j)^{\text{th}}$ row includes the pixel values $f(k, l)$, such that $(k, l) \in S_a(i, j)$ (see Fig. 2). $S_a(i, j)$ is the causal AR support region for pixel (i, j), containing the $m_a n_a - 1$ pixel locations upon which $f(i, j)$ depends. The excitation noise or modelling error, $\boldsymbol{v}$, is assumed to be WGN.

The PDF of the source and blurred image may be found by linear probability transformations of the excitation and noise variables, $\boldsymbol{v} \sim \mathcal{N}(0, \boldsymbol{Q}_v)$ and $\boldsymbol{w} \sim \mathcal{N}(0, \boldsymbol{Q}_w)$. Since $\boldsymbol{v} = (\boldsymbol{I}_{mn} - \boldsymbol{A})\boldsymbol{f}$,

$$\begin{aligned} p\left(\boldsymbol{f} \mid \boldsymbol{a}, \boldsymbol{Q}_v\right) &= \mathcal{N}\left(\boldsymbol{f} \mid 0, \boldsymbol{\Sigma}_f\right) \\ &= \sqrt{\frac{\det|\boldsymbol{I}_{mn} - \boldsymbol{A}|^2}{(2\pi)^{mn} \det|\boldsymbol{Q}_v|}} \exp\left[-\frac{1}{2}\boldsymbol{f}^T \boldsymbol{\Sigma}_f^{-1} \boldsymbol{f}\right], \end{aligned} \tag{3}$$

with covariance $\boldsymbol{\Sigma}_f = \mathbb{E}[\boldsymbol{f}\boldsymbol{f}^T] = (\boldsymbol{I}_{mn} - \boldsymbol{A})^{-1}\boldsymbol{Q}_v(\boldsymbol{I}_{mn} - \boldsymbol{A})^{-T}$. The Jacobian $\det|\boldsymbol{I}_{mn} - \boldsymbol{A}|$ is unity in the case of a causal image model since $\boldsymbol{A}$ is (lower) triangular. $\boldsymbol{I}_p$ is the $p \times p$ identity matrix. $\boldsymbol{Q}_v$ is diagonal, and allows an individual variance, σ^2_{IJ}, in each block. Experimental results show this model gives more accurate blur estimates than a global variance across the whole image.

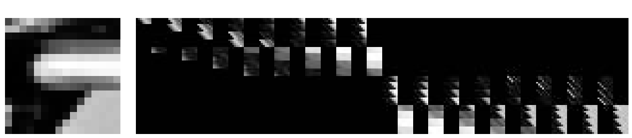

(a) (b)

Fig. 2. (a) A 16×16 pixel image, $\boldsymbol{f}$, and (b) corresponding 60×256 data matrix, $\boldsymbol{F}^T$, formed for $M = N = 2$ and $m_a = n_a = 4$.

2.3. Likelihood for blurred image

The likelihood is obtained by combining (1) and (3), to give [4]:

$$p(\boldsymbol{g}|\boldsymbol{a}, \boldsymbol{h}, \boldsymbol{Q}_v, \boldsymbol{Q}_w) = p(\boldsymbol{H}\boldsymbol{f}) * p(\boldsymbol{w}) = \mathcal{N}\left(\boldsymbol{g} \mid 0, \boldsymbol{\Sigma}_g\right), \tag{4}$$

where $\boldsymbol{\Sigma}_g = \boldsymbol{H}(\boldsymbol{I}_{mn} - \boldsymbol{A})^{-1}\boldsymbol{Q}_v(\boldsymbol{I}_{mn} - \boldsymbol{A})^{-T}\boldsymbol{H}^T + \boldsymbol{Q}_w$, and $\boldsymbol{H}$ is parameterised by $\boldsymbol{h}$, the vector of coefficients defining the PSF $h(k, l)$. Unfortunately, with this form it is difficult to analytically perform the marginalisation in §2.6, due to the presence of $\boldsymbol{Q}_w$ in $\boldsymbol{\Sigma}_g$. Thus, as an approximation $\boldsymbol{Q}_w$ is omitted such that the exponent in the likelihood becomes $\exp[-\frac{1}{2}\hat{\boldsymbol{f}}^T(\boldsymbol{I} - \boldsymbol{A})^T\boldsymbol{Q}_v^{-1}(\boldsymbol{I} - \boldsymbol{A})\hat{\boldsymbol{f}}]$.

