Bayesian Image Estimation from an Incomplete Set of Blurred, Undersampled Low Resolution Images

Javier Mateos[1], Miguel Vega[2], Rafael Molina[1], and Aggelos K. Katsaggelos[3]⋆

[1] Dpto. de Ciencias de la Computación e I.A.
Universidad de Granada, 18071 Granada, Spain
[2] Dpto. de Lenguajes y Sistemas Informáticos
Universidad de Granada, 18071 Granada, Spain
[3] Dept. of Electrical and Computer Engineering
Northwestern University, Evanston, IL 60208-3118

Abstract. This paper deals with the problem of reconstructing a high-resolution image from an incomplete set of undersampled, blurred and noisy images shifted with subpixel displacement. We derive mathematical expressions for the calculation of the maximum a posteriori estimate of the high resolution image and the estimation of the parameters involved in the model. We also examine the role played by the prior model when this incomplete set of low resolution images is used. The performance of the method is tested experimentally.

1 Introduction

High resolution images can, in some cases, be obtained directly from high precision optics and charge coupled devices (CCDs). However, due to hardware and cost limitations, imaging systems often provide us with only multiple low resolution images. In addition, there is a lower limit as to how small each CCD can be, due to the presence of shot noise [1] and the fact that the associated signal to noise ratio (SNR) is proportional to the size of the detector [2].

Over the last two decades research has been devoted to the problem of reconstructing a high-resolution image from multiple undersampled, shifted, degraded frames with subpixel displacement errors (see [3] for a review). Most of the reported work addresses the problem of estimating an $LM \times LN$ high resolution image from at least $L \times L$ low resolution images of size $M \times N$, that is, when the number of available low resolution images is at least equal to L^2, where L is the magnifying factor. In Molina et al. [4] and Abad et al. [5] a method for simultaneously estimating the high resolution image and the associated parameters within the Bayesian model is presented. Kim et al. [6] explore the conditions the shifts of the $L \times L$ low resolution images have to satisfy in order to solve the

high resolution problem, at least from the least squares perspective. Elad and Feuer [7] study the same problem when combining Bayesian, Projection onto Convex Sets and Maximum Likelihood methodologies on high resolution problems. Baker and Kanade [8] also examine the impact of increasing the number of low resolution images, when proposing an alternative approach to the super resolution problem. However, not much work has been reported on the role played by the prior model when the system is incomplete, that is, when we have less than $L \times L$ low resolution images or when the shifts do not satisfy the conditions in [6] or [7].

In our previous work [9] we proposed a new method to solve the high resolution problem when no blurring was present in the observation process. The method was based on the general framework for frequency domain multi-channel signal processing developed by Katsaggelos et al. in [10] (a formulation that was also later obtained by Bose and Boo [11] for the high resolution problem). In this paper we extend the approach in [9] by considering that the low resolution images are also blurred, a case that frequently appears in Astronomy (see [12] for instance) and remote sensing. We also examine how the prior model compensates for the lack of information in the incomplete observation set.

The rest of the paper is organized as follows. The problem formulation is described in section 2. In section 3 the degradation and image models used in the Bayesian paradigm are described. The application of the Bayesian paradigm to calculate the MAP high resolution image and estimate the hyperparameters is described in section 4. Experimental results are described in section 5. Finally, section 6 concludes the paper.

2 Problem Formulation

Consider a camera sensor with $N_1 \times N_2$ pixels and assume that we have a set of q shifted images, $1 \leq q \leq L \times L$. Our aim is to reconstruct an $M_1 \times M_2$ high resolution image with $M_1 = L \times N_1$ and $M_2 = L \times N_2$, from the set of low-resolution observed images.

The low resolution sensors are shifted with respect to each other by a value proportional to $T_1/L \times T_2/L$, where $T_1 \times T_2$ is the size of each sensing element (note that if the sensors are shifted by values proportional to $T_1 \times T_2$ or $q < L \times L$, the high-resolution image reconstruction problem becomes singular). In this paper we assume that the normalized horizontal and vertical displacements are known (see [11, 13] for details). See [14] for an approach where the displacements are assumed unknown and are estimated simultaneously with the high-resolution image.

