FSED - Feature Selective Edge Detection

Magnus Borga
Computer Vision Laboratory
Dept. of Electrical Engineering
Linköping University
SE-585 94 Linköping
Sweden
magnus@isy.liu.se

Helge Malmgren
Dept. of Philosophy
Göteborg University
Box 200
SE-405 30 Göteborg
Sweden
helge.malmgren@phil.gu.se

Hans Knutsson
Computer Vision Laboratory
Dept. of Electrical Engineering
Linköping University
SE-585 94 Linköping
Sweden
knutte@isy.liu.se

Abstract

We present a novel method that finds edges between certain image features, e.g. gray-levels, and disregards edges between other features. The method uses a channel representation of the features and performs normalized convolution using the channel values as certainties. This means that areas with certain features can be disregarded by the edge filter.

The method provides an important new tool for finding tissue specific edges in medical images, as demonstrated by an MR-image example.

1. Introduction

Classification based on gray-scale value can often be useful and relevant e.g. in medical images such as computed tomography (CT) or magnetic resonance imaging (MRI). The simplest form of gray-scale classification is to classify each pixel according to its value. An obvious problem with pixel-based classification is high noise-sensitivity. For that reason spatial operations, such as smoothing and edge detection, are often necessary.

Edge detection is usually performed by a linear filtering operation, i.e. a convolution between the image and a set of filter kernels. A problem, however, not often mentioned with linear filtering is that the result greatly depends on the mapping from the real world objects to the gray-scale value in the image. The root of this problem is that the imaging technique implies a mapping from a high-dimensional feature space down to a one-dimensional gray-scale. This mapping is in a sense arbitrary since there is no unique way of ordering points in a space of more than one dimension and the mapping destroys whatever metric that would be useful in a high-dimensional feature space. This means that

the strengths of two edges, i.e. differences between grayvalues, are not comparable. In other words, the strength of an edge is not related to the importance of the edge.

As an example, two tissues that give similar gray-scale values in an MR-image will give a much weaker edge-filtering response than two tissues that maps to very different gray-scale values. This means that some edges that happens to give weak responses due to the given mapping in the scanning device might disappear in comparison to other edges in the same image.

A non-linear operation such as a threshold could in principle solve this problem but it would give a very noisy result. Another problem with such an approach is that edges between different gray-levels would give the same response, i.e. information about what gray-scale values that lie behind the edges is lost. Such information is, however, very useful when the edges are going to be used in gray-scale classification. If we, for example, want to detect an object with a certain gray-scale value we are only interested in edges between that value and other values but not in edges between other gray-scale values.

Here we propose an approach to solve the problem discussed above by separating the image into a set of grayscale channels and filter certain combinations of channels in order to detect the desired edges and at the same time disregard other edges.

A channel representation expands a one-dimensional value (i.e. grey-value) to an N-dimensional vector. The channels are located such that only one or a few neighbouring channels are active at a time while most of the channels are zero. Each channel can be seen as a response of a filter that is tuned to some specific feature value. In this case the feature is simply the gray-value of the pixel. The channel representation is also known as value encoding [1] or sparse distributed coding [2].

When a pixel is assigned to a certain gray-scale channel, it is associated with a certainty value that indicates how well

it fits to the channel. The certainty values are then utilized in *normalized convolution* (NC) [3, 4], a method for filtering sparse and uncertain data.

To illustrate the proposed method, we look at a very simple one-dimensional example. Consider the signal in figure 1. Assume we are interested in finding edges between value 1 and 2 (dotted lines). It is in general not possible to map the signal values such that the desired result can be obtained with ordinary convolution. Using NC, however, the signal values in which we are not interested can simply be ignored. The result using standard convolution with an edge filter would in principle look like the middle signal in figure 1. If we, however, divide the signals into channels and apply NC using channel 1 and 2 as certainty images, we get a result that looks something like the lower signal in figure 1.




Figure 1. A simple example.

In the following section, the method is explained in more detail. In section 3, some experimental results are shown. Finally, section 4 contains a summary and discussion.

2 The method

The method can briefly be summarised as follows:

- 1. Separate the image into different gray-scale channels.
- **2.** Assign certainties to every pixel according to how well it fits the prototype value of the channels of interest and generate a certainty image.
- **3.** Perform NC on the quantized image using the certainty image.