For clarity, the term $\boldsymbol{H}^{-1}\boldsymbol{g}$, representing an estimate of $\boldsymbol{f}$ (see §2.8), is denoted $\hat{\boldsymbol{f}}$, with coresponding data matrix $\hat{\boldsymbol{F}}$.

2.4. Prior Distributions for source & blur model

To form the hierarchical Bayesian model, priors are placed upon the other model parameters. These must be both representative of our existing beliefs and allow for mathematically tractable results. Presently, a standard Gaussian prior is chosen for $\boldsymbol{a}$:

$$p\left(\boldsymbol{a}\right) = \mathcal{N}\left(\boldsymbol{a} \,\middle|\, \boldsymbol{0}, \delta^2 \boldsymbol{R}_v\right) \tag{5}$$

and an inverse-Gamma (IG) prior for the excitation variance:

$$p(\boldsymbol{Q}_v) = \prod_{IJ} p(\sigma^2_{IJ}|\alpha_{I,J}, \beta_{IJ}), \quad \text{and for each variance}$$

$$p(\sigma^2|\alpha, \beta) = \mathcal{IG}(\sigma^2|\alpha, \beta) = \begin{cases} \frac{\beta^\alpha \sigma^{-2(\alpha+1)}}{\Gamma(\alpha)} e^{-\frac{\beta}{\sigma^2}} & \sigma^2 > 0 \\ 0 & \text{otherwise} \end{cases}. \tag{6}$$

An uninformative prior with large variance is currently used for $p(\boldsymbol{h})$. Hyperparameters δ and $\{\alpha_{IJ}, \beta_{IJ}\}$ may also be chosen to make the priors tend to be uninformative, although δ may be set to model the variance of stable AR coefficients. The $P \times P$ matrix $\boldsymbol{R}_v$, is diagonal and constructed with the variances for each block. The image and blur priors are independent: $p(\boldsymbol{h}, \boldsymbol{a}, \boldsymbol{Q}_v) = p(\boldsymbol{h})p(\boldsymbol{a}|\delta, \boldsymbol{Q}_v)p(\boldsymbol{Q}_v)$.

2.5. Posterior Distribution

The posterior is found by applying Bayes' rule to the likelihood:

$$p(\boldsymbol{h}, \boldsymbol{a}, \boldsymbol{Q}_v|\boldsymbol{g}) = \frac{p(\boldsymbol{g}|\boldsymbol{h}, \boldsymbol{a}, \boldsymbol{Q}_v)p(\boldsymbol{h}, \boldsymbol{a}, \boldsymbol{Q}_v)}{p(\boldsymbol{g})} \tag{7}$$

$$\begin{aligned} &\propto \frac{\det|\boldsymbol{Q}_v|^{-\frac{1}{2}} \det|\boldsymbol{R}_v|^{-\frac{1}{2}}}{(2\pi)^{\frac{(mn+P)}{2}} \det|\boldsymbol{H}|} \exp\left[\frac{\boldsymbol{a}^T\boldsymbol{R}_v^{-1}\boldsymbol{a}}{\delta^2}\right] p(\boldsymbol{Q}_v) \\ &\quad \cdot \exp\left[-\frac{1}{2}(\hat{\boldsymbol{f}} - \hat{\boldsymbol{F}}\boldsymbol{a})^T \boldsymbol{Q}_v^{-1}(\hat{\boldsymbol{f}} - \hat{\boldsymbol{F}}\boldsymbol{a})\right] \\ &\propto \frac{p(\boldsymbol{Q}_v)\det|\boldsymbol{R}_v|^{-\frac{1}{2}}}{(2\pi)^{\frac{(mn+P)}{2}} \det|\boldsymbol{H}| \det|\boldsymbol{Q}_v|^{\frac{1}{2}}} \\ &\quad \cdot \exp\Bigg[-\frac{1}{2}\Big(\boldsymbol{a}^T\{\hat{\boldsymbol{F}}^T\boldsymbol{Q}_v^{-1}\hat{\boldsymbol{F}} + \delta^{-2}\boldsymbol{R}_v^{-1}\}\boldsymbol{a} \\ &\quad - 2\boldsymbol{a}^T\hat{\boldsymbol{F}}^T\boldsymbol{Q}_v^{-1}\hat{\boldsymbol{f}} + \hat{\boldsymbol{f}}^T\boldsymbol{Q}_v^{-1}\hat{\boldsymbol{f}}\Big)\Bigg]. \end{aligned} \tag{8}$$

