

Stochastic Extraction of Elongated Curvilinear Structures With Applications

Vladimir A. Krylov and James D. B. Nelson, *Member, IEEE*

Abstract—The automatic extraction of elongated curvilinear structures (CLSs) is an important task in various image processing applications, including numerous remote sensing, and biometrical and medical problems. To address this task, we develop a stochastic approach that relies on a fixed-grid, localized Radon transform for line segment extraction and a conditional random field model to incorporate local interactions and refine the extracted CLSs. We propose several different energy data terms, the appropriate choice of which allows us to process images with different noise and geometry properties. The contribution of this paper is the design of a flexible and robust elongated CLS extraction framework that is comparatively fast due to the use of a fixed-grid configuration and fast deterministic Radon-based line detector. We present several different applications of the developed approach, namely: 1) CLS extraction in mammographic images; 2) road networks extraction from optical remotely sensed images; and 3) line extraction from palmprint images. The experimental results demonstrate that the method is fairly robust to CLS curvature and can accurately extract blurred and low-contrast elongated CLS.

Index Terms—Curvilinear structure, line extraction, localized Radon transform, conditional random field, mammogram, road extraction, palmprint.

I. INTRODUCTION

THE task of automatic detection of straight lines and, more generally, curvilinear structures (CLS) in images is one of the fundamental problems in image processing and pattern recognition. Typically, CLS extraction is required in specialized line detection problems such as road detection in remote sensing [1], [2] or vessel detection in medical imaging [3]. In some image processing problems CLS detection is a necessary preprocessing stage when complex objects are of interest that comprise distinct combinations of linear features, such as spicule pattern analysis for cancer/mass detection in mammography [4].

Recent decades have seen much interest in the design of line and CLS extraction methods. Several notable multi-scale

filtering techniques and directionally adaptive transforms derived from multi-scale analysis, including steerable filters [5], contourlets [6], [7], directionlets [8], have been proposed for local feature extraction and can be employed to address contours and CLS detection. Furthermore, multi-scale local transforms like beamlets [9], feature-adapted beamlets [10], ridgelets [11] and curvelets [12], have been specifically designed to address the problem of line extraction. Other approaches to line and CLS extraction include isotropic non-linear filtering for wide line detection [13], line profile analysis for CLS detection [14], and skeleton extraction [15]. These methods, which perform well for high-contrast CLS, rely on substantially different models and assumptions.

In this paper we aim to address a general scenario of content-rich images, i.e. where the CLS of interest appear on some non-trivial background. Furthermore, we assume the presence of additional noise factors such as, for instance, those due to the imaging modality, poor illumination and weak contrast between the CLS and non-CLS structures. Indeed, in most natural images the CLS of interest suffer from various kinds of noise and artifacts such as blur, occlusions, and low-contrast. In some cases such as mammographic images the strong blur transforms CLS from lines into ridges. More generally, from the point of view of CLS extraction we consider any non-CLS structure in the image as noise. In the absence of reliable geometrical properties of the considered lines we assume the characteristic property of CLS to be: first, their higher average brightness as compared to the background, and, second, their local linearity. This ridge-like assumption does not hold for an arbitrary kind of CLS, but for numerous line extraction applications this assumption is true and allows the recovery of line structures in the absence of strong geometrical and/or edge-based evidence.

Our work is inspired by the promising results achieved by Radon transform for digital line detection in mammographic [4] and remote sensing [16] applications. Since most of the CLS of interest demonstrate a certain degree of curvature and are considerably shorter than the entire image size, we employ a localized version of the Radon transform similarly to [17]. We perform image partitioning, required for the localized transform, by defining a grid of overlapping rectangular image regions of predefined size. Radon transform is applied separately on each of the grid regions. Alternatively, stochastic geometry approaches allow the analysed regions to change in size and location in a random manner [18], [19]. However, although they offer more flexibility and are free

Manuscript received March 15, 2014; revised August 16, 2014; accepted September 29, 2014. Date of publication October 15, 2014; date of current version November 6, 2014. This work was supported by the U.K. Engineering and Physical Sciences Research Council under Grant EP/J010081/1. V. Krylov has contributed to this work during his term at the Dept. of Statistical Science, University College London. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Wai-Kuen Cham.

V. A. Krylov is with the Department of Electrical, Electronic, Telecommunications Engineering and Naval Architecture, University of Genoa, Genoa 16145, Italy (e-mail: vladimir.krylov@unige.it).

J. D. B. Nelson is with the Department of Statistical Science, University College London, London WC1E 6BT, U.K. (e-mail: j.nelson@ucl.ac.uk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2014.2363612




Fig. 1. Results of Canny edge detection [22] (second row), multi-scale CLS profile extraction [14] (third row) and contourlet shrinkage [7] (fourth row) on mammographic (left), road network (centre), palmprint (right) images.

of undesirable cross-border effects, such approaches require elaborate parametric fine-tuning and typically result in significantly higher computational complexity.

We propose a two-step stochastic approach for the extraction of elongated CLS from images of natural and man-made scenes affected by noise, blur, acquisition artifacts and the presence of non-curvilinear structures. The aim of the first step is to extract a wide set of line segments including the CLS of interest. To this end we apply a localized version of the Radon transform to each of the image regions over an overlapping square grid and extract its first several maxima as line segment candidates. In the second step, we extract continuous chains of line segments by imposing a Conditional Random Field (CRF) structure of local grid dependencies. Optimization is performed stochastically via simulated annealing [20] using a Markov chain Monte Carlo (MCMC) algorithm [20], [21]. We improve performance of the method by parallelizing MCMC implementation via grid partitioning.

The developed stochastic approach is experimentally validated in three distinct applications. First, we consider the problem of CLS extraction from mammographic images. This is a key step for mammogram coregistration, mass detection, and other such medical image processing tasks. As can be seen in Fig. 1 (first column) the CLS that are present in this type of imagery suffer from high levels of blur due to volume projections and tissue density. Consequently, a standard edge detector (Canny's method [22]), CLS extractor

(Steger's method [14]) and hard shrinkage in the contourlet domain [7] fail to accurately identify the CLS locations. The second considered application is road-network extraction from high resolution remotely-sensed optical images. This application plays a key role in urban and agricultural mapping, coregistration, etc. A significant challenge in these type of images is the presence of non-road CLS and geometrical structures such as buildings, fields, tree canopies etc. The characteristic property of roads that allows their extraction on such a background is their continuity and elongatedness. It is immediate from Fig. 1 (second column) that standard techniques do not take these properties into account and therefore yield unsatisfactory results. More specifically, Canny's edge detector suffers from overdetection, and Steger's method results in an excess of discontinuities in the detected road segments. As the third application we consider the palm line extraction from palmprint images. This task is central in biometric identification and matching of palm images. Fig. 1 (third column) illustrates that, whereas Canny's detector is not appropriate for the CLS extraction on these type of images and contourlet shrinkage results in overdetection, Steger's CLS extractor achieves promising results (albeit with some need for further refinement). These three applications present essentially different CLS extraction cases with various kinds of noise.