2.1 Defining channels

There are many ways of choosing the positions and shapes of the channels. The distribution of channels on the gray-scale should optimally be made such that all pixels in an object fall into the same channel. In general, this is of course impossible, but here we are interested in images where gray-scale classification is meaningful and in such images it should be possible to get close to such a distribution. For example an MR image of the brain typically

consists of a relatively small number of typical gray-values; one for gray matter, one for white matter, one for fat, one for water, etc.

To find the gray-values for the different channels we use the histogram of the image. Each gray-value that is typical for a class of objects in the image will give a peak in the histogram and these peaks are used to locate the channels. The gray-value at a certain peak can be seen as the *prototype value* of the corresponding channel. An example of a histogram of an MR image of the brain is shown in figure 2.




Figure 2. A histogram of an MR image of the brain before and after low-pass filtering. The vertical lines indicates local maxima and minima.

In order to find the peaks in the histogram, we start by low-pass filtering the histogram to get rid of small local maxima. Then a differentiating operator is applied and the zero-crossings of the differentiated histogram are detected. By interpolation this gives the position of the histogram peaks with sub-bin precision if necessary. The right plot in figure 2 shows the histogram after low-pass filtering. Local maxima and minima are marked with vertical lines.

The next question is the shape of the channels. One way would be to use Gaussians to model the histogram. Such an approach can be motivated by assuming that each object give a certain gray-value which then is disturbed by Gaussian noise. Here, however, we have chosen a simpler approach. We use non-overlapping channels and the boundaries of the channels are placed at the minima of the histogram. This gives asymmetric channels where a pixel value that corresponds to a local maximum in the histogram gives the value one in one channel and a pixel value that corresponds to a local minimum in histogram gives the value zero. The actual shape of the channel is discussed in the next subsection.

2.2 Assigning certainties

The output of a channel is interpreted as the certainty of the statement that the pixel value is equal to the prototype value of that channel. We get certainty one if the pixel value and the prototype value are identical. But it is not obvious how to choose the mapping from the distance from the prototype value to certainty. As we have seen we want the boundary and values outside to give zero certainty. For the values between the boundary and the prototype value we have used the following mapping:

$$c = 1 - |r|^s \tag{1}$$

where r is the relative distance from the prototype value and $s \ge 0$ is a constant that controls the shape of the channel. Relative distance means that the distance is normalized so that the border of the channel always has the relative distance r=1. The higher the value of s, the wider and more box-like is the shape of the channels. The certainty channels for the histogram in figure 2 for s=2 are shown in figure 3. The vertical lines mark the minima and maxima in the histogram.




Figure 3. An example of channels with s = 2.

2.3 Generating a certainty image

This step is simple. The certainty-channels in which we are interested are simply added together. Since the channels are non-overlapping, only one channel can be active at a certain pixel. This means that certainty is still bounded between zero and one after the addition. The certainty image will then have zero certainty where there are pixel values outside the range of our interest and higher certainty in pixels within the desired gray-scale intervals.

2.4 Normalized convolution

At this stage it is important to remember that we want to detect edges between different gray-scale values and not between different certainty values. In other words, we must be aware of the difference between the value of a pixel and the certainty of that value. As an example, the pixel value zero means black while the certainty zero means that the pixel value is unknown. In the latter case, the pixel should not effect the filtering result.

Ordinary convolution cannot take into account different certainties. To handle this, we use a method called *normal*-

ized convolution (NC) [3, 4] instead, which is a method for filtering sparse and uncertain data.

The aim of this paper is not to explain NC. Furthermore, the space available does not permit a detailed description. However, a brief description is appropriate.

In ordinary convolution, the result in each point is a scalar product between a signal vector \mathbf{f} and a filter vector \mathbf{b} , i.e. $\tilde{x} = \langle \mathbf{f} \mid \mathbf{b} \rangle$. In NC, this scalar product is weighted with diagonal matrices \mathbf{W}_a and \mathbf{W}_c so that $\tilde{x} = \mathbf{b}^T \mathbf{W}_a \mathbf{W}_c \mathbf{f}$. The diagonal of \mathbf{W}_c contains the certainty values and the diagonal of \mathbf{W}_a is called *applicability function* which can be loosely described as a certainty function for the filter. If we have a set of filters we can arrange the filters as columns in a matrix \mathbf{B} and we then get a vector of scalar products as $\tilde{\mathbf{x}} = \mathbf{B}^T \mathbf{W}_a \mathbf{W}_c \mathbf{f}$. If we call \mathbf{B} a *basis*, the scalar products $\tilde{\mathbf{x}}$ are the coordinates in the *dual basis*. To get the coordinates in the filter basis \mathbf{B} , the coordinates in the dual basis can be transformed with the matrix $(\mathbf{B}^T \mathbf{W}_a \mathbf{W}_c \mathbf{B})^{-1}$ This means the we get the coordinates in the filter basis as

