International Journal of Information & Computation Technology.
ISSN 0974-2239 Volume 4, Number 1 (2014), pp. 29-38

http://www. irphouse.com

Application of Local Segmentation in Denoising of Digital Images

Vipin Mishra

Administrator Shri Swami Bhumanand College of Nursing,
Ranipur Jhal, Jwalapur, Haridwar
Email: mrvipinji@gmail.com

Abstract

A unifying philosophy for carrying out low level image processing called "local segmentation" is presented. Local segmentation provides a way to examine and understand existing algorithms, as well as a paradigm for creating new ones. Local segmentation may be applied to a range of important image processing tasks. Using a traditional segmentation technique in intensity thresholding and a simple model selection criterion, the new FUELS denoising algorithm is shown to be highly competitive with state-of-the-art algorithms on a range of images. In an effort to improve the local segmentation, the minimum message length information theoretic criterion for model selection (MML) is used to select between models having different structure and complexity. This leads to further improvements in denoising performance. Both FUELS and the MML variants thereof require no special user supplied parameters, but instead learn from the image itself. It is believed that image processing in general could benefit greatly from the application of the local segmentation methodology.

Keywords: Computer, Information Technology, Digital Image Processing, Images, Digital Images, Image Acquisition, image sensing.

1. Introduction

The local segmentation principle may be used to develop a variety of low level image processing algorithms. In this research it will be applied to the specific problem of denoising greyscale images contaminated by additive noise. The best image denoising techniques attempt to preserve image structure as well as remove noise. This problem domain is well suited to demonstrating the utility of the local segmentation philosophy.

A multilevel thresholding technique will be used to segment the local region encompassing each pixel. The number of segments will be determined automatically by ensuring that the segment intensities are well separated. The separation criterion will adapt to the level of additive noise, which may be supplied by the user or estimated automatically by the algorithm. The resulting segmentation provides a local approximation to the underlying pixel values, which may be used to denoise the image. The denoising algorithm presented is called FUELS, which stands for “filtering using explicit local segmentation”. FUELS differs from existing local denoising methods in various ways. The local segmentation process clearly decides which pixels belong together, and does so democratically, without using the centre pixel as a reference value. If the computed local approximation suggests changing a pixel's value by too much, the approximation is ignored, and the pixel is passed through unmodified. The fact that each local approximation overlaps with its neighbour means that there are multiple estimates for the true value of each pixel. By combining these overlapping estimates, denoising performance is further increased.

FUELS will be shown to outperform state-of-the-art algorithms on a variety of greyscale images contaminated by additive noise. FUELS' worst case error behaviour will be shown to be proportional to the noise level, suggesting that it is quite adept at identifying structure in the image. The denoised images produced by FUELS will be seen to preserve more image structure than algorithms such as SUSAN and GIWS.

2. Global image models

An image model is a mathematical description of the processes affecting the final pixel values in an image. These processes may include atmospheric effects in the scene, noise in the capture device, and quantization of pixel values. It may be possible to construct a model which accounts for all these steps, but it would probably consist of many difficult to determine parameters. Often a simpler model can incorporate the main factors and still achieve good results.

Modeling the steps from acquisition to a final digital image is sufficient, but not necessary, for successful image processing. In many cases the full acquisition history may not even be known. In any case, some assumptions or prior beliefs regarding the properties of the original scene are also required. The characteristics of the light source and the object surfaces will influence greatly how we expect the image pixel values to be distributed spatially and spectrally. The number, distance and size of objects in the scene will also affect the proportion, scale and sharpness of edges in the image.

For example, consider the artificial images in Figure 2.1 [1]. The first is greyscale light intensity image of a simple object. The second image is also of the same object, except that the intensity of each pixel represents the *distance* from the camera to each point in the scene.

This is an example of a *range image*, for which the “light source” is actually a distance measurement device. It has been shown that pixel intensities in a range image usually vary in a spatially linear manner, due to the polyhedral nature of most objects [2]. However this may not be an appropriate assumption to make when analyzing low resolution fingerprint images in a criminal database. There one would expect many ridges and high contrast edges. This thesis focuses on greyscale light intensity images.

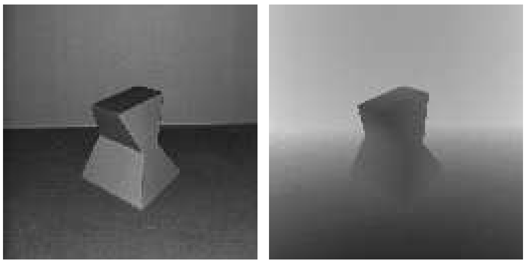

Figure 2.1: Two images of the same scene: (a) light intensity; (b) range, darker is closer.