The contributions of the paper are the design of a general and flexible framework for elongated CLS extraction complemented with MCMC parallelization, and its extensive experimental validation. The designed extraction technique (i) uses a CRF design on a fixed grid to significantly reduce the MCMC complexity and optimization cost and (ii) forms a flexible framework whereby the selection of distinct application-specific unary terms facilitate CLS extraction in images affected by a variety of noise factors, such as blur, occlusion, low-contrast, and background clutter effects. Our approach is significantly faster than full stochastic methods that consider random objects at random locations [16], [18]. A MATLAB implementation of the developed approach is available at the webpages of the authors.

The paper organisation is as follows. In Section II we present the line segment extraction step. In Section III we introduce the CRF structure, MCMC optimization procedure and present the potential and unary energy terms. In Section IV we give the outline of the proposed CLS extraction procedure and MCMC parallelization. In Section V we present the extensive experimental validation on three types of imagery, and in Section VI we summarize the conclusions of this study.

II. LINE SEGMENT DETECTION

The first stage of the proposed CLS extractor establishes an exhaustive set of line segments. The objective at this stage is to detect all of the constituent line segments of all of the genuine CLS present in the image. Hence, this initial step returns many false positives which must be refined in the subsequent stages of processing.

Both the Radon transform and its discrete counterpart, the Hough transform, are popular and efficient tools for the extraction of linear structures [4], [16], [23]. Nevertheless, they




Fig. 2. Orientation-driven clique selection in the 3-by-3 neighbourhood on an overlapping localized grid.

suffer from two major drawbacks when applied to CLS/line extraction in real imagery. Firstly, the standard transforms only address straight line detection,¹ whereas most of the CLS structures present in real images demonstrate a certain (and unknown) degree of curvature. Secondly, these global transforms are applied to the whole image and, thus, naturally favour longer lines over shorter ones. The first drawback results in either complete or partial loss of curved lines, whereas the second restricts the detection solely to longer segments. In order to overcome these shortcomings we choose to employ the localized Radon transform on a fixed grid. In this way, shorter lines receive the same treatment as longer lines and the curved structures can be approximated by a set of shorter line segments.

The continuous Radon transform on a 2D domain \mathbf{X} is defined by an integral of a function $f(x_1, x_2)$ over a straight line defined by ρ - its distance from the origin (intercept) and θ - the angle its normal vector makes with the positive X_1 -axis [23] (which differs from the slope angle by 90 degrees):

$$Rf(\rho, \theta) = \int_{\mathbf{X}} f(x_1, x_2) I[\rho = x_1 \cos \theta - x_2 \sin \theta] dx_1 dx_2, \quad (1)$$

where $I[\cdot]$ is the indicator function. When applied to digital images, the discrete Radon transform is better known as the Hough transform [23]. It sums the intensities along a specific angle θ and intercept ρ [23] in a finite image region.

Since the local Radon transform is implemented over a fixed grid some loss of translational invariance is incurred. This is most severe in situations where a line, of size similar to that of the grid region, is cut in two halves by one of the boundaries. Indeed, in the presence of sufficient noise, the successful detection of such a line will depend entirely upon its position. An overlapping fixed grid, employed here, partially overcomes this problem by allowing the points close to the boundaries to appear in several distinct regions of the grid, cf. overlapping grid structure on Fig. 2. The grid scale should be chosen with respect to the expected minimal size (length) of the CLS present in the image. Note that the finer the scale, the higher the allowed curvature of the detected structures, but also the higher the sensitivity to noise.

Each node of the grid corresponds to a square region in the image over which the discrete Radon transform is taken. As illustrated in Fig. 3(a), the standard Radon transform (1)

¹A generalized Radon transform has been proposed to extend the standard Radon transforms to modeling of curves of arbitrary shapes [24]. This allowed, for instance, the formulation of parabolic and hyperbolic Radon transforms. Such methods entail more complex direct and inverse transformations and are applicable to specific classes of curves.




Fig. 3. Maxima search in an image region (inside orange boxes): (a) correction introduced by \overline{Rf} allows to take into account shorter segments (green against red), (b) circular masks inside regions. (c) Considering the distance and orientation in the neighbourhood (blue regions) to select a segment in the current region.

favours lines that pass near to the centre of the analysed region. This effect is undesirable since we employ the transform to subimages whose location is selected in a fixed manner regardless of the presence of CLS and their positions. To address this issue, we normalize the Radon transform $Rf(\rho, \theta)$ by a factor equal to the length of a segment \mathcal{L} inside its supporting image region at grid-node (i_c, j_c) :

$$\overline{Rf}(\rho, \theta) = Rf(\rho, \theta) / \mathcal{L}(\rho, \theta, i_c, j_c). \quad (2)$$

This normalization provides a fairer comparison between line segments of different lengths. We only consider segments which are at least one third of the region size. This disregards shorter segments which are more vulnerable to noise.

For each node we extract the S -many line segment candidates that correspond to the first S maxima of the normalized transform $\overline{Rf}(\rho, \theta)$ and that are at least $\Delta\rho$ or $\Delta\theta$ apart. This minimal distance criterion between the segment candidates is introduced for the following two reasons. Firstly, CLS are often wider than the one-pixel wide template employed by the transform and are affected by noise/blurring. Accordingly, the criterion treats such CLS as a single segment candidate. Secondly, this arrangement allows the approach to provide a larger variety of segments which can be crucial when recovering poorly contrasted parts of the lines. In order to increase the rotation-invariance we employ circular masks inside subimages as illustrated by Fig. 3(b). Hence, we consider the same number of distinct segments in each direction.

III. LINE STRUCTURE EXTRACTION

The localized Radon transform employed in the first stage extracts the CLS of interest together with many false candidates of various origins, such as overlaps, non-curvilinear geometrical structures, noise, etc. In order to select the relevant segments from a largely redundant set of line candidates we allow not more than one segment to remain per grid node region. Defined on the overlapping grid this choice allows a certain degree of segment overlap that is necessary to model simple line intersections. More complicated intersections can be modelled if an appropriately small scale is chosen.

To refine the segment configuration we consider a CRF model [25] to describe the local dependencies on the grid. A CRF is a form of undirected graphical model that provides a probabilistic framework for labeling dependent data fields. Specifically, a field (X, Y) on observations X

and random variables (labels) Y is called CRF when the random variables $Y|X$ obey the Markov property: $p(Y_v|X, Y_w, w \neq v) = (Y_v|X, Y_w, w \sim v)$, where v and w are vertices in the graph indexing Y , and $w \sim v$ denotes the relationship of neighborhood. The primary advantage of CRFs over hidden Markov models (HMM) is their conditional nature, resulting in the relaxation of the independence assumptions required by HMMs in order to ensure tractable inference. Additionally, CRFs alleviate the label bias problem [26], a characteristic weakness of many conditional Markov models based on directed graphical models, and perform very well in many real labeling problems [25], [26].

