

Normalized Averaging Using Adaptive Applicability Functions with Applications in Image Reconstruction from Sparsely and Randomly Sampled Data

Tuan Q. Pham and Lucas J. van Vliet

Pattern Recognition Group, Delft University of Technology,
Lorentzweg 1, 2628 CJ Delft, The Netherlands
{tuan,lucas}@ph.tn.tudelft.nl

Abstract. In this paper we describe a new strategy for using local structure adaptive filtering in normalized convolution. The shape of the filter, used as the applicability function in the context of normalized convolution, adapts to the local image structure and avoids filtering across borders. The size of the filter is also adaptable to the local sample density to avoid unnecessary smoothing over high certainty regions. We compared our adaptive interpolation technique with conventional normalized averaging methods. We found that our strategy yields a result that is much closer to the original signal both visually and in terms of MSE, meanwhile retaining sharpness and improving the SNR.

1 Introduction

Conventional interpolation techniques often rely on either one or two input signal characteristics such as the signal amplitude (bilinear interpolation), the arrangement of sampled signals (natural neighbor interpolation [1]), or the certainty of signals (normalized convolution [2]) but never all of them. While the existing framework of normalized convolution is efficient in finding a local representation of the signal incorporating the signal uncertainties, it does not take the signal structures into account. In fact, when interpolation of sparsely sampled signals is concerned, the neighborhood's structural content should play an important role in shaping the missing values. This paper points out how adaptive filtering can be used in conjunction with normalized convolution to take advantage of both signal certainty and structural content in image analysis problems.

The structure of the paper is as follows; we start with a short description of normalized averaging. We then introduce adaptive parameters into the applicability function, starting from sizes then shapes of the filter. The sizes of the filter vary according to the local sample density and the shapes are steered towards the local signal structure, which comprises of orientation, anisotropy and curvature. Finally, comparisons are made between our adaptive method and other interpolation techniques with respect to mean square error (MSE), peak signal-noise ratio (PSNR) and a sharpness measure.

2 Normalized Averaging

Normalized convolution (NC) [2] is a method for local signal modeling that takes signal uncertainties into account. While the applications of normalized convolution are numerous, the simplest and most striking example is interpolation of incomplete and uncertain data using the special case called normalized averaging (NA). In this reconstruction algorithm, an interpolated value is estimated as a weighted sum of neighboring values considering their certainty and applicability, in which applicability refers to the influence of the neighbors towards the interpolated value (see equation 1).

$$\frac{(s \cdot c) * a}{c * a} \quad (1)$$

where r denotes the reconstructed version of the measured signal s with the certainty c ; the normalized averaging is done via two convolutions (*) with the applicability function a . This applicability function is usually modeled by a localized function such as the Gaussian function. Since Gaussian filtering has a very fast recursive implementation [3], the NC reconstruction algorithm is efficient while giving reasonable results (see figure 1b).




Fig. 1. (a) original image. (b-d) Results of normalized averaging reconstruction of the image from 10% of the original information, with the applicability function being: (b) the Gaussian function: $\sigma = 1$, (c) the Knutsson applicability function: r^{-3} , (d) the local structure adaptive applicability function: $\sigma_u = \frac{1}{2}(1-A)\sigma_{dens}$, $\sigma_v = \frac{1}{2}(1+A)\sigma_{dens}$.

3 Scale Adaptive Applicability Functions Based on Sample Density

Since interpolation favors the contribution of a closer pixel over a distant pixel, a commonly used applicability is the Gaussian function. However, practical Gaussian kernels decrease rather gradually near the center and they are often truncated outside a certain radius for implementation efficiency. These characteristics impose a constant smoothing factor on the image while the function support is not large enough to recover big missing blobs of data. Sharp functions with large support have been suggested [2][4], however this is not the real solution to the problem. To minimize the unnecessary smoothing and maximize the spatial support, a scale-adaptive filter that shrinks or grows depending on the local sample density should be used.

We define local sample density as $\sigma_{dens}(x, y)$ such that σ_{dens} equals the radius of a pillbox, centered at position (x,y) , that encompasses a total certainty equal to one. For a missing sample image $c : \mathbb{R} \rightarrow \{0,1\}$, σ_{dens} is the distance to the nearest pixel with $c=1$, which relates to the Euclidean distance transform [5]. For images with arbitrary certainty ranging from zero to one, we filter the input certainty by a bank of pillbox filters with increasing radii. The radius corresponding to a response of one is found by a quadratic interpolation of the filtered results.

