

ANNIHILATION-DRIVEN LOCALISED IMAGE EDGE MODELS

Hanjie Pan*, Thierry Blu[†], and Martin Vetterli*

*School of Computer and Communication Sciences
Ecole Polytechnique Fédérale de Lausanne (EPFL)
CH-1015 Lausanne, Switzerland

[†]Electronic Engineering Department
The Chinese University of Hong Kong
Shatin, New Territories, Hong Kong SAR

ABSTRACT

We propose a novel edge detection algorithm with sub-pixel accuracy based on annihilation of signals with *finite rate of innovation* [1, 2]. We show that the Fourier domain annihilation equations can be interpreted as spatial domain multiplications. From this new perspective, we obtain an accurate estimation of the edge model by assuming a simple parametric form within each localised block. Further, we build a locally adaptive global mask function (i.e., our edge model) for the whole image. The mask function is then used as an edge-preserving constraint in further processing. Numerical experiments on both edge localisations and image up-sampling show the effectiveness of the proposed approach, which outperforms state-of-the-art method.

Index Terms— Annihilation equations, sub-pixel edge detection, image up-sampling

1. INTRODUCTION

Edge detection is one of the most intensively studied problems in low-level computer vision, where the goal is to detect sudden changes (or “discontinuities”) present in an image. Numerous approaches have been proposed to detect edges since 1960’s. We refer the readers to [3, 4] for extensive reviews on various edge detection algorithms. The Canny edge detector [5] is probably the most widely used method because of its robustness against noise and the consistent performance in giving one-pixel thin detection results. However, the detected edge positions are limited to pixel accuracy.

Previous efforts to achieve sub-pixel detection accuracy have led to the developments of model based fitting algorithms [6, 7, 8, 9, 10], where the optimal edge is given by minimising the least square errors between the image and an edge model (typically a step function). It has been shown that sub-pixel edge position of an ideal step edge can be derived from the spatial moments of a discrete image [9, 11, 12, 13].

This work was supported in part by an ERC Advanced Grant – Support for Frontier Research – SPARSAM Nr: 247006, and in part by the General Research Fund Nr: CUHK14200114 from the Hong Kong Research Grant Council.

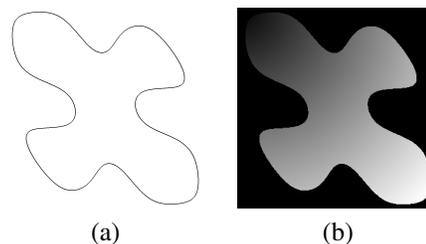


Fig. 1. (a) A curve with finite rate of innovation (1) and (b) its associated edge image I_C , which is analytic everywhere except on the predefined curve.

In this paper, we propose a sub-pixel accurate edge detection algorithm, which is motivated by the work on annihilation for curves with finite rate of innovation (FRI) [2]. By assuming a very simple linear edge model within each block, we obtain edge models efficiently by minimising the *annihilation errors*. We will show that this error is simply the product between the derivative image (e.g., image gradient) and a mask function (which contains our edge modelization).

In a sense, the annihilation-driven approach may be judged similar to model based algorithms [7, 8] with the philosophical difference: our goal is not to identify the large gradients (where edges are usually located) by fitting the image with an edge model but rather reconstruct a mask function that kills (i.e., annihilates) the gradients.

We generalise the same idea to build a continuous domain global edge model for the entire image, which may subsequently be used to assign position-dependent weights in further processing. We show through numerical simulations, that the global edge model is indeed beneficial in preserving sharp edges in image up-sampling and outperforms state-of-the-art algorithm.