In our graphical model each grid node is assigned a label indicating which of the maxima at this location pertains to the extracted CLS configuration. For computational reasons we consider the Potts model of size two. As illustrated in Fig. 2, this corresponds to a 3-by-3 neighbourhood comprising the current grid-location and its eight closest neighbours. To reduce the computational complexity incurred by considering all possible cliques, the relevant clique is selected adaptively based on orientation angle α of the current line segment l . Specifically, we consider four kinds of cliques supporting elongated CLS crossing the central region, cf. Fig. 2:

- horizontally oriented, $\alpha \in [0, 22) \cup [157, 180)$,
- vertically oriented, $\alpha \in [67, 112)$, and
- two diagonal directions, $\alpha \in [22, 67)$ or $\alpha \in [112, 157)$.

Thus, for each segment candidate the considered neighbors are selected differently and depend solely on its orientation angle.

The CRF assumption and Hammersley-Clifford theorem [20], [21] allow the probability of a grid configuration L to be written as a Gibbs distribution, namely

$$E(L) = \frac{1}{Z} \exp\left(-\sum_n E_n\right),$$

where Z is a normalizing constant and the sum in the exponent gives the total energy with the summation taken over all nodes of the grid. The local energy contribution E_n at the n -th grid location (i.e. containing segment l_n) is constructed as follows:

$$E_n = E(l_n|l_n^-, l_n^+) = D(l_n) + \sum_{all\ i} \gamma_i \left[V_i(l_n, l_n^-) + V_i(l_n, l_n^+) \right], \quad (3)$$

where: $D(l_n)$ is the associated unary data term; and V_i are the potential energy terms with weight parameters γ_i controlling their contributions. Note that the neighbouring segments l_n^-, l_n^+ are chosen in accordance with the orientation of the current segment l_n , see Fig. 2.

In the following subsections we introduce potential and unary terms along with the energy minimization procedure. Our flexible framework allows the unitary terms to be designed according to the expected imagery properties such as level of noise, blur, the presence of non-CLS geometry, etc. As described below the potential terms originate from the defining properties of curvilinear structures, namely low curvature, continuity and elongatedness.

A. Potential Data Terms

In most of the image processing applications the initial data suffer a lot from noise, blur, volume artifacts, etc. To recover realistic line structures from such data we have to rely on several basic assumptions. In this work we rely on two central assumptions. First, we assume that the CLS of interest are sufficiently elongated and smooth. This means that at the considered scale such lines should consist of, at least, several segments with similar orientation. Second, we assume that the gaps in the CLS, originating from strong noise or occlusions, are small compared to the sizes of the CLS and that they do not exceed the size of the considered grid region.

To induce realistic CLS configurations we consider interaction terms of two types. Firstly, we want to favour smooth configurations with low degree of curvature, i.e. neighbouring segments should have similar orientations. Secondly, we want to favour continuous line structure, i.e. neighbours should be connected. Hence, we introduce the following two penalties, as below.

- **Orientation penalty:** As discussed above, the consideration of a localized, rather than global, Radon transform allows us to approximate lower order curves by line segments. However, there may exist several line segments with differing orientations, especially in blurry regions. It is therefore sensible to constrain the orientations of adjacent line segments. To this end, we penalize the angle dissimilarity with the following term:

$$V_o(l_n, l_n^+) = \left(\min\left(|\alpha_n - \alpha_n^+|, 180 - |\alpha_n - \alpha_n^+|\right) / 90 \right)^2, \quad (4)$$

where α_n, α_n^+ are orientations (in degrees) of the two line segments l_n and l_n^+ .

- **Distance penalty:** In order to encourage the recovery of missing/poorly-contrasted/occluded line parts and favour continuous CLS, we penalize the segment discontinuities proportional to the distances between the segments that are considered as parts of the same structure:

$$V_d(l_n, l_n^+) = \min_{x \in l_n, y \in l_n^+} dist(x, y). \quad (5)$$

Note that contrary to approaches which allow segment detection at random locations, such as, see [18] and [19], we do not have to penalize overlap of the lines (this is due to the deterministic grid-based generation of line candidates).

These penalties are assigned with weight parameters to control their energy contribution. An example of the effect of these penalties for the case of mammographic CLS extraction is illustrated in Fig. 3(c). In this example appropriate weights of the penalty terms enable the approach to take into account both the distance and orientation in the neighbourhood (blue regions) and to arrive at a suitable choice of segment in the considered region (orange): the third strongest Radon maximum (shown in green) is chosen against the first maximum (red) — which is poorly oriented — and the second maximum (yellow) — which is poorly located (i.e. distant from its neighbouring segments shown in blue).




Fig. 4. Inner and outer (background) regions (at horizontal orientation) employed to calculate the Bhattacharyya distance $d(l)$.

B. Unary Data Terms

The role of the unary data term is to capture the likelihood that a segment exists in the current grid location. This is performed independently for each segment extracted by the localized Radon transform, given the observed data. For flexibility, unary terms can take different forms according to the specific application. In the following we consider several kinds of unary terms and specify where they can be appropriate.

We start with a unary data term that is appropriate for data which contains high levels of noise and little geometric structure. In this scenario, contrast alone cannot be used to distinguish between the background and the ridge-like CLS. Thus, a likelihood based on template matching is appropriate and we first define a quantity

$$d_R(l_c) = \frac{\sum_{\rho \in I_\rho} \sum_{\theta \in I_\theta} \overline{Rf}(\rho, \theta)}{\sum_{\text{all } \rho} \sum_{\text{all } \theta} \overline{Rf}(\rho, \theta)},$$

where (ρ_c, θ_c) are the coordinates of the segment l_c in the Radon transform space, $I_\rho = [\rho_c - \Delta\rho, \rho_c + \Delta\rho]$, $I_\theta = [\theta_c - \Delta\theta, \theta_c + \Delta\theta]$, and the summation in the denominator is taken over *all* possible values of ρ and θ inside the current region (i_c, j_c) . Here, $d_R(l_c)$ estimates the probability of a segment bundle centred in (ρ_c, θ_c) against all possible segments in the analyzed image region. Owing to the normalization term in the denominator, this segment probability assignment is locally contrast-invariant and robust to local histogram stretching. We define the **Radon unary data term** as:

$$D_R(l_c) = 1 - 2d_R(l_c), \quad (6)$$

which takes values in $(-1, 1)$ and is smaller for the strongest lines in the region and greater for weaker lines.

An alternative unary data term is appropriate when the geometry and contrast of the considered lines are sufficiently high, e.g., in the case of road detection in remotely sensed optical images. This unary data term evaluates the dissimilarity between the texture inside the line segment l and an outer (background) region comprising two parallel lines located at distance p on both sides of the current segment position (as in Fig. 4). For example, if the distance p is set to one, then the two parallel lines are adjacent to the current segment. In general, the value of p is chosen such that the outer stripes are sufficiently far from the current segment in order to identify wide lines correctly. The Bhattacharyya distance $d_B(l)$ [27] is used as the dissimilarity metric between pixels in the considered segment and in the outer stripes.