Let $\boldsymbol{g}_{l1,l2}$ be the $(N_1 \times N_2) \times 1$ observed low resolution image acquired by the $(l1, l2)$ sensor. Our goal is to reconstruct $\boldsymbol{f}$, the $(M_1 \times M_2) \times 1$ high resolution image, from a set of q low resolution images $\boldsymbol{g}_{l1,l2}$, with $1 \leq q \leq L^2$. We will denote by $\mathcal{I}$, the set of indices of the available low resolution images.

The process to obtain the observed low resolution image by the $(l1, l2)$ sensor, $\boldsymbol{g}_{l1,l2}$, from $\boldsymbol{f}$ can be modeled as a two stages process as follows. In the first stage,

the optical distortion in the observation process is represented by $\boldsymbol{B}_{l1,l2}$, the $(M_1 \times M_2) \times (M_1 \times M_2)$ point spread function defining a systematic blur, assumed to be known, due , for example, to motion or out of focus blurring, defects in the camera optics, etc. The second stage models the CCD pixel resolution. Let $\boldsymbol{H}_{l1,l2}$ be an $(M_1 \times M_2) \times (M_1 \times M_2)$ integrating matrix that represents the way a set of pixels in the high resolution image affects each low resolution pixel. In this paper we use an $\boldsymbol{H}_{l1,l2}$ representing a linear space-invariant blurring system with impulse response

$$h_{l1,l2}(u,v) = \begin{cases} \frac{1}{L^2} & u,v = -(L-1), \dots, 0 \\ 0 & \text{otherwise} \end{cases}. \tag{1}$$

Let now $\boldsymbol{D}_{l1}$ and $\boldsymbol{D}_{l2}$ be the 1-D downsampling matrices defined by

$$\boldsymbol{D}_{l1} = \boldsymbol{I}_{N_1} \otimes \boldsymbol{e}_l^t, \qquad \boldsymbol{D}_{l2} = \boldsymbol{I}_{N_2} \otimes \boldsymbol{e}_l^t, \tag{2}$$

where $\boldsymbol{I}_{N_i}$ is the $N_i \times N_i$ identity matrix, $\boldsymbol{e}_l$ is the $L \times 1$ unit vector whose nonzero element is in the l-th position and $\otimes$ denotes the Kronecker product operator.

Then for each sensor the discrete low-resolution observed image $\boldsymbol{g}_{l1,l2}$ can be written as

$$\boldsymbol{g}_{l1,l2} = \boldsymbol{D}_{l1,l2}\boldsymbol{H}_{l1,l2}\boldsymbol{B}_{l1,l2}\boldsymbol{f} + \boldsymbol{n}_{l1,l2} = \boldsymbol{W}_{l1,l2}\boldsymbol{B}_{l1,l2}\boldsymbol{f} + \boldsymbol{n}_{l1,l2}, \tag{3}$$

where $\boldsymbol{D}_{l1,l2} = \boldsymbol{D}_{l1} \otimes \boldsymbol{D}_{l2}$, denotes the $(N_1 \times N_2) \times (M_1 \times M_2)$ 2D downsampling matrix and $\boldsymbol{n}_{l1,l2}$ is modeled as independent white noise with variance $\beta_{l1,l2}^{-1}$. We denote by $\boldsymbol{g}$ the sum of the upsampled low resolution images, that is,

$$\boldsymbol{g} = \sum_{u,v \in \mathcal{I}} \boldsymbol{D}_{u,v}^t \boldsymbol{g}_{u,v}. \tag{4}$$

Note that the only, but important, difference between the model in equation 3 and the one used in [9] is the presence of additional blurring in the observation process.

3   Degradation and image models

From Eq. 3, the probability density function of $\boldsymbol{g}_{l1,l2}$, the $(l1,l2)$ low resolution image, with $\boldsymbol{f}$ the 'true' high resolution image, is given by

$$p(\boldsymbol{g}_{l1,l2}|\boldsymbol{f},\beta_{l1,l2}) \propto \frac{1}{Z(\beta_{l1,l2})} \exp\left[-\frac{\beta_{l1,l2}}{2} \parallel \boldsymbol{g}_{l1,l2} - \boldsymbol{W}_{l1,l2}\boldsymbol{B}_{l1,l2}\boldsymbol{f} \parallel^2 \right], \tag{5}$$

where $Z(\beta_{l1,l2}) = (2\pi/\beta_{l1,l2})^{(N_1 \times N_2)/2}$ and $\beta_{l1,l2}$ the inverse of the noise variance.