4 Shape Adaptive Applicability Functions Based on Local Structure

While a variable-sized applicability function solves the problem of excessive image smoothing, it does not guarantee the continuity of the signal structure. This drawback can be seen in figure 3c where the reconstructed image, though visually informative, shows severe edge jaggedness. A natural approach to solve this shortcoming is making the applicability function adaptive to the underlying signals structure. In other words, it is desirable that the reconstructed image patch shows the same oriented pattern as its surrounding neighborhood, if there is any. The corrupted samples therefore should be more similar to the high confidence pixels along its linear orientation. As a result, anisotropic filters should be used as the applicability functions. The problem remains as how to estimate the local structure of an image with missing samples. Fortunately, (differential) normalized convolution [2][6] offers a way to estimate the image gradients from incomplete and uncertain data. We can then construct the gradient structure tensor (GST) [7][8] to analyze the local linear structures.

Local image structure at a certain scale can be characterized by orientation, anisotropy and curvature. These parameters can be extracted using the regularized gradient structure tensor T as shown in equation 2, where the eigenvector \bar{u} corresponding to the largest eigenvalue λ_u determines the local orientation ϕ ; A and \mathcal{K} denote local anisotropy and curvature, respectively. Note that the curvature \mathcal{K} is computed from the differential of the mapping $M(\phi) = \exp(j2\phi)$ rather than on the orientation itself which contains jumps between $\pm \pi$ [9].

$$T = \overline{\bar{\nabla} I \bar{\nabla} I^T} = \lambda_u \bar{u} \bar{u}^T + \lambda_v \bar{v} \bar{v}^T \quad A = \frac{\lambda_u - \lambda_v}{\lambda_u + \lambda_v} \quad \mathcal{K} = \frac{\partial \phi}{\partial \bar{v}} = -\frac{j}{2} M(-\phi) \frac{\partial}{\partial \bar{v}} M(\phi) \quad (2)$$

The idea of a shape-adapted smoothing is not new. Adaptive filtering that shapes the smoothing kernel after the principal components of the inverse GST has been proposed by a number of authors [10][11]. Here, we introduce curvature to the filter by bending the same anisotropic Gaussian applicability kernel. The directional scales of the kernel before curvature bending are given by:

$$\sigma_u = C(1 - A)^\alpha \sigma_{dens} \quad \sigma_v = C(1 + A)^\alpha \sigma_{dens} \quad (3)$$

where the scaling is based on the adaptive scale parameter σ_{dens} (see previous section), an additional term C that may encompass the local signal-to-noise ratio, and the degree of structure enhancement depends on the anisotropy A through an exponent α .

Special attention needs to be given to the application of the gradient structure tensor to incomplete data. Computation of the gradient image can be done in either three ways: (1) reconstruct the image by normalized convolution and compute the gradient from this result, (2) using normalized differential convolution NDC [2], (3) using the derivative of the normalized convolution equation (in contrast to the normalized convolution result) DoNC [6]. An extensive evaluation [6] for natural images, all three methods: gradient from normalized convolution result, NDC and DoNC show comparable performance over a wide range of SNR's. Due to the variable sample density of our input images, we use the first method where the gradient is estimated from an initial image reconstruction using NC with scale-adaptive isotropic Gaussian applicability. The Gaussian scales of the gradient operator and the tensor smoothing for the GST is set to one and three respectively.




Fig. 2. (a-c) Image transformation for the local structure adaptive filtering: curvature bent and rotated to be resampled (at the white dots) by bilinear interpolation. (d) look-up Gaussian kernel with sampled coefficients at the white dots.

4.1 Implementation Details

In the reconstruction experiment given in figure 1, all input samples are lined up properly on the sampling lattice. This makes it very easy to collect the neighbors' positions and look up the corresponding filtering coefficients. However, our adaptive normalized convolution technique also works on randomly sampled signals. The most challenging task is then how to compute the kernel of an off-grid centered filter. Continuous normalized convolution (CNC) [4] approximates the coefficients by truncating the Taylor expansion up to the first order. The method, relying on the first-order partial derivative kernels at every pixel, is therefore not suitable for our varying applicability function. Furthermore, CNC cannot accurately approximate Gaussian kernels of sigma smaller than 1, which is often necessary when little smoothing is required.

Here we present a filter construction approach that uses a large look-up Gaussian kernel. The size of such kernel is chosen such that most commonly used smaller kernels can be easily extracted from the master kernel by pixel-skipping subsampling. For example, a master kernel of size 121x121 hold coefficients of any kernel with sizes from $\{3,5,7,9,11,13,21,25,31,41,61,121\}$. In fact, we only use kernels with a right combination of these sizes such that the distance between adjacent resampled positions

on the input image is less than the resolution grid. In this way, we can minimize the kernel size at each pixel while still satisfy the resampling condition (figure 2). With this simple scheme, Gaussian coefficients need not be recomputed for on-grid centered kernels. Coefficients for off-grid centered kernels can also be estimated from the master kernel using bilinear interpolation. Such an interpolation scheme allows accurate approximation of any Gaussian kernels with sigma greater than 0.3 grid size.