This distance measures the distributional similarity of two continuous random variables. If we assume that the pixel intensities over two separate short line regions originate from two independent Gaussian variables $n_1 \sim N(\mu_{in}, \sigma_{in}^2)$ and $n_2 \sim N(\mu_{out}, \sigma_{out}^2)$, the Bhattacharyya distance is defined as

$$d_B(n_1, n_2) = \frac{1}{4} \frac{(\mu_{in} - \mu_{out})^2}{\sigma_{in}^2 + \sigma_{out}^2} - \frac{1}{2} \ln \frac{2\sigma_{in}\sigma_{out}}{\sigma_{in}^2 + \sigma_{out}^2},$$

where the means μ and variances σ^2 are replaced by their standard sample estimates [27]. Note that the Gaussianity assumption is accurate if we consider two homogeneous regions of an optical image and may be wrong for different acquisition modalities. The choice of metric is motivated by a good track record of performance in various object detection applications [19]. An alternative is, for example, to consider the Student t -test statistic to estimate the dissimilarity in the means [18]; however, this is variance-insensitive. Distance d_B takes values from 0 for the exact same distributions to $+\infty$. We then construct a **Bhattacharyya unary data term** as follows:

$$D_B(l_c) = \begin{cases} 1 - \frac{d_B(X_1, X_2)}{d_0}, & \text{if } d_B(X_1, X_2) < d_0 \\ \exp\left(1 - \frac{d_B(X_1, X_2)}{d_0}\right) - 1, & \text{otherwise,} \end{cases}$$

where X_1 is the set of pixels of the line segment l_c and X_2 - of the two outer lines. Here d_0 is a sensitivity parameter: the higher its value, the more selective the data term is. The distance $D_B(l_c)$ takes values between -1 for perfect radiometric contrast between the line and the background stripes, and 1 for the exact same statistical patterns in the different regions (poor contrast).

C. Energy Minimization

Once the energy (3) with all the involved data terms have been formulated, the refined line structure is recovered by optimizing the CRF configuration, i.e. finding the labels that yield the smallest total energy. Various optimization techniques have been developed in recent decades to address this challenging and time-consuming problem, see [28]. Standard graph-cut techniques based on expansion and swap moves cannot be applied to the energy (3) due to the non-regularity of the considered potential terms (see below). Methods designed for non-submodular energies based on the roof duality and its extensions [29] have empirically been found inappropriate to the considered optimization problem since they fail to label a large portion of nodes. Loopy belief propagation [30] experimentally failed to converge to a stable label configuration. Therefore, we apply a MCMC optimization with simulated annealing [20] to optimize the CRF configuration.

MCMC is a general method for obtaining random samples from a probability distribution for which direct sampling is difficult or unfeasible. For our problem we employ a point processes that uses an MCMC sampler to search for the configuration which minimizes the energy E . The sampler simulates a discrete Markov chain that converges towards an equilibrium state around the minimum energy configuration [20]. In a Metropolis algorithm at each iteration the current label configuration is locally perturbed according to

a predefined proposal density. The resulting configuration is then accepted as a new state of the chain with a certain acceptance probability depending on the energy difference of the two configurations and a relaxation parameter T called temperature. The values of T depend on time t and approach zero as $t \rightarrow \infty$. The described MCMC procedure with a predefined procedure of decreasing T constitute MCMC with simulated annealing.

The employed MCMC procedure is initialized with a configuration that, at each node, gives the maximum value of the unary data term. It then proceeds iteratively by selecting at random (uniformly over all the grid nodes) a grid-location and proposing at random (uniformly over all the segment candidates) a new segment l_n from the list of line segments associated with this grid-location. A new configuration l_n is accepted and replaces the current l_c with the acceptance probability $\delta = \min(1, \exp((E_c - E_n)/T))$. The resulting chain of configurations corresponds to the Metropolis-Hastings procedure [21] with a uniform proposal distribution. One iteration of the approach is completed once all of the nodes have been visited at least once by the iterative process. The temperature parameter T is decreased at each iteration. This encourages more exploratory behaviour during the early stages; i.e., it accepts states with higher energy with non-zero probability for the sake of better exploration. As it decreases, the temperature becomes more prohibitive until, as it approaches zero, the algorithm only accepts states with less energy. The iterative process is stopped when the process stabilizes, i.e. when the proportion of relabelings performed during an iteration becomes small.

IV. CLS EXTRACTION ALGORITHM

The outline of the proposed detector is presented in Fig. 6. The first part of the algorithm (lines 2-5) is directly parallelizable, since the operations are performed independently in each node of the grid. The second part (lines 7-17) cannot be parallelized directly, since the update of sites in MCMC has to be performed sequentially to ensure convergence [20]. This second part of the algorithm is the most time consuming, and in order to accelerate it without violating the convergence conditions we resort to image partitioning. At each MCMC iteration we split the grid into NG similarly-sized grid parts. To be able to run MCMC separately and in parallel we have to ensure the proper independence of nodes in each of the grid parts with nodes from all other parts. To this end we “freeze” the MCMC state in nodes whose segment selection is affected by nodes from more than one grid parts; due to the use of 3-by-3 neighbourhood it suffices to leave out stripes of single-node width, as demonstrated in a simple case of $NG = 4$ on a square grid in Fig. 5. More specifically, at the beginning of each iteration we first randomly select a new seed for partitioning (red in Fig. 5), and then identify the NG sets of nodes, where MCMC can be parallelized (blue, green, yellow, pink in Fig. 5). The selection of blocked nodes (black in Fig. 5) is done in a regular manner to arrive at NG sets of nodes of similar size. Since, typically, multiple MCMC iterations are performed and the seeds of partitioning




Fig. 5. Partitioning of a square grid into $NG = 4$ parts (blue, green, yellow, pink) of similar sizes to parallelize MCMC.

```

1 define an overlapping rectangular grid of square regions on the image;
  // line segment detection
  for all grid nodes  $(i_c, j_c)$  do
2   perform the discrete Radon transform  $\overline{Rf}$ ;
3   find the first  $S$ -many maxima of the transform;
4   calculate the unary data terms for the  $S$  segments;
5   initialize MCMC by setting  $l_c$  as the segment with maximal value
     of unary term;
  // line structure extraction
6 set temperature  $T := T_0$ , acceptance ratio  $\Delta := 1$ ;
  while  $\Delta > M_{stop}$  do
7   reset counters  $changed := 0$ ,  $total := 0$ ;
     while not all grid nodes have been visited do
8      $total := total + 1$ ;
9     randomly select a node  $(i_c, j_c)$  on the grid;
10    randomly select a new candidate  $l_n$  for the node  $(i_c, j_c)$ ;
11    calculate energies  $E_c$  and  $E_n$  in the relevant cliques based on
        $l_c$  and  $l_n$ , respectively;
12    generate a uniform  $u \sim U[0, 1]$ ;
       if  $u < \min(1, \exp((E_c - E_n)/T))$  then
13      accept the candidate  $l_c = l_n$ ;
14       $changed := changed + 1$ ;
15    identify acceptance ratio  $\Delta := changed/total$ ;
16    apply geometric temperature decrease  $T := T * \tau$ ;
  // post-MCMC thresholding
  for all grid nodes  $(i_c, j_c)$  do
17   remove weak segments with  $E_c < M_{thresh}$ .
```

Fig. 6. Pseudocode for the proposed elongated CLS extractor.

are selected randomly on each iteration, the effect on convergence of the preferential treatment of the blocked nodes is negligible.