2.1 The facet model

For greyscale image data, variations on the *facet model* [3] are most commonly used. The facet model assumes that a digital image is a discrete approximation of a piece-wise continuous intensity surface. Each piece is disjoint and is called a facet or segment. The underlying pixel values in each facet are assumed to follow a two dimensional polynomial function. This polynomial fit can be of any order. Figure 2.2 provides an example of a one dimensional discrete signal approximated by increasingly higher order polynomial curves.

The simplest facet approximation is a zeroth order polynomial of the form $f(x, y) = a$, where a is a constant. Each pixel in a constant facet is assumed to have the same value. For greyscale data this would be a scalar representing the intensity, but for colour data it would be an RGB vector. An image containing constant facets referred to as being *piece-wise constant*.

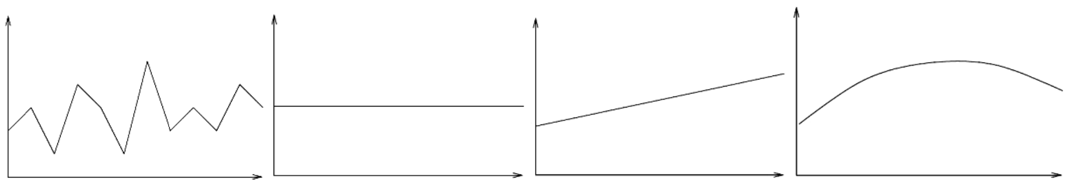

Figure 2.2: One dimensional polynomial approximation: (a) the original signal; (b) constant; (c) linear; (d) quadratic.

Piece-wise constant image models are commonly used in image processing. They only have one parameter to estimate, and are simple to manipulate.

First order polynomial approximation in two dimensions has the mathematical form of a plane, namely $f(x,y) = a + bx + cy$. An image containing facets of this type is *piece-wise planar*. Pixel values in a planar facet are linearly dependent on their position within the facet. Planar facets are more flexible than constant facets, but at the expense of having needing three parameters to be estimated for them. If *b* and *c* are small enough, and we are concerned only with a small area within a larger planar facet, then a constant approximation may be sufficiently accurate.

2.1.1 Generalizing the facet model

The facet model is the basis for the more general *segmentation-based* approach to image modeling. The main difference is that the requirement for a *polynomial* approximation is relaxed. Instead, the interior of a segment may be modeled using any mathematical description which defines a *homogeneity criterion* for the segment. For example, texture is often modeled, not as a functional surface, but as a low order Markov model [4]. A first order Markov model is probability distribution over pairs of pixel values. This model is quite different from a polynomial one, but it can still be used to define a homogeneity criterion. An important attribute of any homogeneity criterion used in image segmentation is whether or not it takes spatial coherence into consideration. Techniques such as thresholding consider only pixel values, and not their relative spatial positions. This can produce unconnected segments, which may or may not be desirable.

2.2 Image sampling

Consider the simple undigitized monochrome scene in Figure 2.3(a). It consists of a square object of uniform high intensity on a dark background of uniform low intensity, and is perfectly piece-wise constant. Figure 2.3(b) illustrates the discretization of this scene into a 9 x 11 digital image. The intensity of each pixel in the final image, shown in Figure 2.3(c), is set to the average light intensity of the area each pixel covers in the original scene. The resulting image is also piece-wise constant, consisting of two distinct segments. This is because the boundary between object and background coincided with the alignment of the pixel grid.

Figure 2.4 shows what occurs when an object in a scene does not align exactly with the sampling grid. The sampling process has produced a range of pixel intensities in the digitized image. Each pixel on the object boundary has received an intensity which is a blend between the intensities of the two original segments.

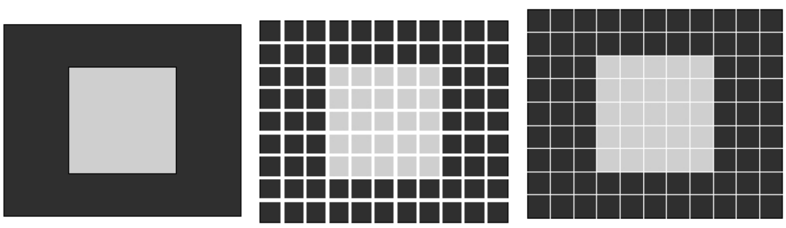

Figure 2.3: Discrete sampling of an object aligned to the pixel grid: (a) original scene; (b) superimposed sampling grid; (c) digitized image.

In fact, there are now seven unique pixel intensities compared to the original two.