Table 1. MSEs and PSNRs of the reconstructed Lena images from 10% original information

	Random 10% of Lena image	Knutsson r^{-3} applicability function, fig. 1c.	Gaussian $\sigma=1$, fig. 1b.	Scale-adaptive isotropic Gaussian	Shape and scale adaptive Gaussian, fig. 1d.
MSE	31449	126.13	88.23	88.32	43.77
PSNR	3.15	27.12	28.67	28.67	31.72

5 Experiments on Image Reconstruction from Missing Samples

In order to compare the result of the adaptive normalized averaging with the traditional normalized averaging method, the same experiments as presented in [2] are carried out. In the first experiment, the famous Lena image has been randomly sampled to a test image containing only 10% of the original information. The reconstruction results have been shown earlier in figure 1. A quantitative comparison is presented in table 1. The second experiment involves reconstruction of the same image having large missing regions (see figure 3 and table 2). While the standard normalized averaging can fill in the ‘holes’ with ‘plausible’ data [2], the adaptive filtering approach really proves its strength by being capable of extending local linear structures into the missing regions.




Fig. 3. (a) Lena ‘hole’ image with large missing blobs. (b-d) Results of normalized averaging with the applicability function as (b) the Knutsson applicability function, (c) the isotropic scale-adaptive Gaussian function, (d) the local structure adaptive Gaussian function (eq. 3).

With the lowest MSE and the highest PSNR (equation 4), adaptive applicability function clearly outperforms other listed applicability functions. It is also interesting to note that the Knutsson applicability function scores poorly for both quantitative tests even though the results are claimed to be better than the Gaussian function [2].

Table 2. MSEs and PSNRs of the reconstructed regions from the Lena ‘hole’ image

	Lena image with holes, fig.3a.	Knutsson r^{-3} applicability function, fig.3b	Gaussian $\sigma=3$	Scale-adaptive isotropic Gaussian, fig.3c	Shape and scale adaptive Gaussian, fig.3d
MSE	7235	133.21	106.11	83.63	41.05
PSNR	9.54	26.89	27.87	28.91	32.00

This is due to the fact that the fixed Knutsson applicability function is only significant in the very center while having near-zero values elsewhere. Though this is desirable for high certainty regions, low certainty regions are not reconstructed very well.

$$MSE = \frac{1}{N} \sum_{j=1}^N (I_j - \hat{I}_j)^2 \qquad PSNR = 10 \log_{10} \frac{MAX^2}{MSE} \qquad (4)$$

To prove the superior in sharpness quality, we use the sharpness comparison measure by Dijk et. al. [12]. The method is described as follows: a scatter diagram composed of a set of data points is plotted such that the coordinates of each data point correspond to the gradient magnitudes of the two tested images at the same image location. Data points below the line $y=x$ correspond to pixels with a higher gradient magnitude in the first image than in the second. Those points favor the sharpness of the first image over the second. Pixels above the line $y=x$, on the other hand, favor the smoothness of the first image over the second. Normally, these two types of data points are always present in a scattergram because an image is filtered for both detail enhancement and noise suppression. As a result, the gradient magnitudes of edge pixels increase while those of noisy background pixels decrease. However, due to the higher gradient energy that edge pixels possess, they form a cluster of data points that extends further away from the origin than the cluster formed by the noisy pixels. A line going through the origin can be fitted to each cluster, resulting in a slope for the degree of edge sharpening or noise smoothing accordingly.

As can be seen from both scattergrams in figure 4, the majority of data points form a cluster below the $y=x$ line. The result from the diagram on the right is not as clear though because what is really reconstructed in the Lena ‘hole’ experiment is the continuation of the image linear structures rather than the image gradient. This sharpness measure agrees with the subjective perception of the better quality of the reconstructed images using our adaptive method over Knutsson’s fixed applicability function.

6 Conclusion and Further Research Directions

We have demonstrated that normalized convolution with adaptive applicability is an attractive framework for signal analysis that takes signal uncertainties as well as local structures into account.

The application of this framework in the interpolation of incomplete and uncertain data image yields a method that sharpens edges, suppresses noise and avoids edge




Fig. 4. (left) Gradient magnitude scattergram of the reconstructed Lena images from a random 10% of the original information using adaptive Gaussian (horizontal axis) and Knutsson applicability function (vertical axis). (right) similar scattergram for the Lena ‘hole’ experiment.

jaggedness by extending local image orientation into the low certainty regions. The comparisons of visual sharpness, MSE and PSNR on the reconstructed images show a significant improvement of the new method over existing normalized convolution based interpolation techniques.