Theoretical considerations require the cooling schedule in MCMC to be logarithmic [20], but as in many applications [18], [21], we employ the geometric descent to accelerate the convergence at the possible cost of worse global exploration. The iterative process is stopped when the configuration stabilizes, i.e. when the proportion of accepted line candidates within a given MCMC iteration goes below a threshold M_{stop} .

Note that in the developed algorithm the segment candidates compete solely with those located at the same node of the grid. Accordingly, after the MCMC procedure each node of the grid contains a line segment. Sporadic, disjoint segments can survive when all the candidates at the given grid node interact weakly with their neighbours due to acquisition noise, low

contrast or absence of lines at the considered scale. To remove these we introduce a post-MCMC thresholding (line 17) to remove weakly interacting segments.

V. EXPERIMENTAL EVALUATION

In this section we experimentally validate the developed CLS extraction approach to confirm its performance and flexibility. To this end we present three case-studies. The first application is in mammographic image analysis, where the extraction of CLS is required for image coregistration and/or spicule detection — both of which facilitate the detection of malignancy. The second application is the extraction of road networks from remotely sensed optical images— this plays an important role in urban planning, traffic control, maps updating, etc. Finally, we present the results of the developed technique in palm line extraction which is one of the critical problems in palm image alignment and the underlying biometric feature extraction.

The three applications, considered here, originate from different domains— medical, remote sensing, and biometrical image analysis. Each case presents different levels and types of noise. Furthermore, the CLS in each case possess different geometrical and statistical properties. Together, these examples thus represent a suitably wide variety of challenges for the proposed CLS extraction approach. In each of the case-studies we begin with a brief application-specific state-of-the-art overview, and then present the experimental results and comparisons.

Throughout the experiments the construction of the energy term is adapted according to the selection of appropriate unary data terms. Therefore, the penalty weights (γ_o, γ_d) will be different in order to achieve a discriminative energy representation. The choice of these parameters poses, in fact, a challenging problem. In case of complete data (ground truth with labels), these parameters can be estimated by stochastic gradient [31]. This algorithm converges locally, i.e. strictly in the neighborhood of the initial set of parameters. In incomplete data scenario, the weight parameters have to be estimated simultaneously with the label configuration that minimizes the total energy E . The standard approach to such problems is the Expectation-Maximization (EM) algorithm. Nevertheless, due to various complexity and feasibility issues the EM algorithm and its variants are not well adapted and there is no guaranty of convergence toward the maximum likelihood estimation. A partial estimation may consist of formulating inequality-based constraints on the γ -parameters to arrive at improved estimates within the same iterative MCMC process, see [18]. In this paper, however, we do not consider the weight parameter estimation problem and take their values as known *a priori*, like in [19] and [32].

Also, the remaining data sensitive parameters— grid scale and post-processing threshold will vary throughout this section. Otherwise, the same parameter configuration, which is robust to resolution and does not require adjustments when changing applications, is employed, namely: square grid with 1/3 overlap of the node regions, as in Fig. 2; $[0, 180)$ angle range for the Radon transform with a step of one degree; distance between Radon maxima of at least $\Delta\rho = 5$ or

$\Delta\theta = 5$; $S = 5$ candidate segments per region; MCMC procedure is set with $T_0 = 1$, $\tau = 0.975$ and the stopping threshold of $M_{\text{stop}} = 0.01$. Finally, we report the computation times achieved in a MATLAB implementation with CPU-parallelized Radon transform/maxima calculation and the described above MCMC parallelization with $NG = 4$ after, typically, 30–40 MCMC iterations on a Core-i7 2.7GHz, 8Gb RAM, Windows 8 system.

A. CLS Extraction From Mammographic Images

Mammography is the most widely spread tool for breast cancer monitoring. The purpose of automatic mammographic image processing is to facilitate the work of medical personnel in detection, classification, and measuring of anatomical features such as masses and CLS. The extraction of the latter poses a challenging image processing problem due to their low contrast, variable widths and partial occlusions caused by volume projections occurring during 2D image acquisition of a 3D object. Accurate CLS extraction is of interest for various mammographic image processing problems including registration [33], mass detection by linear structure removal [34] and microcalcification detection [34]. Extraction of CLS is of particular significance in spiculated mass detection [4], [35], since, as has been noted in [35], almost half of the malignant masses in mammograms are surrounded by a radial pattern of spicules. Various methods for CLS detection have been proposed: based on lines/binning operators [35], template matching [35], global Radon transform [4], Dual-Tree Complex Wavelets [36], contourlets [37], open-ended active contours (snakes) [38], etc. More detailed overviews of approaches and their limitations can be found in [35], [38], and [39].

In this study we focus on CLS extraction; its subsequent classification is outside of the scope of this work and can be achieved by various methods, such as, e.g., principal component analysis [35] or random forests [36]. Mammographic images are not geometrically regular and contain a lot of blurry and non-curvilinear structures [35], [36]. Therefore, we employ the Radon unary data term (6).

Experiments were performed on images from the Digital Database for Screening Mammography [40]. Throughout this subsection (unless specified otherwise) we employ an overlapping grid of square regions with widths equal to 50 pixels, penalty weights $(\gamma_o, \gamma_d) = (2.0, 0.25)$, and the postprocessing threshold of $M_{\text{thresh}} = 1.3$. We analyzed over 50 cranio-caudal and mediolateral-oblique view mammograms obtained at 50 micron resolution, containing various CLS, several results were earlier reported in [41].

1) *Synthetic Images*: We first present a set of synthetic experiments where the goal is to evaluate the accuracy of CLS extraction in the presence of a typical mammographic background. To this end we consider a normalized mammographic image (MI) of healthy fatty tissue, see Fig. 7(a), and a monochrome synthetic image (SI) comprised of elongated CLS with different length and curvature, see Fig. 7(b). We perform CLS extraction on an overlapping grid with square regions of $s_3 = 47$ pixels on the two images separately and then on a weighted mixture $0.9 \times \text{MI} + 0.1 \times \text{SI}$.

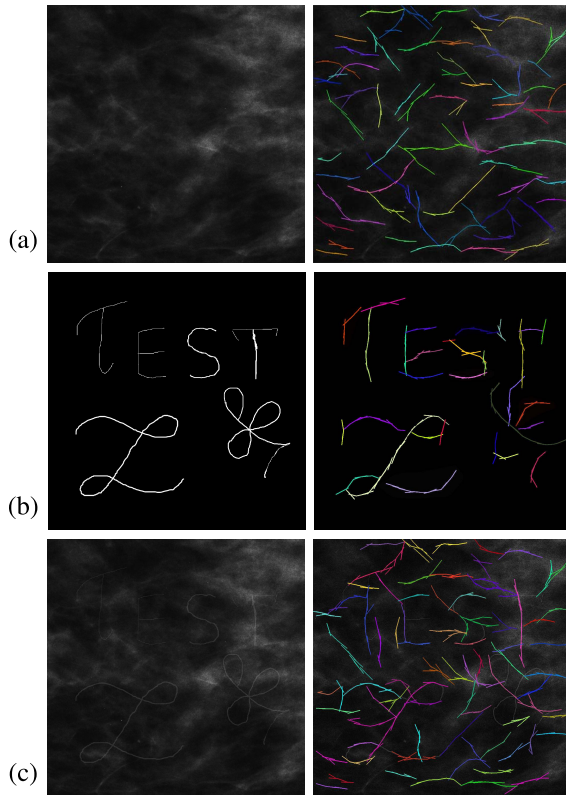


Fig. 7. CLS Extraction results on overlapping grid with square regions of $s_3 = 47$ pixels size on (a) 1000×1000 mammographic image (MI) with healthy tissue, (b) 1000×1000 synthetic image (SI) with CLS, and (c) weighted mixture $0.9 \times \text{MI} + 0.1 \times \text{SI}$.

