TREIZIÈME COLLOQUE GRETSI - JUAN-LES-PINS DU 16 AU 20 SEPTEMBRE 1991

Vector Order Statistics Filtering of Colour Images

Sunil Sanwalka and A. N. Venetsanopoulos

Dept. of Electrical Engineering, University of Toronto
Toronto, Ontario, CANADA M5S1A4

RÉSUMÉ

Une nouvelle approche pour le filtrage d'images en couleurs est proposée, basée sur l'ordre statistique de vecteurs. Les images en couleurs peuvent être considérées comme des champs de vecteurs bidimensionnels, dans lesquels chaque composante vectorielle correspond à un canal couleur. L'utilisation des techniques d'ordonnement vectoriel nous permet d'étendre les concepts des filtres d'ordre statistique monochromes aux traitements multicanaux. L'avantage de cette méthode sur les approches traditionnelles basées sur des traitements scalaires indépendants des composantes couleur est que cette méthode utilise les corrélations du signal et du bruit entre les différents canaux afin d'améliorer les performances de filtrage.

Nous présentons deux filtres d'ordre statistique de vecteurs: le filtre médian et le filtre de moyenne trimmé en α. Ces filtres sont testés avec différents bruits multicanaux dont des bruits gaussien, exponentiel et impulsif. Les performances de ces filtres sont comparées à celles obtenues avec une approche scalaire et des améliorations sont constatées concernant la réduction des bruits et la préservation des contours.

ABSTRACT

A new approach based on vector order statistics is proposed for the filtering of colour images. A colour image can be considered as a two-dimensional vector field, where each vector component corresponds to one of the three colour channels. The use of vector or multivariate ordering techniques allows us to extend the concepts of monochrome order statistics filters to the multichannel or colour case. An advantage of this method over the current approach of applying scalar processing to each of the channels independently, is that it inherently utilizes the interchannel signal and noise correlation to improve the image filtering capablity.

We present two different vector order statistic filters, the median filter and the α-trimmed mean, a type of L-filter. These filters are evaluated using a variety of multichannel noise processes, including guassian, exponetial and impulsive noise. Their performance is compared to the scalar approach and improvements are shown in the noise reduction and edge preservation properties of the filters.

1. INTRODUCTION

The use of colour information has become very important in the fields of image processing and computer vision. Previous work on the processing of colour images involved the application of the techniques developed for monochrome imagery to the different spectral components of the colour image separately. Some good results have been obtained using separable processing, but better results can be obtained by processing the colour image using vector methods and taking advantage of the correlations between the different image components in the algorithms. This type of processing was impractical in the past because of the large storage and processing time requirements, but now becomes feasible due to the availability of increasing computing power and memory at low costs. Recent research in colour image coding, segmentation, and restoration has focused on trying to realize these advantages.[1-3]

In this paper we propose filtering algorithms based on robust estimation theory and order statistics, such as the median filter and L-filters. These will be examined due to their good behaviour in the presence of additive white Gaussian noise, long tailed additive noise, and impulsive noise.[4] It has been found in monochrome digital image processing that order statistics filters are especially suitable because of the failure of classical linear filters in dealing with the nonstationarity of the images and noise. Also, the most important reason for the use of order statistic filters is that they have excellent edge preservation abilities and it is well known that edge information is very important in image perception.[5]

Each colour point or pixel in an image can be described as a vector, in which the components are the three primary colours: red, green and blue. Thus a colour image can be represented as a two dimensional field of vectors. Processing done in a scalar manner ignores correlations between the colours, whereas processing on the vector image inherently includes these correlations and takes advantage of them.

In the proceeding sections of this paper, we will look at a number of topics. Initially, we will discuss scalar and vector median filters. We will then extend these ideas to general L filters by defining a vector ordering scheme. Some of the results of tests of these filters performed on colour images with noise will then be presented with comparisons to scalar processing.

2. Median Filtering

The median is widely used in statistics, often due to the fact that it is the maximum likelihood estimator of location for random variables with the biexponential distribution.

$$f(x) = \frac{1}{\lambda} e^{-\lambda|x-\alpha|} \tag{1}$$

It has found much use in the filtering of images due mainly to its edge preservation properties and its ability to filter out impulsive noise.[6] These noise impulses can be brought about by the recording media, the transmission channel, or the acquisition unit and cannot be removed effectively by conventional linear filters.