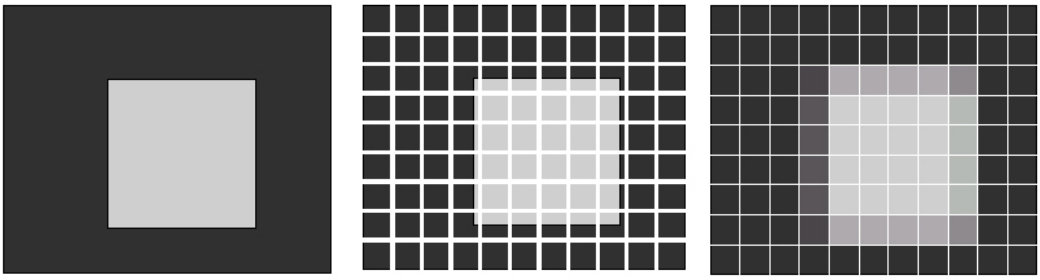

Figure 2.4: Discrete sampling of an object mis-aligned with the pixel grid: (a) original scene; (b) superimposed sampling grid; (c) digitized image.

Most observers would still assert the existence of only two segments, but would have some difficulty assigning membership of each boundary pixel to a specific segment. One interpretation is that those pixels with intermediate intensities have partial membership to both segments. If an application requires a pixel to belong only to one segment, that segment in which it has maximum membership could be chosen. Alternatively, it could be claimed that the image is still piece-wise constant, but now consists of ten segments. These differing interpretations highlight the fact that an image model for a continuous scene may no longer apply to its digitized counterpart.

2.3 Edges and lines

Edges are the boundaries between segments. Lines may be regarded as very narrow segments, having two edges situated very close together. Figure 2.5 gives a one dimensional example of two types of edges and the two corresponding types of lines. The image models used for edges and lines are not unrelated to the models used for the interior of segments.

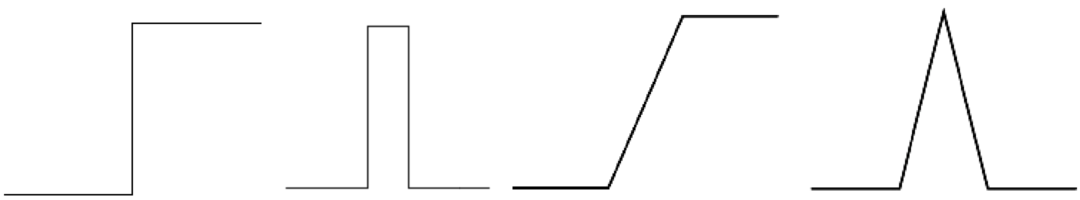

Figure 2.5: (a) step edge; (b) line; (c) ramp edge; (d) roof.

The *step edge* defines a perfect transition from one segment to another. If segments are piecewise constant and pixels can only belong to one segment, then a step edge model is implicitly being used. If a segment is very narrow, it necessarily has two edges in close proximity. This arrangement is called a *line*. An arguably more realistic model for edges is the *ramp edge*. A ramp allows for a smoother transition between segments. This may be useful for modeling the blurred edges created from sampling a scene containing objects not aligned to the pixel grid. Two nearby ramp edges result in a line structure called a *roof*.

Edge profiles may be modeled by any mathematical function desired, but steps and ramps are by far the most commonly used. If a ramp transition occurs over a large number of pixels, it may be difficult to discriminate between it being an edge, or a planar facet. If pixels along the ramp are assigned to a particular segment, the values of those pixels may be dissimilar to the majority of pixels from inside the segment.

2.3.1 A useful noise model

Different types of noise may be introduced at each step of the image acquisition process. The general functional form for the noise component, $n(x, y)$, at each pixel is given in Equation 2.1. Parameters x and y are the spatial position of the pixel in question, $\vec{\theta}$ is a vector of fixed parameters determining some properties of the noise such as its intensity and spread, and **f** is the original image which may be required if the noise term is data dependent. This formulation has scope for a huge number of possible noise functions.

$$n\,(x, y : \vec{\theta}, \mathrm{f}) \tag{2.1}$$

A simple, but still useful and versatile noise model is additive zero-mean Gaussian noise which is independently and identically distributed (i.i.d.) for each pixel. Under this model the noise *adds* to the original pixel value before digitization. The noise term may be written like Equation 2.2, where "$\sim N\,(\mu, \sigma^2)$" denotes a random sample from a normal distribution of mean μ and variance σ^2. Figure 2.6 plots the shape of this noise distribution when $\sigma^2 = 1$.

$$N\,(x, y ; \sigma^2) \sim N\,(0, \sigma^2) \tag{2.2}$$

Because the noise is additive and symmetric about zero, it has the desirable effect, on average, of not altering the mean intensity of the image. It only has one parameter, the variance σ^2, which determines the spread or strength of the noise. Although the work in this thesis assumes that the noise variance is constant throughout the image, it would be possible to vary it on a per pixel basis.