We observe that applied to SI the CLS extraction is quite accurate though slightly angular due to the underlying line segment approximation. In accordance with the selected scale some CLS parts, like those in the leaf (right bottom) and letter ‘s’, are not detected due to high degrees of curvature. When applied to the mixture, we expect most of the CLS comprising SI to be extracted along with CLS present on MI. Note that not *all* the CLS from SI remain CLS in the weighted mixture since some CLS can vanish in regions where MI has bright patches and/or CLS of its own. Nevertheless, to numerically evaluate the (pixel-level) extraction accuracy we take SI as the ground truth for detection.

To arrive at the numerical accuracies of the CLS extraction on Fig. 7(b), (c), we dilate the detected CLS using a disk of radius three, and compare them to the ground truth. To evaluate the accuracy of the results we plot a receiver operating characteristic (ROC), that reports the fraction of true positives out of the total actual positives, called true positive rate (TPR), vs. the fraction of false positives out of the total actual negatives, called false positive rate (FPR), at various parametric settings. To reduce the impact of mammographic lines on the reported FPR we take into account solely the extracted lines located within 10-pixel radius from any white pixels in SI. In Fig. 7(b) the obtained TPR is equal to 0.765 and $\text{FPR} = 0.022$; and on the mixture, Fig. 7(c), $\text{TPR} = 0.741$ and $\text{FPR} = 0.024$. The results of a more representative study with four different weighted mixtures are summarized in Fig. 8. For each of the




Fig. 8. ROC curves of CLS extraction of a synthetic image (SI) from weighted mixtures of SI with a mammographic image (MI), see Fig. 7. Each curve presents six results at different grid scales with square regions with a side of (from left to right on each curve) $s_1 = 65$, $s_2 = 56$, $s_3 = 47$, $s_4 = 38$, $s_5 = 29$, and $s_6 = 20$ pixels.

mixture-images we have evaluated the extraction results at six different scales with regions of sizes in range from 20-by-20 to 65-by-65 pixels. Each accuracy measurement (TPR, FPR) has been averaged over five reruns of the developed CLS extractor; this is necessary due to the stochastic nature of the developed approach. Fig. 8 shows that the extraction accuracy increases with the weight of SI in the mixture— this growth is very fast for SI weight values between 0.05 and 0.1, and is almost saturated with SI proportion above 0.3. Fig. 8 also shows that the finer the scale (smaller region size) the larger the TPR as well as FPR. The results of method [4] are not presented in Fig. 8 due to a different parametrization. Furthermore, the results are quite different from those obtained by our method: for mixtures the average TPR is slightly higher - in range from 0.88 to 0.94 at the expense of significantly higher FPR which is above 0.08. The average observed computational time for the proposed method is in range from (on average) $t_1 = 42$ seconds (scale s_1) to $t_6 = 186$ seconds (scale s_6) on the 1000×1000 images used here.

2) *Real Images*: In Fig. 9 we present four typical results on 1000×1000 pixels regions of interest: (a)-(c) with spiculation/architectural distortions and (d) of healthy tissue with CLS. In the second column we report maps with all the detected line segments after the line segment detection step. Here the intensities are proportional to the value of their Radon energy term. We observe that all CLS that are visually present in the image are contained in the maps along with a number of false positives. This observation empirically validates the applicability of the local Radon transform for segment detection in this type of highly blurred image. In the third column we present the CLS extraction results, where the intensities correspond to the energies (3) scored by the line segments. We observe that the majority of the false positives have been discarded; however some possibly false positives remain. These results can be further processed by removing the shorter line-segment chains if such structures are deemed irrelevant. This is demonstrated in the fourth column where each CLS is assigned a unique colour. To arrive at this result we have employed a deterministic clustering by