Notice that since it is only necessary to find a function that can be maximised by varying the parameters, any constants of proportionality not dependant on the model parameters may be ignored.

It would also be possible in this case to find a representation of $p(\boldsymbol{h}, \boldsymbol{a}, \boldsymbol{Q}_v, \boldsymbol{Q}_w|\boldsymbol{g})$ from the full joint PDF $p(\boldsymbol{g}, \boldsymbol{f}, \boldsymbol{h}, \boldsymbol{a}, \boldsymbol{Q}_v, \boldsymbol{Q}_w)$

by integrating out $\boldsymbol{f}$; this is termed *evidence analysis* in [7]. We apply a similar principle in §2.6 to estimate $\boldsymbol{h}$ alone.

2.6. Marginalisation of Nuisance Parameters

The AR model parameters are not directly relevant to finding the probability of the blur parameters. Integrating the posterior with respect to these *nuisance parameters*, the marginal PDF is:

$$p(\boldsymbol{h}|\boldsymbol{g}) = \int \cdots \int p(\boldsymbol{h}, \boldsymbol{a}, \boldsymbol{Q}_v|\boldsymbol{g})\, \mathrm{d}\boldsymbol{a} \cdot \mathrm{d}\boldsymbol{Q}_v. \tag{9}$$

The integration is done in a similar manner to that described in [8], via standard Gaussian and Gamma integrals. This results in equation (10) overleaf. In order to marginalise the variances, the image is split into its component blocks $\hat{\boldsymbol{f}}_{IJ}$ and the corresponding data matrices $\boldsymbol{F}_{IJ}$ are formed by extracting the appropriate sections of $\boldsymbol{F}$; note that construction of each $\boldsymbol{F}_{IJ}$ depends on the surrounding blocks.

2.7. Blur Parameter Estimation

This function is numerically optimised over the space of possible blurs; it is usual to take the negative log and minimise:

$$\hat{\boldsymbol{h}} = \arg\min_{\boldsymbol{h}} \big(-\ln p\,(\boldsymbol{h}|\boldsymbol{g}) \big). \tag{11}$$

For each blur under test, the procedure essentially consists of finding a restoration $\hat{\boldsymbol{f}}$ of the whole image using this particular $\hat{\boldsymbol{H}}$, then (10) effectively checks how well this conforms to our model. The term

$$\tilde{\boldsymbol{f}}_{IJ} = \hat{\boldsymbol{F}}_{IJ}\hat{\boldsymbol{a}}_{IJ} = \hat{\boldsymbol{F}}_{IJ}\Big\{\hat{\boldsymbol{F}}_{IJ}^T\hat{\boldsymbol{F}}_{IJ} + \delta^{-2}\boldsymbol{I}_{\frac{P}{MN}}\Big\}^{-1}\hat{\boldsymbol{F}}_{IJ}^T\hat{\boldsymbol{f}}_{IJ} \tag{12}$$

may be regarded as a prediction of $\hat{\boldsymbol{f}}_{IJ}$ using an estimate of the AR parameters for this block, $\hat{\boldsymbol{a}}_{IJ}$. Then the model fitting term given by $\hat{\boldsymbol{f}}^T(\hat{\boldsymbol{f}} - \tilde{\boldsymbol{f}})$, as used in the numerator of equation (10), is the prediction error weighted by the local image magnitude.

2.8. Blur Model

Although the method presented should apply to a general linear blur model, such as the commonly used non-causal moving average (MA) PSF model, in our experiments a causal parametric AR blur is used. There are several reasons for this: firstly, the complicated $\det|\boldsymbol{H}|$ term need not be calculated; secondly, this conforms to work in [8] where the method has been extensively tested in the 1D case.

