Wavelet-based Image Segmentation

Joerg Meyer, Zhihe Zhang
Department of Biomedical Engineering
644E Engineering Tower, University of California, Irvine, CA 92697-2625
{jmeyer | zhihez}@uci.edu

Keywords: Wavelet decomposition, image segmentation, 3-D texture-based volume rendering

ABSTRACT

Wavelet decomposition is a powerful tool for image analysis and data compression on multiple levels. As a lossless or lossy data compression method, wavelets together with quantization have the potential to compress large-scale, three-dimensional image data files, such as CT/MRI scans, cryosections, confocal laser microscopies, at various levels of detail, while the quality of the reconstruction and the compression rate are mainly controlled by the quantization level and the depth of the wavelet decomposition. As a by-product of digital archiving, one can use the detail coefficients obtained in a non-standard wavelet decomposition to obtain directional gradient information on multiple scales. By combining the gradients of all spatial dimensions, the contour of an object that exists on a particular scale can be extracted. This way, objects can be sorted by size. The size of the recognized objects or features depends on the underlying basis function. For a discrete pixel image or a three-dimensional volumetric data set, the number of basis functions that are necessary to represent the data is limited by the size of the data.

We found that the Haar basis is the most appropriate basis for digital image processing of discrete pixel or voxel grids. The Haar basis uses box functions with finite support on multiple scales. A weighted linear combination of these basis functions constitutes the reconstructed image. By combining the detail coefficients of all dimensions, the contour of an object can be detected if the object exists in the selected part of the frequency spectrum. Using this method, high-frequency noise can be mostly eliminated, while other, typically larger objects remain visible.

re a mathematical tool for hierarchically decomposing functions in the frequency domain by preserving the spatial domain. Since their introduction [1], wavelets have found more and more applications in computer graphics, such as image compression, digital image processing, and feature detection [2-3]. Using wavelets, an image pyramid can be produced which represents the entropy levels for each frequency. In this case study, we demonstrate how this property can be exploited to segment objects in noisy images based on their frequency response in various frequency bands, separating them from the background and from other objects. We compare our noise-robust Haar wavelet-based technique to other standard image processing methods.

FEATURES OF THE SEGMENTATION TOOLBOX

For this case study, a toolbox has been developed to experiment with various segmentation techniques. The toolbox allows the user to combine traditional thresholding- or contour-based segmentation algorithms with a frequency-based method. The interface of the segmentation toolbox is shown in figure 1. It was implemented using QT, a cross-platform C++ graphical user interface (GUI) application framework [4].

The toolbox contains many digital image enhancement functionalities, including brightness and contrast adjustment, histogram expansion and equalization, Haar wavelet decomposition, filtering, morphological operations, etc. To facilitate the comparison of a modified image with the original one, the original image is shown in the upper-right corner. The histogram of an image provides its gray level distribution information, which is important for the parameter selection in several image enhancement operations. Therefore the histogram of an image under processing is displayed under the original image. The rest of the user interface contains buttons and sliders to choose from a variety of histogram, filtering and wavelet decomposition functions.

The input image can be transformed into several levels of resolution using 2-D Haar wavelets [2-3]. A low-pass filter transformation gives three sets of images at different resolution levels. A high-pass filter provides the detail information. If we choose the number of recursive transformation steps to be 3, then each set consists of three images as shown in figure 1. The horizontal sub images respond to different horizontal high frequencies in the image, whereas the vertical and diagonal sub images respond to different vertical and diagonal high frequencies in the image respectively. The low frequency part is displayed in the upper-left corner. This means that the image has been decomposed into three levels of directional derivatives, or three frequency bands in a preferred direction. From each decomposition step, the lowest frequency detail image (either horizontal, vertical, or diagonal) is chosen, scaled, and then combined with the other images of the set into a single image. The different frequencies are color-coded, using a particular primary color for each frequency level. In our example, the combination process works as follows: enlarge level 2 and 3 images to the size of level 1 images; combine the horizontal, vertical, and diagonal images separately, and apply one of the three additive primary colors (red, green, blue) to each frequency. This means that if an image feature is present in one frequency range, it will take on a particular color. If it is present in more than one frequency band, the resulting color is a combination of the three input colors. If a feature is visible in all three frequency bands, the object appears as white. It should be noted that this visual color comparison works well for three frequency bands mapped to three primary colors. If more than three frequency bands are analyzed, the color distinction becomes more difficult. In this case, and also in the case of three colors, an additional condition can be set so that only those features that appear in all frequency bands are visible (or white), while all other features are eliminated (black). Large features in noisy images are usually visible in all frequency bands and therefore displayed, while small features indistinguishable from noise are usually eliminated. This way, the desired features in a noisy image can be extracted, while noise is eliminated. The method works well in cases where standard threshold- or gradient-based edge detection algorithms would fail, because the noise would mistakenly be misinterpreted as a contour if it exceeds a certain absolute or gradient threshold. The sliders High, Medium, and Low, are used to set thresholds for each decomposition level, i.e., for each frequency band. Intensities below the threshold will be set to zero in the combination process, therefore gray levels less than the threshold will not contribute to the detected features in the combined image. This means that each frequency band can be selected separately, and if a feature is not present in one frequency band, it can be completely eliminated. By setting an appropriate threshold for each frequency band, we can effectively eliminate high frequency noise in an image and highlight dominant features, even if the feature itself is noisy.

