Available online at www.sciencedirect.com

ScienceDirect

Journal of Structural Biology 160 (2007) 93–102

Journal of Structural Biology

www.elsevier.com/locate/yjsbi

Bilateral edge filter: Photometrically weighted, discontinuity based edge detection

Radosav S. Pantelic [a], Geoffery Ericksson [b], Nicholas Hamilton [a], Ben Hankamer [a,*]

[a] *Institute for Molecular Bioscience, The University of Queensland, Brisbane, Qld 4072, Australia*
[b] *Queensland Brain Institute, The University of Queensland, Brisbane, Qld 4072, Australia*

Received 9 November 2006; received in revised form 12 July 2007; accepted 17 July 2007
Available online 25 July 2007

Abstract

Edge-detection algorithms have the potential to play an increasingly important role both in single particle analysis (for the detection of randomly oriented particles), and in tomography (for the segmentation of 3D volumes). However, the majority of traditional linear filters are significantly affected by noise as well as artefacts, and offer limited selectivity. The Bilateral edge filter presented here is an adaptation of the Bilateral filter [Jiang, W., Baker, M.L., Wu, Q., Bajaj, C., Chiu, W., 2003. Applications of a bilateral denoising filter in biological electron microscopy. J. Struct. Biol. 144, 114–122] designed for enhanced edge detection. It uses photometric weighting to identify significant discontinuities (representing edges), minimizing artefacts and noise. Compared with common edge-detectors (LoG, Marr–Hildreth) the Bilateral edge filter yielded significantly better results. Indeed data was of a similar quality to that of the Canny edge-detector, which is considered as a leading standard in edge detection [Basu, M., 2002. Gaussian-based edge-detection methods—a survey. IEEE Trans. Syst. Man Cybern. C Appl. Rev. 32, 252–260]. Compared to the Canny edge-detector the Bilateral edge-detector has the advantages that it only requires the adjustment of a single parameter, is theoretically faster for reasonably sized images, and can be used in selective contrast enhancement of images. The simplicity and speed of the filter for single particle and tomographic analysis are discussed.

Keywords: Electron microscopy; Electron cryo-microscopy; Cryo-electron microscopy; Single particle analysis; Tomography; Edge detection; Feature extraction; Contour tracing; Asymmetric; Noise; Filter; Contrast enhancement

1. Introduction

Edge-detection algorithms have the potential to play an increasingly important role both in single particle analysis (for the detection of randomly oriented particles), and in tomography (for the segmentation of 3D tomographic volumes). Edge-detection based particle-picking algorithms identify particles according to distinctive geometric properties interpreted from structural contours (Zhu et al., 2003; Adiga et al., 2004; Mallick et al., 2004; Yu and Bajaj, 2004; Adiga et al., 2005; Woolford et al., 2007a). Such algorithms are more suited to the detection of asymmetric structures than traditional cross-correlation based methods which require potentially unmanageable numbers of templates (Wong et al., 2004). Edge detection is also being increasingly used in the automated segmentation of tomograms (Baumeister, 2002; Volkmann, 2002; Bartesaghi et al., 2005) as an alternative to manual segmentation. Yet despite their potential, the accuracy and utility of edge-detection based algorithms is currently limited by their sensitivity to noise (Rath and Frank, 2004; Baumeister, 2002; Volkmann, 2002; Hall and Patwardhan, 2004).

Here we present the Bilateral edge filter, a fast, simple (single parameter) filter. It selectively minimizes noise artefacts during the extraction of structural data from low signal-to-noise ratio (SNR) SPA and electron tomography data during particle-picking and segmentation procedures.

* Corresponding author.
E-mail address: b.hankamer@imb.uq.edu.au (B. Hankamer).

1047-8477/$ - see front matter © 2007 Elsevier Inc. All rights reserved.
doi:10.1016/j.jsb.2007.07.005

To explain the concepts underlying the Bilateral edge filter, the nature of edge detection and its limitations are briefly outlined.