Consider a *constant* facet containing pixels with intensity z. After $N(0, \sigma^2)$ noise is added, it is expected that 99.7% of pixels will remain in the range $z \pm 3\sigma$. This is called the 3σ confidence interval for z [5].

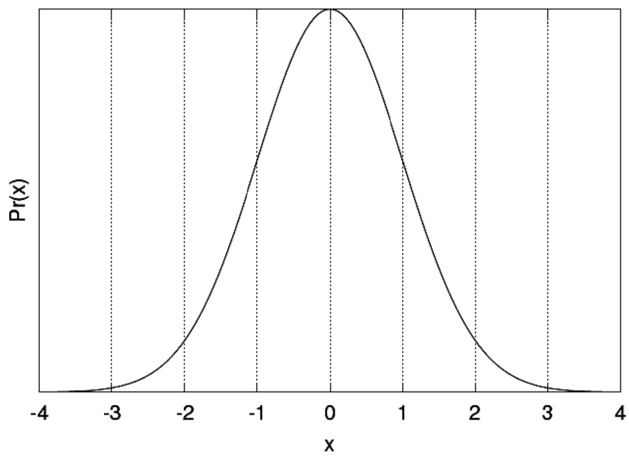

Figure 2.6: The standard Gaussian distribution: z~N(0,1)

Thorough examination of Figure 2.6 shows that very little probability remains for values of z outside the confidence interval. Table 2.1 given below lists the number of standard deviations from the mean for which a given proportion of a normally distributed data is expected to lie.

Table 2.1: Confidence intervals for normally distributed data

Fraction of data (%)	Number of standard deviations from mean
50.0	0.674
68.3	1.000
90.0	1.645
95.0	1.960
95.4	2.000
98.0	2.326
99.0	2.576
99.7	3.000

Figure 2.7 shows the effect of two different noise variances on the square image, which has segments of intensity 50 and 200. When $\sigma = 10$, the two segments remain distinct. When σ is quadrupled, the square obtains some pixels clearly having intensities closer to the original *background* intensity. After clamping, the 3σ confidence interval for the background is 50±120 = (0,170), and 200±120 = (80,255) for the square. These limits have an overlap of 90, or approximately 2σ. On average, this situation would result in about 5% of pixels being closer to the opposing segment mean. For the 11 x 9 square image this would affect about 5 pixels, roughly corresponding with what is observed.

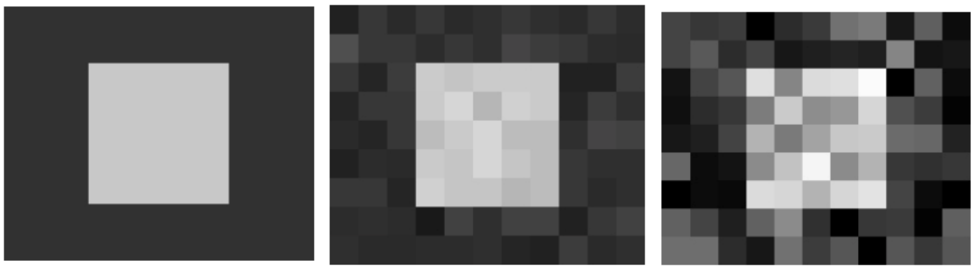

Figure 2.7: The effect of different Gaussian noise levels: (a) no noise; (b) added noise σ = 10; (c) added noise σ = 40.

For piece-wise constant segments, the noise standard deviation defines a *natural level of variation* for the pixel values within that segment. The natural level of variation describes the amount by which pixel values may vary while still belonging in same segment. For the case of planar segments, the natural level of variation depends on both the noise level and the coefficients of the fitted plane. If a global planar segment only has a mild slope, then the variation due to the signal may be negligible for any local window onto that segment. The natural level of variation in the window will be dominated by the noise rather than the underlying image model.

3. Test images

In this section, different denoising algorithms will be compared. It is laborious to provide results for a large set of images at each stage of the discussion. For this reason, one or more images from the small set introduced here will be used consistently throughout this section. The use of a larger set of image test set will be deferred until the final results are considered.

3.1 Square

The square image in Figure 3.1 is an 8 bit per pixel greyscale image of resolution 11 x 9. It consists of a 25 pixel square of intensity 200 atop a 74 pixel background of intensity 50. It will be often used for illustrative purposes, and for *subjectively* examining the effect of various techniques.