2. MOTIVATION: FRI CURVE ANNIHILATION

Recent work in [2] extends the traditional FRI sampling and reconstruction framework [1, 13, 14, 15, 16] to a specific class of curves, which are defined implicitly as the roots of certain

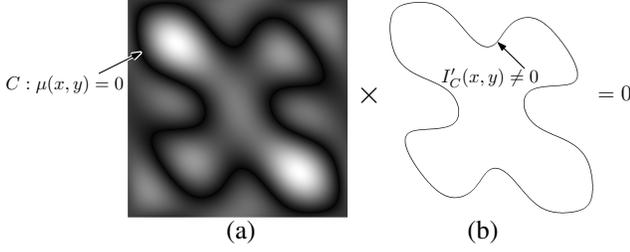


Fig. 2. Spatial domain interpretation of the annihilation equation (2). The mask function $\mu(x, y)$, whose roots defines the curve, annihilates whatever is different from zero in the derivative image I'_C .

mask function (Fig. 1(a)):

$$C : \underbrace{\sum_{k=-K_0}^{K_0} \sum_{l=-L_0}^{L_0} c_{k,l} e^{j \frac{2\pi k}{\tau_x} x + j \frac{2\pi l}{\tau_y} y}}_{\text{mask function: } \mu(x,y)} = 0, \quad \text{for } \begin{cases} 0 \leq x < \tau_x, \\ 0 \leq y < \tau_y. \end{cases} \quad (1)$$

Here τ_x and τ_y are some positive real numbers that specify the periods along the x and y directions, respectively. The curve is uniquely specified by the set of coefficients $c_{k,l}$, which is of finite dimension. Hence, the family of curves also has finite rate of innovation.

For each curve defined implicitly this way, we have an associated edge image, which is analytic almost everywhere except on the curve where it becomes discontinuous (Fig. 1(b)). It can be shown that the Fourier transform of the edge image satisfies the annihilation equation [2]:

$$\sum_{k=-K_0}^{K_0} \sum_{l=-L_0}^{L_0} c_{k,l} \hat{I}_C \left(\omega_x - \frac{2\pi k}{\tau_x}, \omega_y - \frac{2\pi l}{\tau_y} \right) = 0, \quad \forall \omega_x, \omega_y, \quad (2)$$

where $\hat{I}_C = (\omega_x + j\omega_y) \hat{I}_C(\omega_x, \omega_y)$. Observe that the annihilation equation (2), which is a Fourier domain convolution, is simply a multiplication in the spatial domain:

$$\sum_{k,l} c_{k,l} \hat{I}_C \left(\omega_x - \frac{2\pi k}{\tau_x}, \omega_y - \frac{2\pi l}{\tau_y} \right) \xrightarrow{\mathcal{F}^{-1}} \mu(x, y) \cdot I'_C(x, y).$$

Here $I'_C \stackrel{\text{def}}{=} \left(\frac{\partial}{\partial x} + j \frac{\partial}{\partial y} \right) I_C$ is the derivative image, which is different from zero only on the curve C where the edge image I_C is discontinuous (Fig. 2).

Thus, we may treat the function $\mu(x, y)$ as a “mask”, which will automatically *annihilate* whatever is different from zero in the derivative image. Thanks to the spatial domain interpretation, we can apply the annihilation idea to cases where the curve model (1) is not satisfied exactly. For instance, we may obtain a curve model approximately by minimising the annihilation errors in the spatial domain directly:

$$\min \| \mu \cdot I'_C \|, \quad (3)$$

where $\| \cdot \|$ is a norm to be specified (e.g., ℓ_2). Note that here we are no longer restricted to the mask parametrisation (1) in general.

In our previous work [2], it was difficult to generalise the exact FRI curve framework to more practical cases, e.g., with natural images. The curve model for image edges were obtained by minimising the Fourier domain annihilation errors (2) in the least square sense. Because (1) is a *global* model, we had little control over the geometry of the reconstructed curve model. The model size (i.e., K_0 and L_0) was specified in an ad hoc manner to achieve a good balance between representing image edges and not over-fitting the given data. This motivated the development of a new curve model (or equivalently the associated mask function) that is locally adaptive.