Given n observations x_i , $i=1, \ldots, n$, the median of the set of x_i's denoted by x_{med} is given by

$$x_{med} = \begin{cases} x_{(\nu+1)} & n = 2\nu+1 \\ \tfrac{1}{2}(x_{(\nu)}+x_{(\nu+1)}) & n = 2\nu \end{cases} \tag{2}$$

where the $x_{(i)}$ are the ordered sequence of observations x_i

i.e. $\quad x_{(1)} \le x_{(2)} \le x_{(3)} \le \cdots \le x_{(n)}$

The $x_{(i)}$ is also called the i^{th} *order statistic* .

An alternative definition of the median is given as the x_i that minimizes the L_1 error norm to the set of x_i.

$$f(x) = \sum_{i=1}^{n} | x_i - x | \tag{3}$$

$$f(x_{med}) \le f(x_i) \;\; \forall x_i \quad , x_{med} \in \{x_i \, , \, i=1,\ldots,n\}$$

This definition of x_{med} is not exactly equivalent to the previous one, because when $n=2v$ both $x_{(v)}$ and $x_{(v+1)}$ will minimize the L_1 error norm. This difference is unimportant in image filtering applications because n is generally odd.

Now to extend the definition of the scalar median to the vector or multivariate median, the second definition presented above will be used as there is no natural way to absolutely order multi-dimensional data. Ordering of multidimensional data will be further considered in the next section. The multivariate median of n observations of $\mathbf{x}_i$, $i=1,\ldots,n$, where the $\mathbf{x}_i$ are k dimensional vectors $[x_{i_1},x_{i_2},x_{i_3},\ldots,x_{i_k}]^T$, is given by

$$f(\mathbf{x})=\sum_{i=1}^{n}|\mathbf{x}_i-\mathbf{x}| \tag{4}$$

$$f(\mathbf{x}_{med})\leq f(\mathbf{x}_i)\ \forall\, \mathbf{x}_i\quad,\ \mathbf{x}_{med}\in\{\mathbf{x}_i\ ,\ i=1,...,n\}$$

the choice of the distance metric is arbitrary and we will choose the Euclidean distance.

3. Ordering of Multivariate Data

Order statistics have played an important role in the robust analysis of data contaminated with outlying observations and it is this robustness that also makes them one of the most useful families of image filters. For univariate or one dimensional data, the definition of order statistics is well defined and has been thoroughly studied.[7] Let the random variables $X_1,X_2,\ldots,X_n$ be arranged in ascending order of magnitude and written as

$$X_{(1)}\leq X_{(2)}\leq\cdots\leq X_{(n)} \tag{5}$$

Then $X_{(i)}$ is the so-called i^{th} *order statistic*. Important order statistics are the minimum $X_{(1)}$, the maximum $X_{(n)}$, and the median $X_{(n/2)}$. Generally, the X_i are independent identically distributed (iid) random variables, having cdf $F(x)$. Thus the cdf $F_r(x)$ of the r^{th} order statistic $X_{(r)}$ is given by

$$F_r(x)=\sum_{i=r}^{n}\binom{n}{i}F^i(x)[1-F(x)]^{n-i} \tag{6}$$

Additional details of the statistical analysis of one dimensional order statistics can be found in [7].

Now to extend these concepts to multivariate data, let us denote $\mathbf{X}$ as a p-dimensional multivariate or multiple random variable,

$$\mathbf{X}=[X_1,X_2,\ldots,X_p]^T$$

where $X_1,\ldots,X_p$ are random variables. Now, let $\mathbf{x}_1,\ldots,\mathbf{x}_n$ be n random samples from the multivariate $\mathbf{X}$. Each of the $\mathbf{x}_i$ are p-dimensional vectors of observations.

$$\mathbf{x}_i=[x_{i_1},x_{i_2},\ldots,x_{i_p}]^T$$

We would like to arrange the $\mathbf{x}_1,\ldots,\mathbf{x}_n$ in some sort of order. The notion of data ordering which is very natural in the one dimensional case does not extend in a straightforward way to multivariate data. There is no unambiguous, universally agreeable total ordering of n samples $\mathbf{x}_1,\mathbf{x}_2,\ldots,\mathbf{x}_n$, but there have been many ways proposed to attempt multivariate ordering. These so called sub-ordering principles can be separated into 4 categories; marginal ordering, reduced or aggregate ordering, partial ordering, and conditional ordering.