3.2 Lenna

The lenna image [6] in Figure 3.2 has become a *de facto* standard test image throughout the image processing and image compression literature. Its usefulness lies in the fact that it covers a wide range of image properties, such as flat regions, fine detail, varying edge profiles, occasional scanning errors, and the fact that so many authors produce results using it as a test image. One interesting feature is that the noise in lenna seems to be inversely proportional to the brightness, perhaps a legacy of having been scanned from a negative.

Figure 3.1: The 11 x 9 8 bpp square test image.

Figure 3.2: (a) the 512 x 512 8 bpp lenna image; (b) histogram.

3.3 Montage

Figure 3.3 shows a greyscale test image called montage [7], and its histogram. It has resolution 512x512 and uses 8 bits per pixel. The image consists of four quadrants. The top left is the middle 256x256 section of a smoothed version of the lenna image. The bottom right is a right hand fragment of a German village street scene, which contains some small amounts of natural noise. The top right is a synthetically created image which is perfectly piece-wise constant. It covers a range of segment boundary shapes and intensity differences between adjacent segments. The bottom left is the same as the top right, except that the segments are piece-wise planar, covering a range of horizontal and vertical gradients.

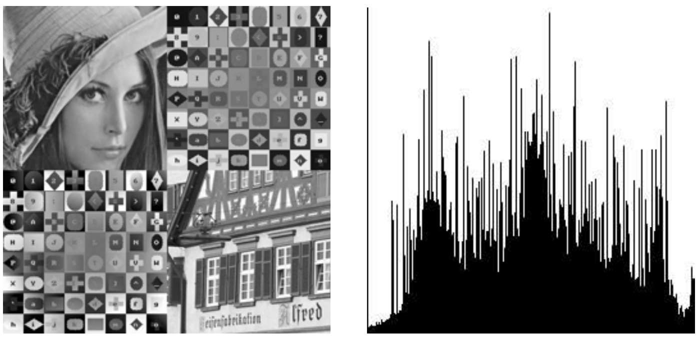

Figure 3.3: (a) the 512 x 512 8 bpp montage image; (b) histogram

It is hoped that this image covers a wide range of image properties, while also being very low in noise. The low noise level is important for experiments in which synthetic noise will be added. The image has features such as constant, planar and textured regions, step and ramp-like edges, fine details, and homogeneous regions. The montage image is used for measuring *objectively* the RMSE performance of various techniques.

4. Conclusions

It has been shown that the principles of local segmentation can be used to develop effective denoising algorithms. After many analyses, the FUELS algorithm for denoising greyscale images contaminated by additive noise was presented. FUELS has an efficient implementation, and only requires one parameter, the level of noise in the image. This can be supplied by the user, or FUELS can determine it automatically. FUELS was shown to outperform existing methods, like SUSAN and GIWS, for a variety of images and noise levels.

Examination of images in terms of local segmentation has led to a better understanding of image processing on a small scale, particularly for the commonly used 3 x 3 configuration. At conversations and seminars, I have often heard the off-hand comment that “only 10% to 20% of images are edges”. Analysis of the distribution of k values chosen for images in the research suggest that only around 50-60% of pixels are locally homogeneous, 20-30% consist of two segments, and 10-20% tend to be difficult to model well. Perhaps the speakers were confusing edges with those pixels which are difficult to predict or model. The success of FUELS suggests that attempts to model these difficult blocks does not significantly improve denoising performance.

5. References

[1] Range Image Segmentation Comparison Project at the University of South Florida. http://marathon.csee.usf.edu/range/seg-comp/images.html

[2] Klaus Köster and Michael Spann. MIR: An approach to robust clustering – application to range segmentation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 22(5):430–444, May 2000.

[3] R. M. Haralick and L.Watson. A facet model for image data. *Computer Graphics and Image Processing*, 15:113.129, 1981.

[4] N. Huang. Markov model for image segmentation. In *22nd Allerton Conference on Communication, Control, and Computing*, pages 775.781, October 3.5 1984.

[5] Frederick Mosteller and Robert E. K. Rourke, editors. *Sturdy Statistics: Nonparametric and Order Statistics*. Addison-Wesley, 1974.

[6] Playmate of the Month: Lenna Sjööblom. *Playboy: Entertainment for Men*, pages 134–141, November 1972.

[7] Bernd Meyer and Peter E. Tischer. TMW – a new method for lossless image compression. In *Picture Coding Symposium (PCS'97)*, pages 533–538, Berlin, Germany, 1997. VDE-Verlag GMBH.