Since we have multiple low resolution images, the probability density function of $\boldsymbol{g}$ given $\boldsymbol{f}$ is

$$p(\boldsymbol{g}|\boldsymbol{f},\underline{\beta}) = \prod_{(l1,l2)\in\mathcal{I}} p(\boldsymbol{g}_{l1,l2}|\boldsymbol{f},\beta_{l1,l2})$$

$$\propto \frac{1}{Z_{noise}(\underline{\beta})} \exp\left[-\frac{1}{2} \sum_{(l1,l2)\in\mathcal{I}} \beta_{l1,l2} \parallel \boldsymbol{g}_{l1,l2} - \boldsymbol{W}_{l1,l2}\boldsymbol{B}_{l1,l2}\boldsymbol{f} \parallel^2 \right], \tag{6}$$

where $\underline{\beta} = (\beta_{l1,l2} | (l1, l2) \in \mathcal{I})$, and $Z_{noise}(\underline{\beta}) = \prod_{(l1,l2) \in \mathcal{I}} Z(\beta_{l1,l2})$.

As prior model for $\boldsymbol{f}$ we use a simultaneous autoregression (SAR) [15], that is

$$p(\boldsymbol{f}|\alpha) = \frac{1}{Z_{prior}(\alpha)} \exp\{-\frac{1}{2}\alpha\, \boldsymbol{f}^t \boldsymbol{C}^t \boldsymbol{C}\, \boldsymbol{f}\}, \tag{7}$$

where the parameter α measures the smoothness of the 'true' image, $Z_{prior}(\alpha) = (\prod_{i,j} \lambda_{ij}^2)^{-1/2} (2\pi/\alpha)^{(M_1 \times M_2)/2}$ and $\lambda_{ij} = 1 - 2\phi(\cos(2\pi i/M_1) + \cos(2\pi j/M_2))$, $i = 1, 2, \ldots, M_1$, $j = 1, 2, \ldots, M_2$ and $\boldsymbol{C}$ is the Laplacian operator.

4   Bayesian Analysis

Having defined the degradation and image models, the Bayesian analysis is performed to estimate the hyperparameters, α and $\underline{\beta}$, and the high-resolution image. In this paper we use the following two steps:

Step I: Estimation of the hyperparameters

$\hat{\alpha}$ and $\hat{\underline{\beta}} = (\hat{\beta}_{l1,l2} | (l1, l2) \in \mathcal{I})$ are first selected as

$$\hat{\alpha}, \hat{\underline{\beta}} = \arg\max_{\alpha, \underline{\beta}} \mathcal{L}_{\boldsymbol{g}}(\alpha, \underline{\beta}) = \arg\max_{\alpha, \underline{\beta}} \log p(\boldsymbol{g}|\alpha, \underline{\beta}), \tag{8}$$

where $p(\boldsymbol{g}|\alpha, \underline{\beta}) = \int_{\boldsymbol{f}} p(\boldsymbol{f}|\alpha)p(\boldsymbol{g}|\boldsymbol{f}, \underline{\beta})d\boldsymbol{f}$.

The solution to this equation is obtained with the EM-algorithm with $\mathcal{X}^t = (\boldsymbol{f}^t, \boldsymbol{g}^t)$ and $\mathcal{Y} = \boldsymbol{g} = [\boldsymbol{0}\ \ \boldsymbol{I}]^{\text{t}}\mathcal{X}$

Step II: Estimation of the high-resolution image

Once the hyperparameters have been estimated, the estimation of the high-resolution image, $\boldsymbol{f}_{(\hat{\alpha},\hat{\underline{\beta}})}$, is selected to minimize

$$\hat{\alpha}\parallel \boldsymbol{C}\boldsymbol{f} \parallel^2 + \sum_{(l1,l2)\in\mathcal{I}} \hat{\beta}_{l1,l2} \parallel \boldsymbol{g}_{l1,l2} - \boldsymbol{W}_{l1,l2}\boldsymbol{B}_{l1,l2}\boldsymbol{f} \parallel^2, \tag{9}$$