Furthermore, an inverse filtering operation is used to find $\hat{\boldsymbol{f}} = \boldsymbol{H}^{-1}\boldsymbol{g}$. This does not incorporate any regularisation; however with an AR blur model, $\boldsymbol{H}^{-1}$ is well defined, and we are effectively estimating an MA restoration filter which does not severely amplify noise. To include regularisation in the model, a variant of the method using Gibbs sampling can be used and will be discussed in future work. A regularised estimate for $\hat{\boldsymbol{f}}$ could be used here, but as it is not part of the model, parameter estimates may not be optimal.

In practice, the $\boldsymbol{H}$ terms implicit in (11) are replaced by $(\boldsymbol{I}_{mn} - \boldsymbol{H}')^{-1}$, where this new $\boldsymbol{H}'$ is generated in a similar way to the original $\boldsymbol{H}$, and parameterised by a vector $\boldsymbol{h}'$. In the first order case, $\boldsymbol{h}' = [h_1, h_2, h_3]$, and the defining kernel h' used to construct $\boldsymbol{H}'$ is

$$h'(k,l) = \begin{bmatrix} h_3 & h_2 \\ h_1 & \bullet \end{bmatrix}, \tag{13}$$

where the $\bullet$ marks the center of the kernel image.

3. EXPERIMENTAL RESULTS

3.1. Synthetic Image, Synthetic Blur

The algorithm is first validated using synthesised data generated according to the BSAR model. Since there is no equivalent to the fundamental theorem of algebra in dimensions higher than one [1], it is not generally possible to produce a factorisation of the polynomial space to produce a pole-zero plot as in the 1D case [8]. Thus for visualisation on a 2D plot, we first consider only two blur parameters.

Experiments with various block sizes and model orders have been tried. In the example shown in Fig. 3(a), a first order AR process is used for the blur and the source image. WGN is used to drive the source, whose AR parameters $\hat{\boldsymbol{a}}_{IJ}$ for each 16×16 pixel block are estimated from the 256×256 *Cameraman* image. Noise is added to the blurred image at 40dB blurred-image SNR (BSNR) (see Fig. 3(b)). In Fig. 4, (10) is evaluated across a grid of points centred on the true blur parameters, in this example chosen as $\boldsymbol{h}' = [0.45, 0.35, 0.05]$. The first two parameters are chosen as the unknowns, and the last is assumed known, to enable plotting the PDF. The estimate of the blur parameters is chosen at the location of the minimum on the plot, in this case correctly estimated as $[\hat{h}_1, \hat{h}_2] = [0.45, 0.35]$.

The experiment is repeated, but with optimisation across all three parameters of $\boldsymbol{h}'$, using the deterministic nonlinear Nelder-Mead Simplex method, producing $\hat{\boldsymbol{h}}' = [0.442, 0.347, 0.058]$. The restored image using the inverse filter $(\boldsymbol{I} - \boldsymbol{H}')$ is shown in Fig. 3(c).

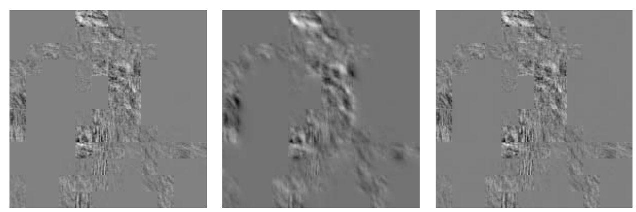

(a) $\boldsymbol{f}$ (b) $\boldsymbol{g}$ (c) $\hat{\boldsymbol{f}}$

Fig. 3. (a) 1st order BSAR synthetic image with 16×16 pixel blocks (b) Blurred image, at 40dB BSNR (c) Restored image. (central crops)

