APPLICATION OF SVD BASED SPATIAL FILTERING TO VIDEO SEQUENCES

Konstantinos Konstantinides and Gregory S. Yovanof

Hewlett-Packard Laboratories
P.O. Box 10490
Palo Alto, CA 94303

ABSTRACT

Video capture devices, such as CCD cameras, are a significant source of noise in image sequences. Preprocessing of video sequences with spatial filtering techniques usually improves their compressibility. In this paper we present a block-based, non-linear filtering algorithm based on the theories of SVD and compression-based filtering. A novel noise estimation algorithm allows us to operate on the input data without any prior knowledge of either the noise or signal characteristics. Experiments with real video sequences and an MPEG codec have shown that SVD based filters preserve edge details and can significantly improve nearly-lossless compression ratios by 15%.

1. INTRODUCTION

Several video processing and high-quality freeze-frame applications such as image analysis and video printing require that the input signal be processed by a spatial filter in order to improve its visual quality. Spatial filtering has also been applied to video data as a means of improving its compressibility. A significant source of contamination for image sequences is the additive noise introduced by the video capture device (photonic noise in the case of ccd cameras, or film-grain, when the video stream is scanned-in from analog data originally stored on film). Traditional linear filtering techniques [1], [2] have often been used to suppress additive noise in image sequences, mainly due to their simplicity. The major drawback of such techniques is that they tend to blur the original signal. Arguably, non-linear filters excel in achieving noise removal while preserving fine texture and edge details.[3]

One class of non-linear filtering techniques that has been successfully used in the past in various noise removal and image analysis applications is based on the Singular-Valued-Decomposition (SVD) method [4] [5]. Traditional techniques usually apply SVD to the whole image in a single, compute-intensive step, and they don't address the problem of distinguishing between significantly small and insignificantly large singular values. Furthermore, all filtering techniques proposed so far require some prior knowledge of the noise and image characteristics. In several applications this information may not be available and could be difficult to estimate from the input data.

In this paper we present a block-based, non-linear filtering technique based on SVD that employs an efficient method for estimating the noise power from the input data without the need for additional a priori information. The noise estimation method is based on a scheme recently introduced by Natarajan [6] for reducing additive noise from signals using data compression. Natarajan's scheme, stemming from the observation that noise is hard to compress, allows one to filter random noise using data compression. Based on that work, we show how one can derive a threshold to distinguish between significant and insignificant singular values. By applying SVD sequentially into blocks of the image we significantly reduce the computational requirements.

Experimentation with real video sequences has shown that the proposed SVD based filter in the case of high bit rate, nearly lossless, video coding can effectively suppress noise while preserving edge details, and can improve the compression ratio achieved by an MPEG codec by roughly 15%. These improvements are achieved without utilizing any prior knowledge of either the image or noise characteristics. To improve the overall performance, the proposed spatial filtering scheme can also be combined with motion-compensated temporal filtering [7].

2. SVD-BASED FILTERING

In the theory of SVD, any m x n real-valued matrix A can be decomposed as $A = U \Sigma V^T$, where U, V are orthogonal matrices and $\Sigma = diag(\alpha_1, \alpha_2, \ldots, \alpha_n)$ is a diagonal matrix. The elements of Σ are called the singular values of A. In theory, the rank of A is the number of its non-zero singular values. In practice, under an additive noise model, we observe a matrix

0-7803-2431-5/95 $4.00 © 1995 IEEE

$B = A + E$, where E is a noise perturbation matrix of full rank. In that case, the last singular values of B will be small, but not necessarily zero. Let $\beta_1 \geq \beta_2 \geq .... \geq \beta_{n-1} \geq \beta_n$ be the singular values of B listed in non-increasing order. We define the effective rank of B as r if $\beta_r \geq \epsilon_1 \geq \beta_{r+1}$, where $1 \leq r \leq min(m,n)$ and $\epsilon_1 = \|E\|_2$ is the 2-norm of E [8]. Let $B = U_B \Sigma_B V_B^T$ be the SVD of $B = A + E$. Given r, let $\Sigma'_B = diag(\beta_1, \ldots, \beta_r, 0, 0, ..., 0)$. For SVD based filtering, the filtered matrix is defined as $B' = U_B \Sigma'_B V_B^T$.

The main steps in our SVD based filtering technique [9] are: a) Divide an M x N image **B** into sub-blocks B_i. b) Perform SVD on each sub-block and evaluate its effective rank using a threshold ϵ. c) Set to zero the "non-significant" singular values, and replace B_i with B'_i, as defined before. By setting to zero the non-significant singular values, in effect we perform a lossy compression on each block. Based on recent work by Natarajan [6], the noise and the loss cancel out, and the reconstructed block is closer to the original block.