In marginal or M-ordering, the multivariate samples are ordered along each one of the p-dimensions independently. For colour signals this is equivalent to the separable method where each one of the colours is processed independently. The i^{th} marginal order statistic is the vector

$$\mathbf{x}_{(i)}=[x_{(i)_1},x_{(i)_2},\ldots,x_{(i)_p}]^T.$$

where $x_{(i)_r}$ is the i^{th} largest element in the r^{th} dimension or channel. The marginal order statistic $\mathbf{x}_{(i)}$ may not correspond to any of the original samples $\mathbf{x}_1,\ldots,\mathbf{x}_n$ as it does in one dimension.

Reduced or R-ordering is also called aggregate ordering. With this type of ordering each multivariate observation $\mathbf{x}_i$ is reduced to a single value d_i by means of some combination of the component values. The metric that is employed is frequently the generalized distance to some point α.

$$d_i=(\mathbf{x}_i-\alpha)^T\Gamma^{-1}(\mathbf{x}_i-\alpha) \tag{7}$$

The samples are then arranged in ascending order of magnitude of the associated metric values d_i.

In partial or P-ordering the objective is to partition the data into groups or sets of samples, such that the groups can be distinguished with respect to order, rank, or extremeness. This type of ordering can be accomplished by using the notion of convex hulls, but the problem with this is that determination of the convex hull is difficult to do in more than 2 dimensions. The other drawback is that there is no ordering within the groups and thus it is not easily expressed in analytical terms. These properties make P-ordering infeasible for implementation in digital image processing.

In conditional or C-ordering the multivariate samples are ordered conditional on one of the marginal sets of observations. Thus the i^{th} order statistic would be

$$\mathbf{x}_{(i)}=[x_{(i)_1}\ ,\ x_{[i]_2}\ ,\ \cdots\ ,\ x_{[i]_p}]^T$$

where $x_{(i)_1}$ is the marginal order statistic of the first dimension, and $x_{[i]_j}$,$j=2,..,p$ are the quasi-ordered samples in the other dimensions conditional on the i^{th} ordered sample in the first dimension. The marginal samples used for the ordering may be the original ones, or those derived from some preliminary co-ordinate transformation.

The ordering scheme that we propose to use is a variation of R-ordering, where the distance metric used is the L_1 error norm to the set of $\mathbf{x}_i$ or equivalently the aggregate distance (d) of each point from all the other points.

$$d_i=\sum_{k=1}^{n}|\mathbf{x}_k-\mathbf{x}_i| \tag{8}$$

The d_i ,$i=1,..,n$ are then arranged in order of magnitude and the associated vectors will be correspondingly ordered.

$$d_{(1)}\leq d_{(2)}\leq\cdots\leq d_{(n)}$$

$$\mathbf{x}_{(1)}\leq\mathbf{x}_{(2)}\leq\cdots\leq\mathbf{x}_{(n)}$$

Here the $\mathbf{x}_{(i)}$ have a one to one relationship with the samples $\mathbf{x}_i$ unlike marginal ordering. This ordering scheme was chosen for a number of reasons. The first is that by using this distance metric, $\mathbf{x}_{(1)}$ is the vector median of the data samples. Secondly, large values for the aggregate distance (8) give the most natural and accurate description of outliers or extreme values.[8] Thus we get some sense of an absolute order for the data vectors from the median to the most outlying value without any need for apriori information about the signal and noise distributions. The ordering can be entirely based on the data and is independent of an origin or fixed point in space.