3. SUB-PIXEL EDGE DETECTION

If we magnify an image by a sufficiently large factor, then we can well approximate image edges with a straight line: $ax + by + c = 0$, where a , b and c are some (unknown) parameters and x , y are the horizontal and vertical coordinates. In the annihilation framework, we can treat the edge model as the roots (i.e., the zero-crossings) of a mask function:

$$\mu(x, y) = ax + by + c.$$

Based on the spatial domain interpretation in the previous section, the best linear edge model is therefore the one such that the annihilation error is minimised (3).

The simple line approximation only makes sense when we consider the problem in a local region. Specifically, the edge parameters (a, b, c) within a block centred around (x_0, y_0) is:

$$\min_{a,b,c} \iint |\mu(x, y) \cdot I'(x, y)|^2 w(x - x_0, y - y_0) dx dy \quad (4)$$

subject to $a^2 + b^2 = 1$,

where $w(\cdot, \cdot)$ is a localised window, e.g., an isotropic Gaussian with a certain standard deviation σ : $\frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2+y^2}{2\sigma^2}\right)$.

The quadratic constraint is to avoid the trivial solution where $a = b = c = 0$. In practice, we may locate the blocks around the Canny edge points and repeat the same process in a block-wise manner to obtain an edge model for the entire image. Notice that (4) is a simple quadratic minimisation subject to one quadratic constraint. The associated Lagrangian is

$$\mathcal{L}(\mathbf{d}, \lambda) = \mathbf{d}^H \mathbf{A} \mathbf{d} + \lambda(1 - \mathbf{d}^H \mathbf{B} \mathbf{d}),$$

where

$$\mathbf{d} = \begin{bmatrix} a \\ b \\ c \end{bmatrix} \quad \text{and} \quad \mathbf{B} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Here \mathbf{A} is a 3×3 Hermitian symmetric matrix, whose entries can be obtained directly from the derivative image I' with *linear filtering*. By setting the derivatives of the Lagrangian to zero, (4) reduces to an eigendecomposition problem:

$$\frac{\partial \mathcal{L}(\mathbf{d}, \lambda)}{\partial \mathbf{d}} = 2(\mathbf{A} - \lambda \mathbf{B})\mathbf{d} = \mathbf{0},$$

whose closed form solution is

$$\mathbf{d} = \begin{bmatrix} A_{1,3}(A_{2,2} - \lambda) - A_{1,2}A_{2,3} \\ (A_{1,1} - \lambda)A_{2,3} - A_{1,3}A_{1,2} \\ A_{1,2}^2 - (A_{1,1} - \lambda)(A_{2,2} - \lambda) \end{bmatrix},$$

where $A_{i,j}$ for $i, j = 1, 2, 3$ are entries of \mathbf{A} . The Lagrange multiplier λ are solutions of a quadratic equation¹

$$\alpha_1 \lambda^2 + \alpha_2 \lambda + \alpha_3 = 0, \quad (5)$$

where $\{\alpha_i\}_{i=1,2,3}$ are constants that can be computed directly from \mathbf{A} 's entries.

Once the line parameters \mathbf{d} are reconstructed, we can further define an edge point within each block. Specifically, an edge point should be located on the reconstructed line while having the shortest distance to the block centre:

$$(x_e, y_e) = \arg \min_{x,y} (x - x_0)^2 + (y - y_0)^2 \\ \text{subject to } \mu(x - x_0, y - y_0) = 0,$$

i.e., $x_e = x_0 - \frac{ac}{a^2+b^2}$, and $y_e = y_0 - \frac{bc}{a^2+b^2}$. Note that the edge model within each block is directly related to the entries of \mathbf{A} , which are given by filtering the derivative image I' . In the implementation, we can reconstruct the mask parameters for the entire image efficiently with a few 2D convolutions.

It is possible to further refine the sub-pixel accurate edge detection by adapting the shape of the window based on the estimated edge model, e.g., using an anisotropic Gaussian window that is aligned with the edge orientation.

4. A LOCALLY ADAPTIVE GLOBAL MASK

The localised linear edge model, which gives reliable estimations of the edge slope and position, may already be useful in its own right (see an example in Section 5.1). But is it possible to combine these local edge models and build a single *global* mask function? We may then use the global mask to assign position-dependent weights in subsequent image processing [2], e.g., image up-sampling.