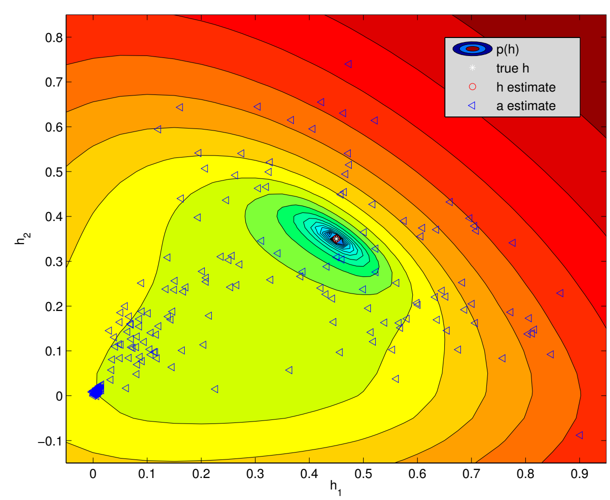

Fig. 4. Probability for AR blur parameters, $p(\boldsymbol{h}'|\boldsymbol{g})$, over the space of $[h_1, h_2]$. Also shown are the first two source AR parameters, $[a_1, a_2]_{IJ}$ for each block (triangles). Contours have been compressed to give more detailed coverage in flat parts of the PDF.

$$p(\boldsymbol{h}|\boldsymbol{g}) \propto \frac{1}{\det|\boldsymbol{H}|} \prod_{I=1}^{M} \prod_{J=1}^{N} \left(\frac{\left(\hat{\boldsymbol{f}}_{IJ}^{T} \left(\hat{\boldsymbol{f}}_{IJ} - \hat{\boldsymbol{F}}_{IJ} \{ \hat{\boldsymbol{F}}_{IJ}^{T} \hat{\boldsymbol{F}}_{IJ} + \delta^{-2} \boldsymbol{I}_{\frac{P}{MN}} \}^{-1} \hat{\boldsymbol{F}}_{IJ}^{T} \hat{\boldsymbol{f}}_{IJ} \right) + 2\beta_{IJ} \right)^{-\frac{\frac{mn}{MN}+1+2\alpha_{IJ}}{2}}}{\det \left| \hat{\boldsymbol{F}}_{IJ}^{T} \hat{\boldsymbol{F}}_{IJ} + \delta^{-2} \boldsymbol{I}_{\frac{P}{MN}} \right|^{\frac{1}{2}}} \right) \tag{10}$$

3.2. Real Image, Synthetic Blur

The 256×256 pixel *Cameraman* image is blurred with the same parameters, and noise added at 35dB BSNR, shown in Fig. 5. The blur is estimated using the same deterministic optimisation method, using 8×8 pixel blocks, with hyperparameters $\alpha_{IJ}, \beta_{IJ} = 0$, $\delta = 1$. The image model was found to need some modification however.

Firstly, unlike true AR signals, real images are inherently nonzero mean. Removing the global mean from the image is not enough; the estimated local sample mean of each block is required to more accurately model the true image. Thus in the estimation procedure, after $\hat{\boldsymbol{f}}$ is found, the mean of each block is subtracted. Care is required when subtracting these means from the matrix $\hat{\boldsymbol{F}}$, since it is the local mean for each block that should be used; as such any data elements providing boundary conditions for a block must use the mean of their neighbour and not their own block mean, and these will have different values depending on the block in question. This procedure avoids discontinuities in the data.

Furthermore, experimentation has shown that a 1st order model does not well represent a real image. Tests with different block sizes and model orders indicated 8 to 16 pixel blocks with 2nd to 8th order models tend to give good results. A 3rd order source model is used for the present experiment. With these changes — despite exclusion of the noise from the model — a successful restoration is still possible. At lower BSNRs, the restoration is still sharp, but the amplified noise becomes more prominent.