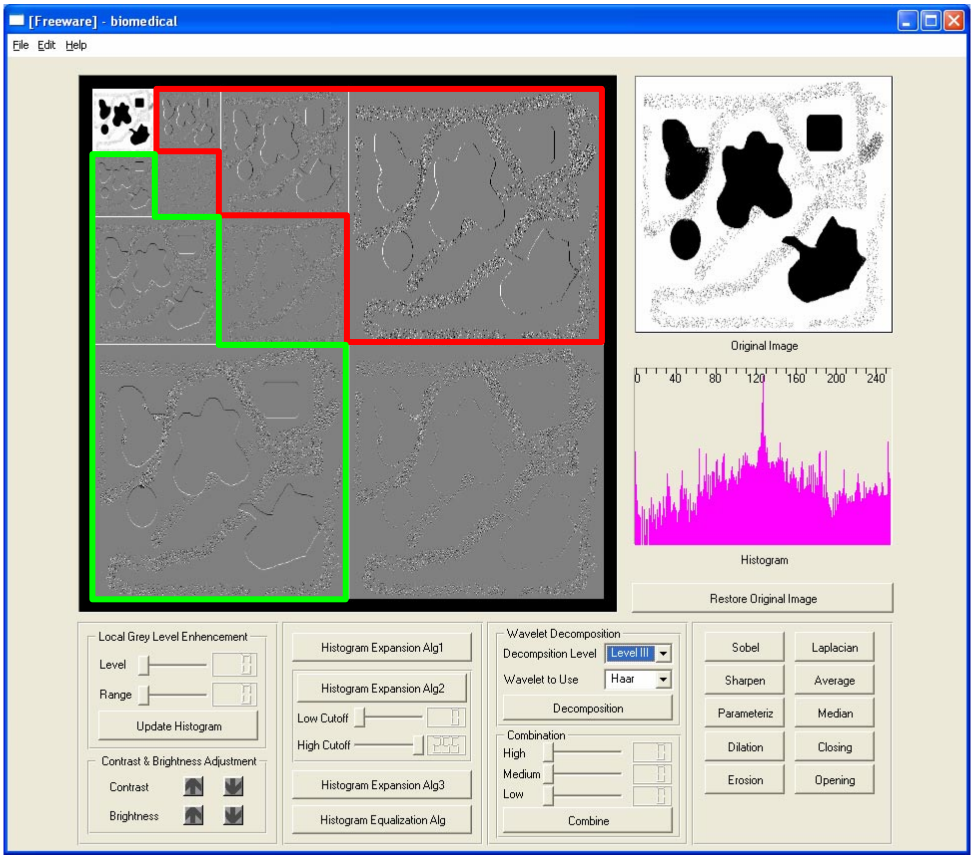

Figure 1 Interface of the segmentation toolbox.

RESULTS AND DISCUSSIONS

Figure 2 shows one of the segmentation results for a test image similar to the one shown in the GUI (figure 1). For comparison, results from Sobel and Laplacian filters are also shown in figure 2. The input data ranges from 0 to 255. The threshold setting for figure 2d is [0 0 0], which are the threshold values for High, Medium, and Low, respectively. It can be seen that all frequency components, from high to low, are present in the final image. The threshold setting for figure 2e is [0 0 10]. Now we can see a significant difference between figure 2a and figure 2f: the noise has been greatly reduced in this image, while the outlines of the features remain relatively unchanged. The threshold setting for figure 2f is [0 10 30]. In this image, the noise is totally eliminated, and only the desired feature contours that have a strong presence in all three frequency bands are visible.