A naive approach to build a global mask function is to sum up the linear edge models over all blocks:

$$\mu(x, y) = \sum_{n=1}^N (a_n(x - x_n) + b_n(y - y_n) + c_n)w(x - x_n, y - y_n), \quad (6)$$

where (a_n, b_n, c_n) are the line coefficients obtained within each of the N blocks from (4) and $w(\cdot, \cdot)$ is an isotropic Gaussian window. The straightforward approach suffers from two problems: (i) The global mask function (6) vanishes in positions that are *not* in the neighbourhood of any edge points (x_n, y_n) ; (ii) The coefficients (a_n, b_n, c_n) , which are obtained without considering inter-block interferences, are suboptimal in terms of minimising the annihilation error $\|\mu \cdot I'\|^2$ for the global mask function (6).

¹The two solutions λ_{\min} and λ_{\max} correspond to the cases where the objective function is minimised and maximised, respectively.

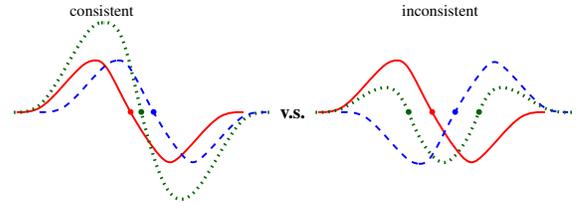


Fig. 3. 1D illustration of the mask functions in the adjacent blocks (solid red and dashed blue curve). Dotted green curve is the combined mask function. We would like the adjacent blocks to be “sign-consistent” and avoid cancelling each other.

4.1. Mask Parametrisation

One way to address the issue that the mask function reduces to zero in smooth areas, is to choose the line coefficients as $a = 0$, $b = 0$, and $c = 1$ *explicitly* for blocks in smooth areas (as we do not expect to find any edges in smooth areas):

$$\mu(x, y) = \sum_{n=1}^N (a_n(x - x_n) + b_n(y - y_n) + c_n)w(x - x_n, y - y_n) \\ + \underbrace{\sum_{m=1}^M w(x - x_m, y - y_m)}_{\mu_{\text{smooth}}(x, y)}. \quad (7)$$

Here we may choose (x_n, y_n) as the coordinates of the Canny edge pixels; while (x_m, y_m) correspond to the non-edge pixels. Note that μ_{smooth} is independent of the edge parameters (a_n, b_n, c_n) for $n = 1, \dots, N$. The total degrees of freedom of the global mask function is still determined by the number of edge points.

4.2. Global Optimisation

Recall that the objective function $\|\mu \cdot I'\|^2$ is a quadratic function of the edge parameters $\{a_n, b_n, c_n\}_{n=1}^N$. If we follow the same formulation as in the block-wise process (4), we would like to minimise the annihilation error subject to the quadratic constraint $a_n^2 + b_n^2 = 1$ for each of the N constituent blocks. The resultant optimisation problem is a quadratic programming with N quadratic constraints (QCQP). In practice, N is usually very large (in the order of 10^3), which precludes finding the closed form solution in general.

Alternatively, we may replace the quadratic constraints with a set of *linear* ones, e.g., $c_n = 1$ for $n = 1, \dots, N$. One potential issue is that the adjacent blocks might end up having the opposite signs over the same area, thus cancelling each other (as is illustrated in Fig. 3). To alleviate this issue, we can use the block-wise estimation (4) in the n -th block c_n^{prev} as a reference solution:

$$\min_{\substack{a_1 \dots a_N \\ b_1 \dots b_N \\ c_1 \dots c_N}} \iint |\mu(x, y) \cdot I'(x, y)|^2 dx dy \\ \text{subject to } \text{sign}(c_n^{\text{prev}})c_n = 1 \quad \text{for } n = 1, \dots, N,$$

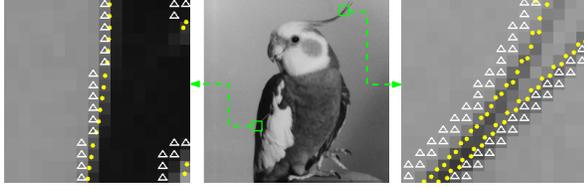


Fig. 4. Comparisons of the edge localization results with Canny edge detector [5] (white triangles) and the proposed annihilation-driven approach (yellow dots).