which results in

$$\boldsymbol{f}_{(\hat{\alpha},\hat{\underline{\beta}})} = \boldsymbol{Q}\left(\hat{\alpha}, \hat{\underline{\beta}}\right)^{-1} \sum_{(l1,l2)\in\mathcal{I}} \hat{\beta}_{l1,l2}\boldsymbol{B}_{l1,l2}^t\boldsymbol{W}_{l1,l2}^t\boldsymbol{g}_{l1,l2}, \tag{10}$$

where $\boldsymbol{Q}(\hat{\alpha}, \hat{\underline{\beta}}) = \hat{\alpha}\boldsymbol{C}^t\boldsymbol{C} + \sum_{(l1,l2)\in\mathcal{I}} \hat{\beta}_{l1,l2}\boldsymbol{B}_{l1,l2}^t\boldsymbol{W}_{l1,l2}^t\boldsymbol{W}_{l1,l2}\boldsymbol{B}_{l1,l2}$.

Note that the prior model in Eq. 7 plays an important role in the estimation of the high-resolution image and the hyperparameters. If we examine the matrix $\boldsymbol{Q}(\alpha, \underline{\beta})$ in Eq. 10 we note that when fewer than $L \times L$ low resolution observations are available or when the shifts in those low resolution images do not satisfy the conditions in [6] and [7] this matrix would not be invertible without the presence of $\boldsymbol{C}$. It is therefore important to examine the quality of the reconstruction and

also the accuracy of the estimated hyperparameters as a function of the number of low resolution observations, q. This is done experimentally, as described in detail in the next section.

It is important to note that the calculations involved in finding $\hat{\alpha}, \hat{\underline{\beta}}$ and $\boldsymbol{f}_{(\hat{\alpha},\hat{\underline{\beta}})}$ can be performed using the general framework for frequency domain multi-channel signal processing developed in [10].

5   Experimental Results

A number of experiments were performed with the proposed algorithm over a set of images to evaluate its behavior as a function of the number of available low resolution images. Results are presented in Fig.1.

The performance of the proposed algorithm was evaluated by measuring the peak signal-to-noise ratio (PSNR) defined as PSNR= $10 \times \log_{10}[M_1 \times M_2 \times 255^2 / \parallel \boldsymbol{f} - \boldsymbol{f}_{(\hat{\alpha},\hat{\underline{\beta}})} \parallel^2]$, where $\boldsymbol{f}$ and $\boldsymbol{f}_{(\hat{\alpha},\hat{\underline{\beta}})}$ are the original and estimated high resolution images, respectively.

According to Eq. 3 the high resolution image in Fig. 1a was blurred using a motion blur of length 10. Then, the integrating function in Eq. 1 was applied to the blurred image obtaining $\boldsymbol{u} = \boldsymbol{HBf}$. This high resolution image, $\boldsymbol{u}$, was downsampled with $L = 4$, thus obtaining a set of 16 low resolution images, $u_{l1,l2}(x,y) = u(L_1 x + l1, L_2 y + l2)$, $x, y = 0, \ldots, \frac{M_1}{L} - 1$, $l1, l2 = 0, \ldots, 3$. Gaussian noise was added to each low resolution image to obtain three sets of sixteen low resolution images, $\boldsymbol{g}_{l1,l2}$, with 20, 30 and 40dB SNR. The noise variances for the 30dB set of images are shown in Table 1.

Table 1. Noise variances for the low resolution image set with SNR of 30dB.

$\beta_{l1,l2}^{-1}$	0	1	2	3
0	3.44	3.44	3.45	3.46
1	3.44	3.44	3.45	3.46
2	3.42	3.42	3.43	3.43
3	3.42	4.42	3.43	3.43

In order to test the performance of the proposed algorithm we ran it on different sets of q randomly chosen low resolution images with $1 \leq q \leq 16$. For comparison purposes, Fig. 1b depicts the zero-order hold upsampled image of $\boldsymbol{g}_{0,0}$ for 30dB SNR (PSNR=13.68dB) while the bilinear interpolation of $\boldsymbol{g}_{0,0}$ is shown Fig. 1c (PSNR=14.22dB). The estimated high-resolution images using 1, 2, 4, 6 and 16 low resolution images are depicted in Fig. 1d–h, respectively. A visual inspection of the resulting images shows that the proposed method clearly outperforms zero-order hold and bilinear interpolation even when only one low resolution input image is used and the quality of the high resolution estimated