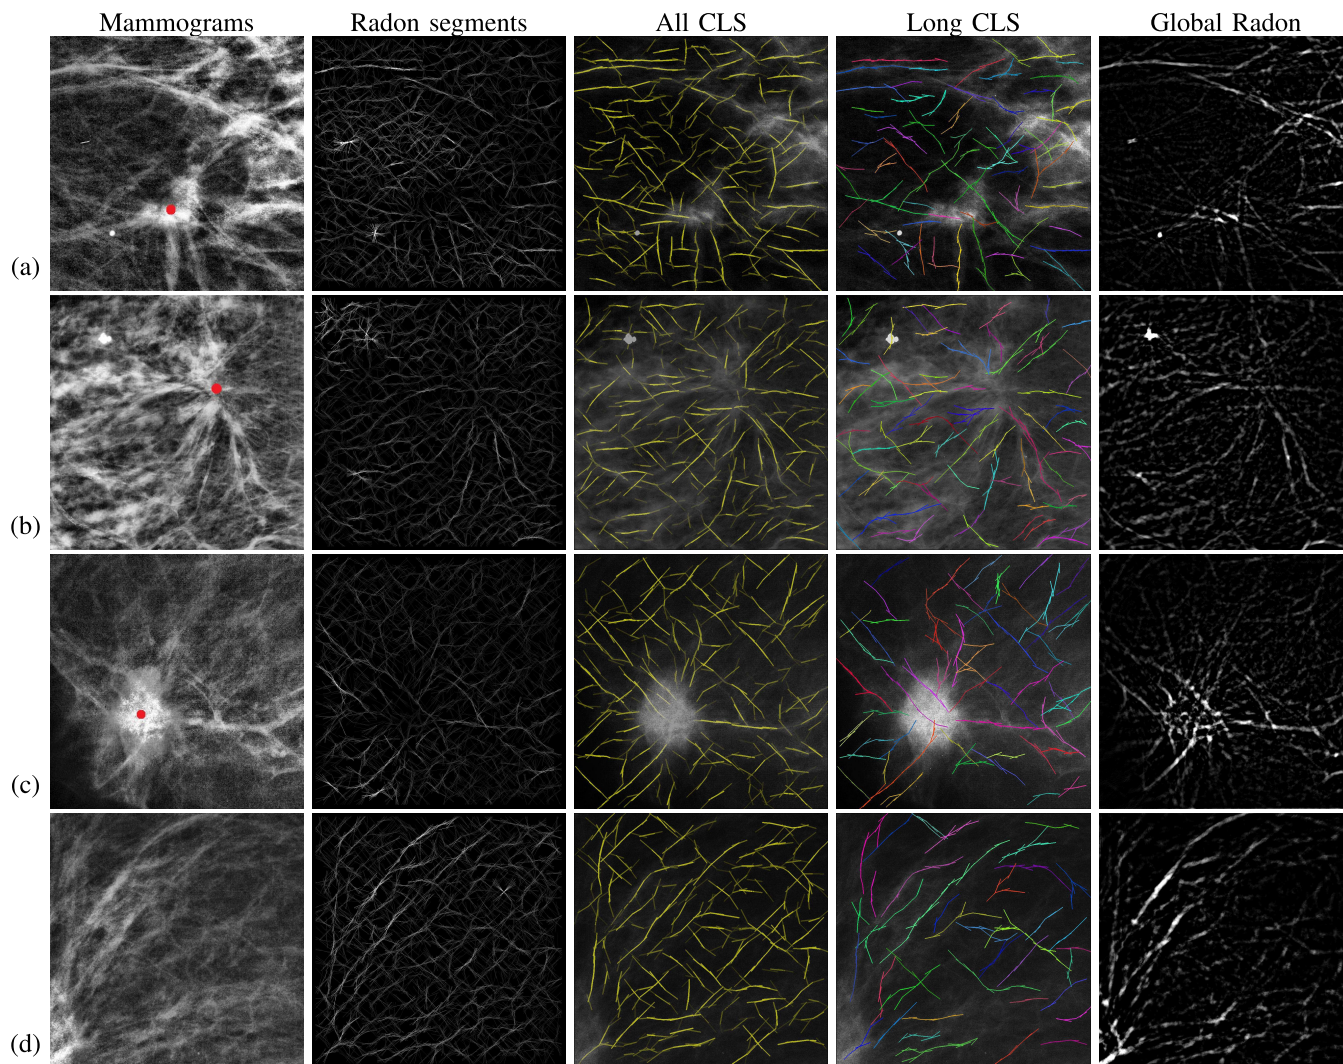


Fig. 9. CLS extraction on 1000×1000 pixels mammographic images (a)-(c) containing spiculated masses or architectural distortions (with red dots marking their centres) and (d) healthy tissue with CLS, from left to right: images after local histogram enhancement (for visualization only), all extracted Radon segments, all extracted CLS, long CLS with colour coding, result of the global Radon transform enhancement [4].

combining segments at adjacent grid positions with penalties $V_o(l_1, l_2) \leq 0.1$ and $V_d(l_1, l_2) \leq 5$, and retaining solely the CLS constructed of five or more segments. We observe that this manipulation has significantly reduced the noise but might be undesirable in some cases since it removes short true positives.

Visually, the star-like CLS clusters forming around the mass centres (depicted schematically by red dots in Fig. 9(a)-(c)) are quite evident. At the same time in the healthy tissue image, Fig. 9(d), the CLS do not extend from any single point and are almost parallel, slowly converging to the nipple, which is characteristic of healthy linear features on mammograms. For comparison, in the fifth column we present the results of the global Radon filtering method [4]. Note that due to the similar methodological component employed (Radon transform) the method in [4] is the most relevant comparison. Furthermore, methods in [35]–[37] do not recover elongated connective CLS, and active contours in [38] require manual initialization in the form of mass/CLS position specification. It is immediate that the proposed method performs much better in extracting

continuous and low-contrast CLS and in that sense suffers significantly less from the detection noise. Both the proposed method and that in [4] have a single ‘physical’ parameter: in our method - the expected minimal length (i.e. grid scale), and for [4] - the expected width of the CLS. From an operational point of view the minimal length parameter is less restrictive and detection of structures at different distances from the sensor. Specifically, for our method, the same grid specification was used for all images, whereas we had to tune the width parameter of the global Radon method in [4] from $w = 12$ to $w = 25$ because the detection results vary appreciably.

B. Road Network Extraction

The problem of road network extraction from aerial or satellite imagery plays an important role in image processing for various applications, such as image coregistration, building detection, urban planning, and agricultural and forestry mapping. All of these applications can benefit greatly from the development of unsupervised, reliable, and computationally fast road network extraction methods. The development of

such techniques is hindered by several problems, namely: the heterogeneous nature of road materials results in different radiological patterns of roads; the various types of occlusions (shadows, buildings, tree canopies); and the varying road width present in the same scene. As such, standard line, edge and ridge detection techniques [23] are inappropriate for road detection, see in [42] and [43].

A representative range of approaches have been designed to extract road networks: from region growing [1] and dynamic programming [42] to endpoint tracking [44], and Hough transform-based detection with Gabor-filtering [16]. Junction-based road-network extraction represents another promising research direction undertaken recently to extract roads from optical urban (i.e. highly structured) scenes [45] and SAR imagery [2]. A completely different stochastic approach considers an image as a stochastically generated configuration of a certain kind of geometrical primitives, such as lines, circles, ellipses, etc. This approach has given rise to a range of stochastic techniques, such as active contours [32] and marked point [18], [19], [46] processes. Some of these techniques report very accurate results but they suffer from the need for operator input [1] or they require very computationally expensive stochastic optimization [18], [32]. For example, processing a typical 1000×1000 pixels scene by stochastic techniques can take from several minutes [46] to over an hour [18].

The typical characteristic of the road in urban imagery is its high contrast, thinness, connectivity and elongatedness as compared to the other objects present in the scenes, such as trees, buildings, smaller objects, etc. Thus we employ the Bhattacharyya unary data term which highlights highly contrast line segments. We present experiments on road network extraction in semi-urbanized zones on several images from Google Maps (© Google) of size 400×500 and with approximately 0.5 metre-per-pixel ground resolution. In this subsection we employ an overlapping square grid with side equal to 30 pixels, energy weights $(\gamma_o, \gamma_d) = (0.75, 0.5)$, unary data term sensitivity parameter $d_0 = 0.3$, and a post-processing threshold $M_{\text{thresh}} = -1$. We have performed experiments on over 20 high resolution images with various road networks and three typical extraction results are presented in Fig. 10. We observe that the segment extraction via Radon transform with Bhattacharyya unary data term $D_B(l)$ identify the road CLS well with only a few undetected segments. Note that some of these undetected roads originate from occlusions or low contrast and might be identified by varying the grid scale. A comparison with the general CLS extraction technique [14] confirms high appropriateness of the developed method to road extraction. Specifically, we observe that our method tolerates small gaps and higher degrees of curvature than the benchmark [14]. The average computation time is just below 5 seconds of which around 1.5 seconds is needed for the line segment extraction step, lines 1-5 in Fig. 6, and 3 – 3.5 seconds for the MCMC iterations.