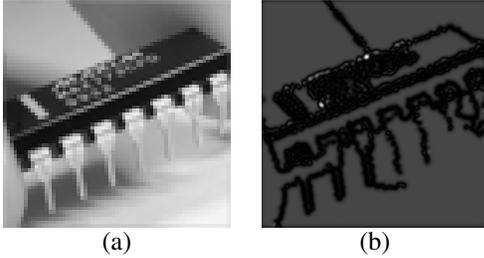


Fig. 5. Global mask function reconstructed from the given low-resolution image. (a) Low-resolution image (size: 65×65). (b) Continuous domain mask function (plot size: 1000×1000).

which has a closed form solution.

5. EXPERIMENTAL RESULTS

5.1. Edge Detection

We applied the sub-pixel accuracy edge detection algorithm in Section 3 to grey scale images. The detected edge positions with the Canny edge detector and the annihilations-driven approach are compared in Fig. 4. We have achieved sub pixel accurate edge detections with the proposed algorithm.

5.2. Application to Image Up-sampling

We tested the effectiveness of the global mask function in preserving image edges with image up-sampling. In experiments, the ground truth high resolution image is first low-pass filtered with bicubic filter and down-sampled by a factor of 3. The global edge model is then estimated from the *low-resolution* image (Fig. 5). Experimentally, we found that a more accurate edge model can be obtained when the algorithm is applied to an image that is interpolated from the given low-resolution one, e.g., the solution of (9). One explanation is that more edge control points, i.e., larger N in (7), are allowed in that case.

The global mask is then enforced softly as a regularisation term in the up-sampling process [2]:

$$\begin{aligned} \min_I \quad & \|\Delta I\|_2^2 + \ell \overbrace{\|\mathbf{MDI}\|_1}^{\text{annihilation}} \\ \text{subject to} \quad & \Phi I = I_{LP} \end{aligned} \quad (8)$$

Here I is the up-sampled image and I_{LP} is the given low-resolution image; Δ and \mathbf{D} are convolution matrices asso-

Table 1. Comparisons of different up-sampling images with the ground truth high-resolution image (PSNR).

Image	w/o annihilation constraint (9)	w/ annihilation constraint (8)	learning-based algorithm [17]
<i>chip</i>	27.61dB	28.95dB	28.29dB
<i>peppers</i>	29.62dB	30.95dB	30.43dB
<i>bank</i>	22.33dB	23.12dB	22.91dB
<i>MRI</i>	25.25dB	25.85dB	25.40dB

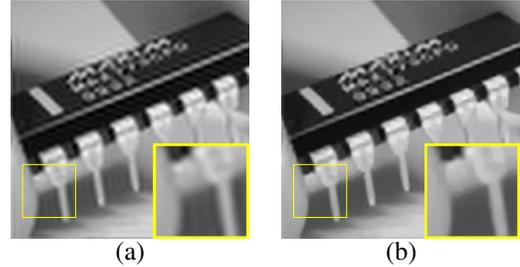


Fig. 6. Comparisons of different image up-sampling results (up-sampling factor: 3). (a) Up-sampled image without annihilation constraint (PSNR = 27.61dB). (b) Up-sampled image with annihilation constraint (8) (PSNR = 28.95dB).

ciated with the discrete Laplacian filter and the first order derivatives respectively: \mathbf{M} is a diagonal matrix with diagonal entries specified by the mask function (7); Φ represents the sampling process (low-pass filtering and down-sampling). We have chosen the regularisation weight $\ell = 9$ for all the images. Table 1 summarises the PSNR's between the up-sampled images and the ground truth high resolution images. In our comparisons, we have included one state-of-the-art learning based algorithm [17] as well as a standard approach, which minimises the smoothness regularisation subject to the data-fidelity constraint only:

$$\begin{aligned} \min_I \quad & \|\Delta I\|_2^2 \\ \text{subject to} \quad & \Phi I = I_{LP} \end{aligned} \quad (9)$$

With the annihilation constraint, the up-sampled images have sharper edges and much less ringing artifacts (Fig. 6).