image increases with the number of images. Note also that the visual quality of the estimated high resolution images obtained using 6 and 16 low resolution input images (depicted in Fig. 1g and h, respectively) are almost indistinguishable, which means that the prior model assists in accurately recovering the high resolution image even when we have little information.
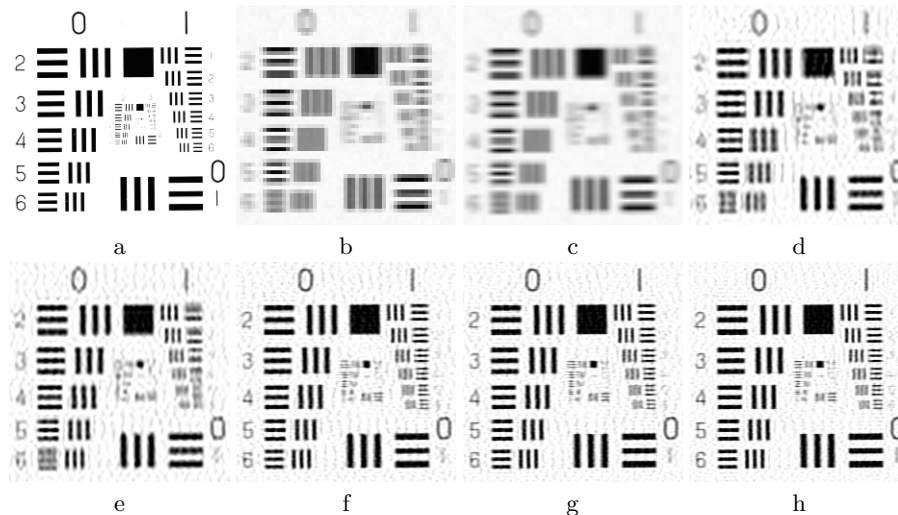

Fig. 1. (a) original image (b) zero order hold, (c) bilinear interpolation, (d)-(h) results with the proposed method using 1, 2, 4, 6, and 16 low resolution images.

The estimated noise parameters, $\hat{\underline{\beta}}$, using the proposed algorithm on the 30dB low resolution image set are shown in Table 2. From our experiments we conclude that the proposed method produces accurate estimations for all low resolution image noise variances especially when the number of input images is high. The proposed algorithm also provides good results when only a few input images are considered, even when only one low resolution input image is used.

PSNR evolution against the number of low resolution input images is shown in Fig. 2 for all the 20, 30 and 40dB low resolution images sets. Numerical results show that the proposed method provides a clear improvement even in the case when severe noise is present, although higher improvements are obtained as the noise decreases. Note that the proposed algorithm always outperforms bilinear interpolation even when only one image is used. For example, for the 30dB SNR low resolution images set, the PSNR for the reconstructed image using just one low resolution input image is equal to 16.74dB (see Fig. 1d), and it increases monotonically to 21.96dB with the number of images. Note also that most of the improvement is achieved when a low number of input images is used. This

Table 2. Estimated noise variances for the low resolution image set with SNR of 30 dB.

1 image	$\hat{\beta}_{0,0}^{-1}$
	4.46

2 images	$\hat{\beta}_{0,0}^{-1}$	$\hat{\beta}_{0,1}^{-1}$
	3.97	3.56

4 images	$\hat{\beta}_{0,0}^{-1}$	$\hat{\beta}_{0,1}^{-1}$	$\hat{\beta}_{2,2}^{-1}$	$\hat{\beta}_{3,1}^{-1}$
	3.38	3.50	4.72	4.54

6 images $\hat{\beta}_{l1,l2}^{-1}$	0	1	2	3
0	3.43	3.38	N/A	N/A
1	N/A	N/A	N/A	4.25
2	N/A	N/A	3.46	N/A
3	N/A	4.47	3.64	N/A

16 images $\hat{\beta}_{l1,l2}^{-1}$	0	1	2	3
0	3.48	3.41	3.52	3.44
1	3.41	3.43	3.49	3.43
2	3.34	3.25	3.53	3.41
3	3.48	3.56	3.44	3.42

makes clear the importance of the prior model in the information recovering process.

6 Conclusions

A new method to estimate a high resolution image from an incomplete set of blurred, undersampled low resolution images has been proposed. The approach followed can be used with any number of low resolution images from 1 to L^2, since the prior model accurately recovers the high resolution image even in the case where just one or very few input images are provided. The proposed method has been validated experimentally.