$$\mathbf{x} = (\mathbf{B}^T \mathbf{W}_a \mathbf{W}_c \mathbf{B})^{-1} \mathbf{B}^T \mathbf{W}_a \mathbf{W}_c \mathbf{f}. \tag{2}$$

If **B** is an *orthonormal basis* and we have constant certainties and applicabilities, i.e. \mathbf{W}_a and \mathbf{W}_c are unitary matrices, the coordinates in the basis and in the dual basis are identical. But in general, this is not the case.

3 Experimental results

In figure 4 we present result from an MR-image. This is the same image that generated the histogram in figure 2. The gray-scale is separated into four channels. If a standard edge detection is performed we get the result in the middle image. This is the sum of the magnitudes of the the convolutions between the image and four directional quadrature filters. In the lower image we see the result from NC of the quantized image using channels 2 and 3 as certainty (\mathbf{W}_c). Channels 2 and 3 correspond roughly to gray and white matter respectively.

The filter basis **B** used in this experiment consisted of one constant (DC-filter) and one filter which, together with the applicability \mathbf{W}_a , constitutes a directed quadrature filter with 15×15 coefficients. NC was performed with filters in four directions, 0° , 45° , 90° and 135° and the magnitudes of the convolution results were then added together.

Here we have also lowered the certainty were there are edges in the original image. This is to avoid a problem caused by the low spatial resolution of the MR-image: When there is an edge between e.g. channel 1 and channel 3, the edge pass through channel 2 because of volume effects. This generates a thin stripe of channel 2 at the border between channel 1 and channel 3. To avoid false edges at these positions, the certainty is lowered where there are edges in the original image.

4 Summary and discussion

We have presented a new method for finding edges between certain gray-levels in images. It can find edges between a number of specified gray-levels without being disturbed by other edges in the image. The method separates the gray-scale into a set of channels, assigns certainties to the pixels in accordance to their fit to the channels and then performs normalized convolution.

The separation into gray-scale channels is of course a kind of pixel-based gray-scale classification and hence, it might seem like we have a chicken and egg problem here: In order to make a robust gray-scale classification we need edge information, but in order to obtain the edge information we need to do gray-scale classification. But the separation into channels serves only as a soft, preliminary classification and the proposed method provides a means to combine spatial and gray-level information in an efficient way.


The method is particularly well suited for medical images such as e.g. MR-images where the histogram have a limited number of distinct peaks. In such images, certain tissues maps to certain gray values and the detection of edges between certain gray-values should be useful in segmentation and classification.

As the title of the paper suggests, other features than gray-value could be used, e.g. local orientation, local frequency or colour. A colour image is already divided into three channels (e.g. RGB). Usually, edge detection is applied on one channel at a time. But if we, for instance, only want to find edges between red and green, a filtering of a single colour-channel will not help us. The method proposed in this paper can, however, straightforward be used to solve that problem.

References

- D. H. Ballard. Vision, Brain, and Cooperative Computation, chapter Cortical Connections and Parallel Processing: Structure and Function. MIT Press, 1987. M. A. Arbib and A. R. Hanson, Eds.
- [2] D. J. Field. What is the goal of sensory coding? *Neural Computation*, 1994. in press.
- [3] H. Knutsson and C.-F. Westin. Normalized and Differential Convolution: Methods for Interpolation and Filtering of Incomplete and Uncertain Data. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 515–523, New York City, USA, June 1993. IEEE.
- [4] C.-F. Westin. A Tensor Framework for Multidimensional Signal Processing. PhD thesis, Linköping University, Sweden, SE-581 83 Linköping, Sweden, 1994. Dissertation No 348, ISBN 91-7871-421-4.








Figure 4. Result on an MR image. Top: Original image. Middle: Result from a normal edge detection. Bottom: Result using NC with gray-level channels 2 and 3 as certainty.