4. ORDER STATISTIC FILTERS

L-filters, also sometimes called order statistic filters, are based on L-estimators. An L-estimator has the following definition for one dimensional data:

$$T_n=\sum_{i=1}^{n}a_ix_{(i)}\qquad,\text{where}\ \sum_{i=1}^{n}a_i=1 \tag{9}$$

where the a_i's are a set of weights that define the estimator's performance and are normalized. This can now be applied to image filtering as a local area operation. The definition of the 2 dimensional L-filter is

$$y_{ij}=\sum_{k=1}^{n}a_kx_{(k)} \tag{10}$$

where y_{ij} is the output pixel, $x_{(k)}$, $k=1,..,n$ are the ordered samples contained in the input window centered at pixel (i,j), and a_k, $k=1,..,n$ are the filter coefficients that still must adhere to (10). For an appropriate choice of the filter coefficients, the L-filter can be defined as the median and the α-trimmed mean, respectively

$$a_k = \begin{cases} 1 & k=\nu \qquad n=2\nu+1 \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

$$a_k = \begin{cases} \dfrac{1}{n(1-2\alpha)} & k=\alpha n+1, \ldots, n-\alpha n \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

Similarly, we can define the multi-dimensional L-filter, using the p-dimensional vector image definition, as:

$$\mathbf{y}_{ij} = \sum_{k=1}^{n} \mathbf{a}_k^T \mathbf{x}_{(k)} \tag{13}$$

where now $\mathbf{y}_{ij}$ is the output pixel vector, $\mathbf{x}_{(k)} = [x_{(k)_1}, \ldots, x_{(k)_p}]^T$, $k=1,..,n$ are the ordered vectors contained in the input window centered at (i,j), and $\mathbf{a}_k = [a_{k_1}, \ldots, a_{k_p}]^T, k=1,..,n$ are the filter coefficient vectors. The $\mathbf{a}_k$ must now satisfy

$$\sum_{i=1}^{n} \mathbf{a}_i = \mathbf{e} = [1,1,...,1]^T \tag{14}$$

We can now define the α-trimmed mean for p-dimensional vector images

$$\mathbf{y}_{ij} = \sum_{k=1}^{k=n(1-2\alpha)} \mathbf{a}_k^T \mathbf{x}_{(k)} \tag{15}$$

$$\text{where } \mathbf{a}_k^T = \left[\frac{1}{n(1-2\alpha)}, \frac{1}{n(1-2\alpha)}, \ldots, \frac{1}{n(1-2\alpha)} \right]^T$$

The α-trimmed mean as above defined will reject 2α% of the outlying samples and the arithmetic mean of the remaining samples will be calculated. This filter will have good performance in the presence of Gaussian noise and will still preserve edges well.

5. Multichannel Noise Processes

To test these proposed filters and evaluate their performance, we need to model some noise processes. Since the original signal is in our case multichannel, the noise must also be multichannel in nature. The problem of modeling noise in real digital colour images (i.e. digital television , colour digital photography, satellite photography, etc) is currently unexplored and beyond the focus of this paper, thus we will consider general classes of noise probability distributions, limiting ourselves to additive noise. To fully examine the performance of the filters, noise processes with short tailed , long tailed , and extremely long tailed distributions would be useful.

A useful short tailed noise distribution is the multivariate gaussian. Its pdf for 3 dimensions is

$$f(\mathbf{x}) = \frac{1}{(2\pi)^{3/2} |\Sigma|^{1/2}} \exp\{-\frac{1}{2}(\mathbf{x}-\eta)^T \Sigma^{-1} (\mathbf{x}-\eta)\} \tag{16}$$

where $\eta = [\eta_1, \eta_2, \eta_3]^T$ is the vector of means, and Σ is the covariance matrix. If the distribution is zero mean and uncorrelated with equal variance, i.e. $\eta = [0,0,0]^T$ and $\Sigma = \mathbf{I}$, then (16) reduces to

$$f(x_1, x_2, x_3) = \frac{1}{(2\pi)^{3/2}} e^{-1/2(x_1^2+x_2^2+x_3^2)} \tag{17}$$

For long tailed noise we will use a distribution that is a multivariate extension of the scalar biexponential distribution given in (1). This 3 dimensional exponential distribution can be written as

$$f(\mathbf{x}) = K e^{-\lambda \|\mathbf{x}\|} \ , \ \lambda \geq 0 \tag{18}$$

or

$$f(x_1, x_2, x_3) = \frac{\lambda^3}{8\pi} e^{-\lambda \sqrt{x_1^2+x_2^2+x_3^2}} \tag{19}$$