The efficiency of this algorithm depends on the accuracy of the estimate of the threshold ϵ. An estimate of this threshold can be obtained via the following calibration scheme: a) Apply the filtering algorithm on a representative noisy frame from the video sequence for various values of the threshold ϵ. b) Compress each filtered image using the lossless JPEG compression algorithm. c) Plot the compressed size as a function of ϵ. d) Let ϵ^* be the knee-point of the plot. Use $\epsilon = \epsilon^*$ as the threshold on the SVD filtering algorithm for the other frames. If s denotes the compressed size of a filtered image, the knee-point is defined as the point at which the second derivative $d^2s/dlog(\epsilon)^2$ is maximum.

3. EXPERIMENTAL RESULTS

As an example, we applied the SVD based spatial filter to the "football" video sequence from the MPEG test suite. This sequence consists of 150 CCIR-601 (720 x 480) interlaced frames. Following the noise estimation phase algorithm, we filtered three frames (No. 1, 25, and 75) for various values of the threshold ϵ and we compressed them using a lossless JPEG image compression algorithm. Fig. 1 shows a plot of the size of the compressed filtered images as a function of ϵ. Even though different frames have different compressed sizes, lossless compression seems to level-off after a threshold of $\epsilon = 12$ in all of them. Indeed, the second derivative of each of these plots

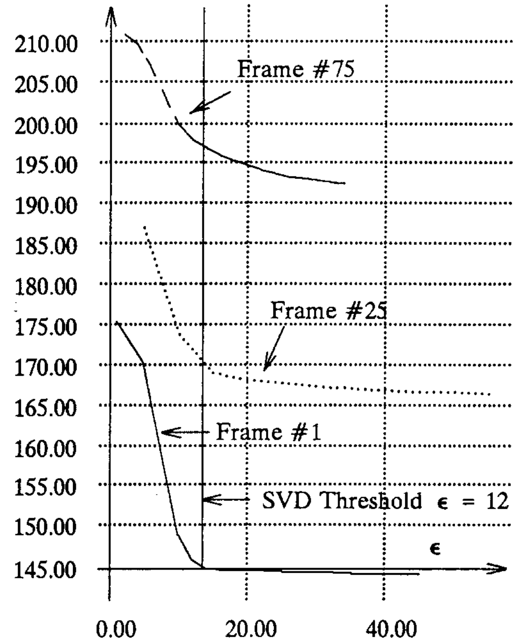

Fig. 1: Compressed size of filtered frames vs ϵ.

showed knee-point values in the range of 11 to 15. As a result, we selected $\epsilon^* = 12$. (We estimate that this corresponds to additive white noise of zero mean and a variance of roughly 2.25). Fig. 2 shows a plot of the size of the compressed filtered frame 75 as a function of ϵ, and its second derivative with respect to $log(\epsilon)$. The maximum of the second derivative occurs at $\epsilon = 12$. (Units on the vertical axis are normalized from 0 to 1).

For *nearly-lossless* MPEG-1 compression, both the original and the filtered (at $\epsilon = 12$) sequences were then compressed using the MPEG video encoder from the University of California at Berkeley. The bit rate, and thus the quality of the compressed video is controlled by three scale factors: the ISCALE, the PSCALE, and the BSCALE, which determine the effects of DCT quantization in I, P, and B frames respectively. In all of our experiments, we used a group of pictures size of 15 with the coding pattern of IBBPBBPBBPBBPBB. Motion estimation was performed in a [+15,-15] search range using logarithmic search for the P-frames and "CROSS2" search for the B-frames. Table 1 shows the improvements achieved in nearly-lossless compression using SVD-based filtering for various scale factors. At "low" bit rates (i.e. below 9.36 Mbits/s for 30 frames/s) the decompressed frames in both sequences were very blocky, and there was very little file size improvement between the original and the filtered compressed sequences. This was

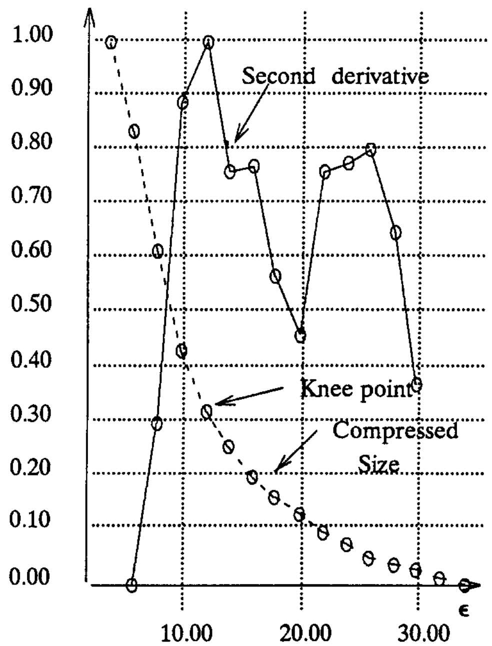

Fig. 2 : Compressed size of filtered frame #75 vs ε, and its second derivative.