References

1. Aizawa, K., Komatsu, T., Saito, T.: A scheme for acquiring very high resolution images using multiple cameras. In: IEEE Conference on Audio, Speech and Signal Processing. Volume 3. (1992) 289–292
2. Stark, H., Oskoui, P.: High-resolution image recovery from image-plane arrays, using convex projections. Journal of the Optical Society A **6** (1989) 1715–1726
3. Borman, S., Stevenson, R.: Spatial resolution enhancement of low-resolution image sequences. A comprehensive review with directions for future research. Technical report, Laboratory for Image and Signal Analysis, University of Notre Dame (1998)
4. Molina, R., Vega, M., Abad, J., Katsaggelos, A.K.: Parameter estimation in Bayesian high-resolution image reconstruction with multisensors. Submitted to IEEE Trans. Image Processing (2002)
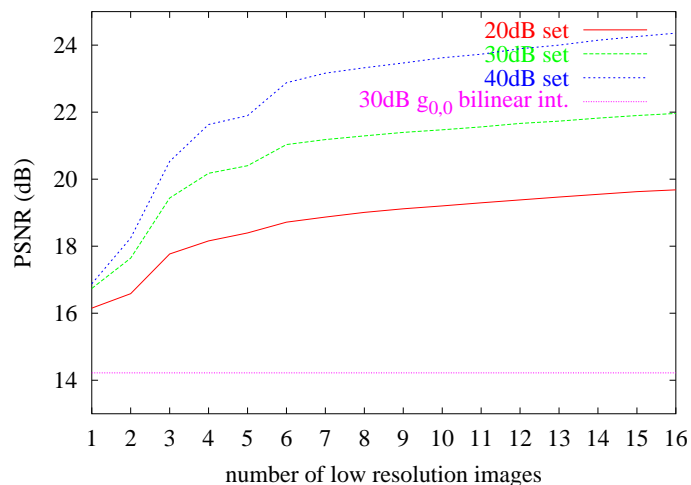
Fig. 2. PSNR evolution with the number of low resolution images.

5. Abad, J., Vega, M., Molina, R., Katsaggelos, A.K.: Parameter estimation in superresolution image reconstruction problems. To appear in Proc. International Conference on Acoustic, Speech and Signal Processing (2003)
6. Kim, S.P., Bose, N.K., Valenzuela, H.M.: Recursive reconstruction of high resolution image from noisy undersampled multiframes. IEEE Transactions on Acoustics, Speech and Signal Processing **38** (1990) 1013–1027
7. Elad, M., Feuer, A.: Restoration of a single super-resolution image from several blurred, noisy, and undersampled measured images. IEEE Trans. on Image Processing **6** (1997) 1646–1658
8. Baker, S., Kanade, T.: Limits on super-resolution and how to break them. IEEE Transactions on Pattern Analysis and Machine Intelligence **24** (2002)
9. Mateos, J., Molina, R., Katsaggelos, A.K.: Bayesian high resolution image reconstruction with incomplete multisensor low resolution systems. To appear in Proc. International Conference on Acoustic, Speech and Signal Processing (2003)
10. Katsaggelos, A.K., Lay, K.T., Galatsanos, N.P.: A general framework for frequency domain multi-channel signal processing. IEEE Trans. Image Processing, **2** (1993) 417–420
11. Bose, N.K., Boo, K.J.: High-resolution image reconstruction with multisensors. Int. Journ. Imaging Systems and Technology **9** (1998) 141–163
12. Molina, R., Nez, J., Cortijo, F., Mateos, J.: Image restoration in Astronomy. A Bayesian review. IEEE Signal Processing Magazine **18** (2001) 11–29
13. Ng, M.K., Yip, A.M.: A fast MAP algorithm for high-resolution image reconstruction with multisensors. Multidimensional Systems and Signal Processing **12** (2001) 143–164
14. Tom, B.C., Galatsanos, N.P., Katsaggelos, A.K.: Reconstruction of a high resolution image from multiple low resolution images. In Chaudhuri, S., ed.: Super-Resolution Imaging. Kluwer Academic Publishers (2001) 73–105
15. Ripley, B.D.: Spatial Statistics. John Wiley (1981)