Finally, for extremely long tailed noise we will use the multivariate gaussian noise as defined above (17) contaminated with α% positive and negative impulses. This distribution can be written separably as

$$f(x_1, x_2, x_3) = f(x_1) f(x_2) f(x_3) \tag{20}$$

$$f(x_i) = (1-\alpha) e^{-x^2/2} + \frac{\alpha}{2}\delta(x+255) + \frac{\alpha}{2}\delta(x-255)$$

where $\delta(\cdot)$ is the delta function.

6. Simulation Results

A set of experiments was performed to assess the performance of the vector order statistic filters. Two different filters were tested; the vector median (VM) and the vector α-trimmed mean (VαTM) filter. These filters' ability to remove additive noise from colour images was assessed. Six noise distributions were used: gaussian noise with variance 30 and 50 , exponetial noise also with variance 30 and 50 and α contaminated gaussian noise with variance 30 and α = 5%, 10%. To evaluate the performance of these filters the mean square error (MSE) is calculated and is tabulated in Tables 1 and 2 . In these tables comparison of the vector filter is made with a similar filter applied in the separable or scalar manner. The percentage change in the MSE is given, where a negative percentage means that there was a reduction of the MSE by using the vector filter. This percentage change is the average of the change observed for a number of different colour images.

Percentage Change in MSE for VM versus SM		
mse +/- %	3×3 window	5×5 window
gaussian, variance=30	18.2	10.6
gaussian, variance=50	21.6	17.8
exponential, variance=30	1.4	-2.0
exponential, variance=50	4.1	1.0
5%-contaminated gauss	20.0	12.3
10%-contaminated gauss	21.7	14.2

Table 1

Percentage Change in MSE for VαTM versus SαTM		
mse +/- %	3×3 window α=11.1%	5×5 window α=12%
gaussian, variance=30	1.2	-8.0
gaussian, variance=50	7.4	-.5
exponential, variance=30	-9.2	-10.9
exponential, variance=50	-8.0	-9.3
5%-contaminated gauss	-3.1	-6.7
10%-contaminated gauss	-4.1	-5.4

Table 2

We present these quantitative results here, but we have found that these results do not correspond well to the subjective evaluation of the image quality. This is because the MSE does not properly take into account the nonlinearity of visual perception. This is especially true for colour images where the perceptual colour space is highly nonlinear.[9]

Table 1 shows the results for the vector median filter (VM). We can see that almost all the results are positive, i.e. there is an increase in the MSE for the vector median over the scalar median. There is even an increase for the case of exponential noise, even though the vector median is the maximum likelihood estimator for that distribution. From observing the filtered images we see that in the flat areas of the image, the scalar median (SM) provides more smoothing, but in the edge areas the vector median better preserves the sharpness of the edges. Overall, the performance of the two filters in many cases is visually indifferent unless examined closely.

From table 2 we can see that the MSE is reduced in most cases for the VαTM filter with the highest reductions for exponential noise. One of the most important observations that can be made from viewing the filtered colour images is that the vector filters are superior in preserving the edges and texture in the image. We can see this clearly in Figure 1 below, which shows a surface plot for one channel of a colour step image, that has been filtered by the VαTM filter and the SαTM filter respectively.

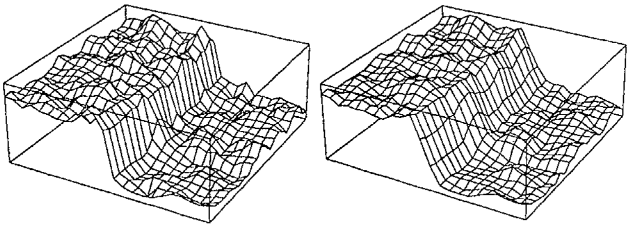

Figure 1

Also very noticeable in both the exponential and α-contaminated gaussian noise cases is the presence of impulses or spikes in the scalar filtered images. From figures 2, 3 we can see that the vector filter removes these impulses more effectively. The VαTM also maintains the colour of the original image more truly, whereas in the scalar filter colour artifacts are created such as green dots at edges in the image.