Table 1
Improvement in output bit rates using nearly-lossless MPEG and SVD-based filtering.

Quant. Scales (I,P,B)	Compression Improvement	Bit Rate at 30 fps (Mbits/s)
1,1,1	16.01%	42.26
1,2,2	10.78%	24.72
1,3,3	6.50%	18.17
2,5,5	1.98%	11.47

expected, since in these cases the quantization error from the compression algorithm is far greater than the background noise we try to filter out. However, we measured 0.36 dB improvement in SNR in the filtered sequence. For bit rates between 18 Mbits/s and 42 Mbits/s, the quality of the compressed video was very good and we measured a size improvement of 5-16% between the original and the filtered compressed sequences. At these rates, the average SNR improvement was close to 1 dB.

To examine the effects of SVD-based filtering on a video sequence with higher background noise, the original "football" video sequence was corrupted with additive white noise of zero mean and $\sigma^2 = 32.58$ variance. This noise level corresponds to a peak SNR of 33 dB. After the noise estimation phase of the algorithm, we selected a new threshold $\epsilon^* = 35$. Using again the Berkeley MPEG coder, Table 2 shows the effects of SVD based filtering on the compressibility of the noise-corrupted video.

Table 2
Improvement in MPEG bit rates after SVD-based filtering on noise-corrupted "football" sequence.

Quant. Scales (I,P,B)	Compression Improvement	Bit Rate at 30 fps (Mbits/s)
1,2,2	54.72%	36.41
1,3,3	56.89%	26.31
2,5,5	57.88%	15.21

Table 2 shows that the output bit rate increases with the background noise level. However, SVD-based filtering yields now close to 55% improvement in MPEG-1 output bit rate. For comparison, we also tested the performance of a 3x3 median filter. As shown in Figs. 4-6, the median filter causes image blurring and thus it is not suitable for our application (notice the blurring of the letters in Fig. 5). In contrast, the SVD filter (see Fig. 6) preserves edge details and overall picture fidelity.

Fig. 3: Original frame #75 (one field)

4. CONCLUSIONS

We presented a novel noise filtering algorithm for video sequences based on the theories of SVD and data compression. Simulation results show that the technique can effectively filter noisy images with no

Fig. 4: Frame #75 processed with a 3x3 median filter.

Fig. 5: Frame #75 processed with the SVD filter and threshold $\epsilon^* = 12$.

prior knowledge of either the signal or the noise characteristics. This results in increased compressibility when the filtered data is subsequently processed by a video compression scheme like MPEG. Experiments have shown a 16% improvement in the compression ratio achieved by nearly-lossless MPEG or equivalently, a visual quality improvement of 1 dB at the same rate. This scheme can be used in conjunction with traditional motion-compensated temporal filtering techniques to further improve the overall performance of a high-quality video processing system.

Acknowledgment

Thanks to Balas K. Natarajan and Vasudev Bhaskaran for discussions, comments, and suggestions.

REFERENCES

1. A. Rosenfeld, A.C. Kak, "*Digital Picture Processing,*" Academic Press, 2nd edition, 1982.
2. M.K. Ozkan, A.T. Erdem, M.I. Sezan, A.M. Tekalp, "*Efficient multiframe Wiener restoration of blurred and noisy image sequences,*" IEEE Trans. Image Processing, vol. 1 (4), pp. 453-476, Oct. 1992.
3. T. Viero and Y. Neuvo, "*3-D Median structures for image sequence filtering and coding,*" in Motion Analysis and Image Sequence Processing, M.I. Sezan and R.L. Lagendijk, eds, Kluwer Academic Publishers, Norwell, MA 1993.
4. H.C. Andrews and C.L. Patterson, "Singular value decompositions and digital image processing," IEEE Trans. on ASSP, Vol. ASSP-24, No. 1, Feb. 1976, pp. 26-53.
5. Hsien-Che Lee et al, "*Digital Image Noise Suppression method Using SVD Block Transform,*" US Patent No. 5010504, Apr. 23 1991.
6. B.K. Natarajan, "Filtering random noise via data compression," Proc. IEEE Data Compression Conf., Snowbird, Utah, 1993, pp. 60-69.
7. B.K. Natarajan and V. Bhaskaran, "Effective nearly lossless compression of digital video sequences via motion-compensated filtering," Proc. IEEE ICASSP-94, pp. V-441-443, Adelaide, 1994.
8. K. Konstantinides, and K. Yao, "Statistical analysis of effective singular values in matrix rank determination," IEEE Trans. on ASSP, Vol. 36, No. 5, May 1988, pp. 757-763.
9. K. Konstantinides and G. Yovanof, "Improved compression performance using SVD-based filters for still images," Proc. IS&T/SPIE Symp. on Electronic Imaging: Science and Technology, Still-Image Compression, Vol. 2418, Feb. 1995, San Jose, CA.