1.1. Pre-filtering

Pre-filtering is widely used in conjunction with edge detection to suppress noise, and impose a progressive continuity in pixel intensity within which significant discontinuities contributed by edges can be more accurately identified. Indeed most if not all segmentation and particle-picking applications incorporate pre-filtering at some level (Sherman et al., 1996; Banks et al., 2003; Rath et al., 2003; Roseman, 2003, 2004; Hall and Patwardhan, 2004; Huang and Penczek, 2004; Plaisier et al., 2004; Volkmann, 2004; Adiga et al., 2005; van der Heide et al., 2007; Woolford et al., 2007a). The degree of filtering (and consequent choice of filters) depends upon the type of information to be retrieved, and should therefore be considered separately from edge detection. Simple filters such as the Gaussian, mean and median have been commonly used, as well as more complex methods such as anisotropic diffusion (Adiga et al., 2005; Yu and Bajaj, 2004). However, choice of pre-filter should not detract from the speed, simplicity and therefore effective use of subsequent edge-detectors. Iterative median filtering has been shown to produce results on par with more complicated and slower anisotropic methods (van der Heide et al., 2007).

1.2. Edge-detection algorithms

An edge is by definition, a significant change in intensity along one or multiple axes. Discontinuity based detectors use this property to locate and demarcate edges within an image. The quality of a given edge-detector is defined by its ability to (Basu, 2002):

(1) Accurately detect edges masked by noise.
(2) Handle multiple edge orientations (directionality).
(3) Minimize localization errors (shifting of detected edges from their actual position).
(4) Reduce interference (where inaccurate localization of a given edge interferes with the localization and detection of other edges).
(5) Ensure accurate contour connectivity.
(6) Avoid the formation of artefacts such as double edges and closed loops.

One of the most widely used edge-detection filters is the Laplacian of Gaussian (LoG), which is based on a convolution of the Gaussian (smoothing) and Laplacian (discontinuity detection) kernels. Edges are traced at zero-crossings in second derivative space, which are defined by the junctions of positive and negative sign-values corresponding to each pixel. An extension of the LoG filter involving sign-value thresholding prior to zero-crossing tracing (Gonzalez, 2002), improves the localization accuracy of the filter, and also minimizes interference and artefacts (Basu, 2002). In a more recent paper, an arbitrary zero-crossings approach was developed to allow the dynamic adjustment of contours at values that differ from zero. In essence this allows the user to expanded or shrink a contour to facilitate dynamic contouring (Woolford et al., 2007b).

Selective edge detection and suppression of noise has been traditionally achieved by varying the filter size. Small filter sizes preserve high-resolution detail but consequently include inhibitive amounts of noise, while larger sizes effectively suppress noise but only preserve low-resolution detail (Basu, 2002). Multi-scale edge detection was proposed as a solution for effective noise reduction, and involves the combination of edge maps generated at different filter sizes (Basu, 2002). The Marr–Hildreth filter is such a method. It is based on the assumption that true edges are indicated by the "coincidence" of zero-crossings between different filter sizes (Marr and Hildreth, 1980). The fundamental principals of the Marr–Hildreth approach can be approximated by combining zero-crossings from different LoG filter sizes. However, the underlying assumption only holds for a limited number of closely distributed filter sizes in practice (since the edge localization shifts between filter sizes) (Basu, 2002).

The Canny edge-detector (Canny, 1986) is a more accurate, though complex filter. It requires the optimization of three independent parameters controlling Gaussian smoothing and the tracing of edges from a generated contour gradient map using non-maximal peak suppression (according to two thresholds which control connectivity (hysteresis)). The filter's failure to check whether discontinuity is contributed by noise or true edges, make it susceptible to spurious and unstable boundaries (Basu, 2002). However, this can be compensated for by the correct adjustment of hysteresis.

Anisotropic diffusion is a method of smoothing, which facilitates noise reduction *within* and not *between* regions. This has the effect of smoothing areas of background while maintaining the boundaries (edges) delineating adjacent regions (Bajaj and Xu, 2003). The Bilateral filter is an adaptation of anisotropic diffusion and has been shown to be more accurate than, purely linear spatial filters such as the Gaussian (Jiang et al., 2003). As its name suggests, the Bilateral edge filter presented here is an adaptation of the original Bilateral filter (Jiang et al., 2003) to edge detection. The filter is capable of a level of accuracy on par with that of the Canny edge-detector through the adjustment of a single (photometric) parameter rather than three (Canny). By distinguishing between discontinuities contributing noise and those representing edges, the Bilateral filter also complements the Canny by addressing one of its known issues (Basu, 2002). The selectivity and simplicity of the filter make it easy to operate and well suited to the tracing of randomly oriented single particles and tomographic slices at potentially higher speeds for a reasonably sized image.