Compared to low frequency components (the objects that are to be segmented in figure 2), high frequency components (visible as noise in figure 2) naturally have lower entropies at higher wavelet transformation levels, i.e., in the lower frequency bands. It is this property that we use to separate noise from object features in an image. Since noise typically has lower entropy in the low frequency bands, the noise can be easily eliminated by setting a threshold that is larger than the noise level. As we discussed before, entropy levels less than the threshold are set to zero, and do not contribute to the combined image. However, the effectiveness of the above method depends on the fact that large objects, i.e., low frequency components, even if the area in the image where they occur appears to be noisy, have relatively constant entropies at all wavelet transformation levels.

It can be seen from the images in figure 2 that though the high frequency noise can be completely eliminated, the useful features, i.e., the contours of the segmented objects, are sometimes also affected, and the algorithm does not necessarily produce closed contours. A closed contour would be essential for a flood-fill algorithm that separates the object from the rest of the image (segmentation). This is an issue that needs to be studied further. One way to solve the problem is to use morphological filters to close the discontinuous contours.

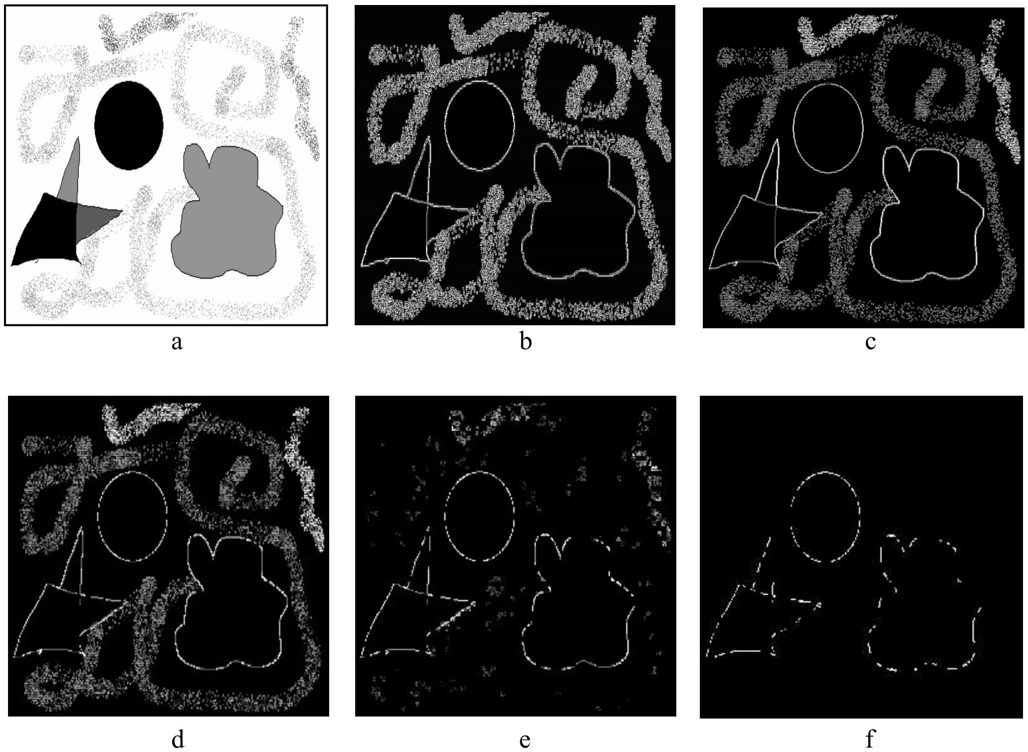

Figure 2 a) original image, b) Sobel filter result, c) Laplacian filter result, d)-f) combined images from different settings of thresholds [high medium low]; d) [0 0 0], e) [0 0 10], f) [0 10 30].

CONCLUSION

We investigated a new way of applying wavelets to digital image processing, and demonstrate that high frequency noise can be successfully removed without affecting contour detection by selectively combining the various frequency bands of wavelet-transformed images. The reliability of current contour-detection methods for image segmentation (e.g., Sobel, Laplace) can be improved by combining it in a weighted fashion with our wavelet-based decomposition and selective recombination technique. Using a higher weight on the wavelet-based method makes the segmentation more robust against interference from noise.

REFERENCES

[1] Mallat, S. 1987. "A compact multiresolution representation: the wavelet model." *Proc. IEEE Computer Society Workshop on Computer Vision,* IEEE Computer Society Press, Washington, D.C., p.2-7.

[2] Stollnitz, E. J.; Derose, T. D.; Salesin, D. H.; 1996. *Wavelets for Computer Graphics*. Morgan Kaufmann Publishers, Inc., San Francisco, CA.

[3] Gonzalez, R. C.; Woods, R. E. 2002. *Digital Image Processing*. Prentice Hall, Inc., Upper Saddle River, New Jersey.

[4] http://www.trolltech.com, official website of Trolltech Inc., the manufacturer of QT.