Apparently, normalized averaging is not the only technique for image reconstruction. Image inpainting [13] also reconstruct missing sample images by extending isophote lines from the exterior into the missing regions. However, different from our non-iterative adaptive normalized averaging method, image inpainting is a diffusion process that takes many iterations to propagate orientation and image grayscale into the restored areas. The information to be propagated is computed from the intact surrounding area. As a result, image inpainting is only suitable for the restoration of damaged photographs and movies or the removal of small objects where the majority of image information is present. For example, the algorithm would give a satisfactory




Fig. 5. Comparison of our image reconstruction with Sapiro’s image inpainting: (a) our reconstruction from 10% original information, (b) an unsuccessful inpainting result for the same experiment in (a), (c) our reconstruction for one of the missing holes in the Lena ‘hole’ experiment, (d) same hole reconstruction by inpainting. Notice the continuation of the hair structure reconstructed by our algorithm as opposed to the mean gray level filling-in from inpainting

result in the Lena 'hole' image experiment, but not when 90 per cent of the image data is unavailable (see figure 5). In addition, due to the discrete implementation of diffusion, the computation of the gradient vector fields is not as robust to noise as done with the gradient structure tensor. Image inpainting also cannot handle irregular sample input as is possible in our adaptive filtering approach.

Future research will focus on the application of adaptive normalized convolution in other problems such as image denoising, video restoration, data fusion, optic flow computation and compensation and super-resolution. The local structure adaptive filtering scheme presented in this paper can also be extended to the local structure neighborhood operators, in which more robust estimators than the weighted mean (as used in normalized convolution) can be used to approximate the local data.

Acknowledgements. This research is partly supported by the IOP Beeldverwerkings project of Senter, Agency of the Ministry of Economic Affairs of the Netherlands.

References

1. A. Okabe, B. Boots, K. Sugihara, *Spatial tessellations concepts and application of Voronoi diagrams*, Wiley, 1992.
2. H. Knutsson, C-F Westin, Normalized and differential convolution: Methods for interpolation and filtering of incomplete and uncertain data, in *CVPR'93*, 1993. 515–523.
3. I.T. Young, L.J. van Vliet, Recursive implementation of the Gaussian filters, *Signal Processing*, 44(2), 1995. 139–151.
4. K. Andersson, H. Knutsson, Continuous Normalized Convolution, *ICME'02*, 2002. 725–728.
5. P.E. Danielsson. Euclidean distance mapping. *CGIP*, 14, 1980. 227–248.
6. F. de Jong, L.J. van Vliet, P.P. Jonker, Gradient estimation in uncertain data, *MVA'98 IAPR Workshop on Machine Vision Applications*, Makuhari, Chiba, Japan, Nov. 1998. 144–147.
7. L.J. van Vliet, P.W. Verbeek, Estimators for orientation and anisotropy in digitized images, in J. van Katwijk et.al. (eds.), *ASCI'95, Proc. First Annual Conf. of the Advanced School for Computing and Imaging*, Heijten, the Netherlands, 1995. 442–450.
8. H. Knutsson, Representing local structure using tensors, *SCIA'89*, Oulu, Finland, 1989. 244–251.
9. M. van Ginkel, J. van de Weijer, L.J. van Vliet, P.W. Verbeek, Curvature estimation from orientation fields, *SCIA'99*, Greenland, 1999. 545–551.
10. M. Nitzberg, T. Shiota, Nonlinear Image Filtering with Edge and Corner Enhancement. *IEEE Trans. on PAMI* 14(8), 1992. 826–833.
11. A. Almansa and T. Lindeberg, Enhancement of Fingerprint Images using Shape-Adapted Scale-Space Operators, in J. Sporring, M. Nielsen, L. Florack and P. Johansen (eds.), *Gaussian Scale-Space Theory*, Kluwer Academic Publishers, 1997. 21–30.
12. J. Dijk, D. de Ridder, P.W. Verbeek, J. Walraven, I.T. Young, L.J. van Vliet, "A new measure for the effect of sharpening and smoothing filters on images", in *SCIA'99*, Greenland, 1999. 213–220.
13. C. Ballester, M. Bertalmio, V. Caselles, G. Sapiro, J. Verdera, Filling-in by Joint Interpolation of Vector Fields and Gray Levels, *IEEE Transactions on Image Processing*, 10 (8), 2001. 1200–1211.