Lastly we have also applied these filters to an noisy image that has been digitized off video tape. The results of filtering this image are shown in figures 4, 5. We can see that the vector filter removes the streaks very effectively and also that the edges are sharper and the colour more intense.

7. Conclusions

A method for ordering multivariate or vector data has been proposed that allows for the extension of scalar order statistics to m-dimensions. These vector order statistics were applied to the filtering of additive noise from colour images. Two vector L-filters were tested and compared to similar scalar filters applied to the image components separately. These filters were assessed using multichannel noise distributions with varied characteristics. It has been found that the vector median filter exhibits very similar characteristics to the scalar median and that for many noise distributions the output of the two filters is almost visually indistinguishable. In contrast, the vector α-trimmed mean filter has shown significant improvements in edge preservation and impulsive noise removal.

8. References

[1] N. P. Galatsanos and R. T. Chin, "Digital restoration of multichannel images", IEEE Trans. Acoust. Speech Signal Proc. vol ASSP-37(3),415-421 (1989).

[2] R. Nevatia, "A color edge detector and its use in scene segmentation", IEEE Trans. Systems, Man, and Cybernetics. vol SMC-7,no. 11,820-826 (1977)

[3] O. Faugeras, "Digital color image processing within the framework of a human visual model", IEEE Trans. Acoust. Speech Signal Proc. vol ASSP-27, no. 4, 380-393 (1979)

[4] I. Pitas and A. N. Venatsanopoulos, Nonlinear digital filters: principles and applications. Kluwer Academic, 1990.

[5] W. K. Pratt, Digital Image Processing. John Wiley & Sons, 1978.

[6] J. Astola, et al., "Median Type Filters for Colour Signals" , IEEE International Conference on Circuits and Systems, Helsinki 1988, pgs 1753-1756

[7] H. A. David, Order Statistics. J. Wiley , 1980.

[8] V. Barnett, "The Ordering of Multivariate Data", J. R. Statistical Society A, 1976, Volume 139, Part 3, pgs 318-343.

[9] A. K. Jain, Fundamentals of Digital Image Processing. Prentice Hall, 1989.

Figure 2: SαTM 3×3 window, 10%-contaminated guassian noise

Figure 3: VαTM 3×3 window, 10%-contaminated guassian noise

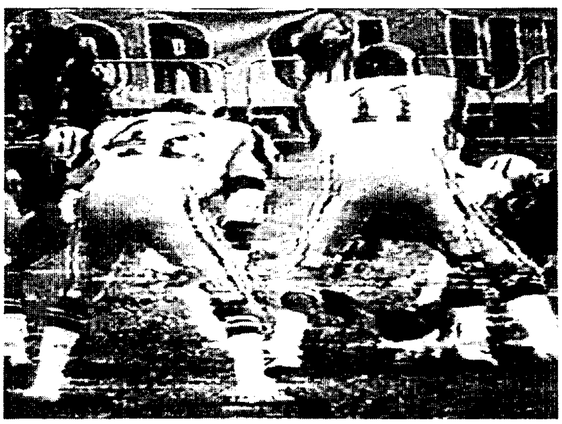

Figure 4: SαTM 3×3 window, corrupted video image

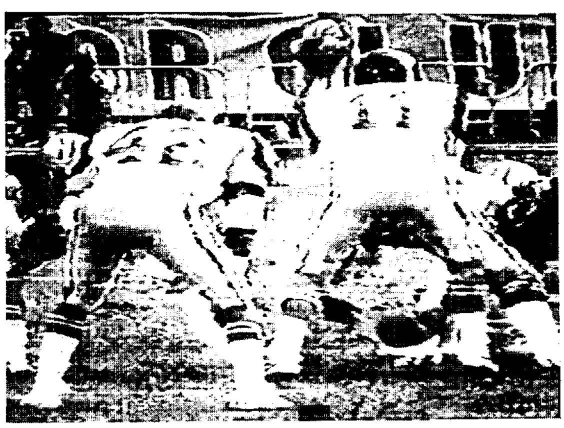

Figure 5: VαTM 3×3 window, corrupted video image