6. CONCLUSION

We proposed an efficient edge detection algorithm based on the annihilation idea. By assuming a simple linear edge model within each block, we achieved sub-pixel accuracy edge detections. We also generalised the same idea to build a global edge model for the whole image, which was shown to be beneficial in preserving sharp edges.

The current algorithms rely on the Canny edge detector to indicate where an edge point should locate approximately. In the future work, we may make use of the estimated edge information, e.g., the ratio between the two eigenvalues from (5), to decide whether an edge block or a smooth area block should be placed for a particular position in (7).

7. REFERENCES

- [1] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1417–1428, 2002.
- [2] H. Pan, T. Blu, and P. L. Dragotti, "Sampling curves with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 62, no. 2, pp. 458–471, 2014.
- [3] V. Torre and T. A. Poggio, "On edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 2, pp. 147–163, 1986.
- [4] D. Ziou and S. Tabbone, "Edge detection techniques—An overview," *Pattern Recognition and Image Analysis*, vol. 8, pp. 537–559, 1998.
- [5] J. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, pp. 679–698, 1986.
- [6] R. Kakarala and A.O. Hero, "On achievable accuracy in edge localization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 7, pp. 777–781, Jul 1992.
- [7] K. Jensen and D. Anastassiou, "Subpixel edge localization and the interpolation of still images," *IEEE Transactions on Image Processing*, vol. 4, no. 3, pp. 285–295, 1995.
- [8] P. J. L. van Beek, *Edge-based image representation and coding*, Ph.D. thesis, Delft University of Technology, 1995.
- [9] L. Baboulaz and P. L. Dragotti, "Exact feature extraction using finite rate of innovation principles with an application to image super-resolution," *IEEE Transactions on Image Processing*, vol. 18, no. 2, pp. 281–298, 2009.
- [10] C. Chen, P. Marziliano, and A. C. Kot, "2D finite rate of innovation reconstruction method for step edge and polygon signals in the presence of noise," *IEEE Transactions on Signal Processing*, vol. 60, no. 6, pp. 2851–2859, 2012.
- [11] A. J. Tabatabai and O. R. Mitchell, "Edge location to subpixel values in digital imagery," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 188–201, 1984.
- [12] E. P. Lyvers, O. R. Mitchell, M. L. Akey, and A. P. Reeves, "Subpixel measurements using a moment-based edge operator," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 12, pp. 1293–1309, 1989.
- [13] P. Shukla and P. L. Dragotti, "Sampling schemes for multidimensional signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 55, no. 7, pp. 3670–3686, 2007.
- [14] I. Maravić and M. Vetterli, "Exact sampling results for some classes of parametric nonbandlimited 2-D signals," *IEEE Transactions on Signal Processing*, vol. 52, no. 1, pp. 175–189, 2004.
- [15] T. Blu, P. L. Dragotti, M. Vetterli, P. Marziliano, and L. Coulot, "Sparse sampling of signal innovations," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 31–40, 2008.
- [16] J. Berent, P. L. Dragotti, and T. Blu, "Sampling piecewise sinusoidal signals with finite rate of innovation methods," *IEEE Transactions on Signal Processing*, vol. 58, no. 2, pp. 613–625, 2010.
- [17] R. Zeyde, M. Elad, and M. Protter, "On single image scale-up using sparse-representations," *Curves and Surfaces*, vol. 6920, pp. 711–730, 2012.