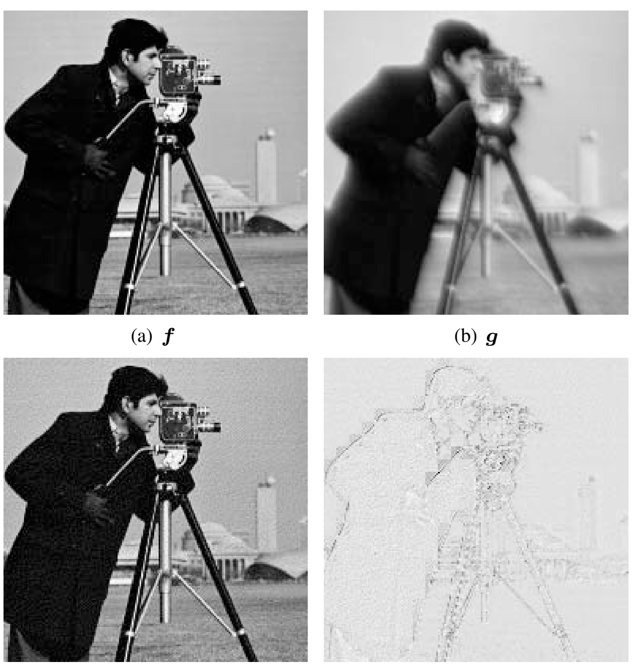

(a) $\boldsymbol{f}$ (b) $\boldsymbol{g}$ (c) $\hat{\boldsymbol{f}}$ (d) $\hat{\boldsymbol{f}}^T(\hat{\boldsymbol{f}} - \hat{\boldsymbol{A}}\hat{\boldsymbol{f}})$

Fig. 5. Cameraman (a) source; (b) blurred; (c) restored and (d) weighted prediction error images. Central crop shown.

The estimated blur is found as $\hat{\boldsymbol{h}}' = [0.473, 0.377, 0.011]$ in the 35dB case. Due to the slight error, a small number of superwhite and super-black pixels are produced with intensities outside the range of the source image (giving longer tails in the image histogram). Thus to display the image correctly, these tails should be clipped, or histogram specification may be used. The image estimate after clipping is shown in Fig. 5(c). The model fitting term (§2.7) is shown in Fig. 5(d).

4. CONCLUSIONS

This paper is primarily concerned with the use of a nonstationary image model, and a robust framework for blur identification using this model, rather than a particular optimisation approach. We have shown that the MMAP method provides a means to reliably estimate blur parameters, with the parameter-space dimensionality reduction provided by marginalisation, and nonstationary BSAR image model both reducing ambiguity in these estimates.

We note that while the causal AR blur model does not exactly match real blurs, the method should be extendible to other blur models. We are also investigating more advanced image models, with non-rectangular regions and local means, and Gibbs sampling and Markov chain Monte Carlo (MCMC) methods for optimisation under more complicated degradation models, including observation noise.

5. REFERENCES

[1] D. Kundur and D. Hatzinakos, "Blind image deconvolution: An algorithmic approach to practical image restoration," *IEEE Sig. Proc. Mag.*, vol. 13, no. 3, pp. 43–64, May 1996.

[2] Y.-L. You and M. Kaveh, "Blind image restoration by anisotropic regularization," *IEEE Trans. Image Processing*, vol. 8, no. 3, pp. 396–407, 1999.

[3] L. Chen and K.-H. Yap, "A soft double regularization approach to parametric blind image deconvolution," *IEEE Trans. Image Processing*, vol. 14, no. 5, pp. 624–633, 2005.

[4] R. L. Lagendijk, J. Biemond, and D. E. Boekee, "Identification and restoration of noisy blurred images using the expectation-maximization algorithm," *IEEE Trans. Acoustics, Speech, and Signal Proc.*, vol. 38, no. 7, pp. 1180–1191, Jul. 1990.

[5] M. Cannon, "Blind deconvolution of spatially invariant image blurs with phase," *IEEE Trans. Acoustics, Speech, and Signal Proc.*, vol. 24, no. 1, pp. 58–63, 1976.

[6] A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice Hall, New Jersey, first edition, 1989.

[7] R. Molina, A. K. Katsaggelos, and J. Mateos, "Bayesian and regularization methods for hyperparameter estimation in image restoration," *IEEE Trans. Im. Proc.*, vol. 8, pp. 231–246, 1999.

[8] J. R. Hopgood and P. J. W. Rayner, "Blind single channel deconvolution using nonstationary signal processing," *IEEE Trans. Speech and Audio Proc.*, vol. 11, no. 7, pp. 476–488, Sep. 2003.