We perform a further comparison with five state-of-the-art line-detection techniques: multi-scale CLS profile extractor [14], active contours [32], two reversible jump MCMC-based techniques [18], [19] and a jump-diffusion




Fig. 10. Road network extraction results on Google Maps images (©Google). From top to bottom: initial images, ground truth (manual extraction) of the road network, detected Radon segments maps weighted with $(1 - D_B(l))/2$, extraction results weighted with $\exp(-E_n)$, and multi-scale CLS profile extraction [14] result.

approach [46] on a 650×900 ‘Road’ image (©IGN) [19], courtesy of F. Lafarge, available at <http://www-sop.inria.fr/members/Florent.Lafarge/benchmark/>, see Fig. 11. The latter three techniques employ stochastic configurations, and their characteristic difference with the approach developed in this paper is the random number of objects in the analysed scene. Whereas after the MCMC process convergence this number is equal to the number of grid nodes in the developed approach. This may be reduced later on if none of the candidate segments at any of the grid locations is assigned a sufficiently strong energy (cf. thresholding in line 17 in Fig. 6). On this image we employ the same parametric setting except for a smaller size of square grid (15 pixels) and an increased weight of the orientation term $\gamma_o = 1.2$ due to high connectivity and curvature of the processed road network. The lines in the outer regions are taken at distance $p = 3$ pixels, which corresponds to the average road width of 4-6 pixels.

A numerical comparison is presented in Table I, where the results of the benchmark techniques are reported as in [19].




Fig. 11. (a) ‘Road’ image (©IGN) from [19], (b) the ground truth map and (c)-(h) detection results.

TABLE I
QUANTITATIVE RESULTS OBTAINED ON THE ‘ROAD’ IMAGE

Algorithm	TPR	FPR	Time
Proposed method	0.709	0.053	13 sec
Steger [14]	0.611	0.048	2 sec
Verdie et al. [19]	0.637	0.004	15 sec
Lafarge et al. [46]	0.658	0.013	381 sec
Lacoste et al. [18]	0.812	0.006	155 min
Rochery et al. [32]	0.490	0.013	60 min

Note that the developed approach does not estimate the widths of line segments. Our method detects some lines inside the road which are aligned accurately (due to the orientation penalty) and that are sufficiently long, i.e., at least several times longer than the side of the employed localized Radon transform window. The designed method is not appropriate for detecting the side lines of the roads (edges). To present an adequate comparison with methods [18], [19], [32], [46] that estimate the road widths the detection results reported by our method and Steger’s approach [14] were dilated with radius $r = 4$ at the post-processing stage. Table I shows that the proposed method achieves a higher TPR at the expense of an increased FPR. This difference is partially due to the above mentioned issue of CLS widths: template matching with various widths is more efficient and accurate at the expense of a sharp increase of computation complexity. Indeed, the most relevant improvement can be seen in computation time. Note that the computation times of the benchmark techniques are taken from [19], although they were obtained on a different hardware system and can therefore only serve as a rather rough comparison. It is immediate that the proposed approach performs significantly faster than methods [18], [46]. The technique in [19] gives a comparable computation time; however, its results were obtained in a massively parallelized CUDA-implementation with a specialized GPU card. Whilst a similar kind of implementation is possible for the developed approach, it is beyond the scope of this work. It is also

worth noting that the method in [19] employed a preliminary partitioning of the image that specified the probability of finding objects in its regions. This helped to further reduce the computational load. The method proposed in this paper has been employed without any preliminary classification and, thus, we believe that the increase of FPR is partially due to this difference.

C. Palmprint CLS Extraction

In this section we consider the third application for the developed CLS extractor — line detection on palmprint images. This task plays the key role in biometric person verification and identification based on palmprints [47]. Line information (principal lines and, more generally, CLS) is generally regarded as a good discriminative feature for palmprint-based identification [47], [48]. Indeed, robust palmprint verification systems based solely on CLS have been developed. As such, the extraction of palm CLS and principal lines is a problem area which continues to enjoy much activity and interest. Zhang et al. [49] built a principal line detector based on local filtering with thick-line operator filters. Duta et al. [50] used a series of morphological operations to detect principal lines from feature points. Zhang et al. [51] employed overcomplete wavelet representation together with directional context modeling approach to recover principal line features, Wu et al. [48] applied directional filters with smoothing to palm line feature extraction. Huang et al. [47] designed a method based on modified finite Radon transform filtering, followed by direction and energy-based post-processing, in order to identify principal lines.

Since the focus of this paper is CLS extraction, the further problem of palm print matching or verification will not be addressed here. In fact, once the CLS extraction is performed, the procedures employed in, say, [47] can be adapted to perform principal line selection and matching. Only minor post-processing is required for the adaptation: since the designed framework extracts lines (object level) rather than line pixels, we suggest to detect the latter by rank order thresholding (similar to that presented below) in the proximity of the detected line segments.

The CLS contained in palmprint images correspond to the intensity valleys rather than ridges as was the case in the previous applications. Therefore, to avoid modifying the framework we invert the palmprint images as a preprocessing step. The palmprint images do not suffer from any particularly strong source of noise. Nevertheless, as can be seen in Fig. 12 (left and middle columns), the use of both the Bhattacharyya and the Radon unary term results in excessive over-detection inclusive of weak and short secondary lines (wrinkles). To improve on a better performing Radon term result, we constrain the segment selection by considering the perpendicular profiles of each line segment. In particular, we take the average values of a set of parallel segments (at distances $1, \dots, L$ on both sides of every detected segment, where L is the segment length), find the rank of the current segment in the set, and exclude it from the set of candidates if it lies below a predefined threshold (0.8 in this section). Indeed this is close in spirit to the




Fig. 12. Identified long CLS (top) from a palmprint image and the extracted candidate line segments (bottom) obtained with Bhattacharyya unary data term (left), Radon unary term (middle), and Radon term after rank thresholding (right).

Radon maxima selection employed in the first step. However, the analysis of ranks takes into account the relative strength of the segments across grid region borders and discards weak, local maxima. Moreover, the rank-based thresholding reduces the number of segments in the input of the second extraction step by at least a half. This reduction provides a significant acceleration in the optimization step. Some results obtained using rank thresholding are presented in Fig. 12 (right column) and empirically confirm the improvement in CLS extraction performance. In the following, we employ the Radon unary term after rank thresholding.

For the experimental validation we employ the high resolution images from the COEP palmprint database [52]. We run elongated CLS extraction on centrally positioned regions of interest with the following parameters: overlapping square grid with 30-by-30 pixels regions, penalty weights $(\gamma_o, \gamma_d) = (1.0, 0.2)$, and the postprocessing threshold of $M_{\text{thresh}} = 1$. Three typical results are presented in Fig. 13. In the second row, the extraction results are overlaid on the inverted initial palmprint images with segment intensities proportional to the final energy $\exp(-E_n)$. The third row shows the long CLS extracted by the clustering introduced in Section V-A. We observe that the principal lines are accurately extracted with a noticeable presence of weaker CLS detection (wrinkles) on the more challenging second image. Note that a complete removal of wrinkles in complicated cases cannot be obtained without application-specific postprocessing, like the directional thresholding proposed in [47]. In the fourth row of Fig. 13 we present a comparison with the multi-scale line profile extraction method [14] and in the fifth the contourlet hard shrinkage result [7]. It can be seen that the proposed method yields the most satisfactory results. Compared to the pixel-level extractor [47], the designed framework requires more computation due to a different detection goal, namely object extraction. The results in Fig. 13 are obtained in approximately 2 sec (first two images) and 2.5 sec (third image), whereas in [47] the reported time is within 1 sec. Methods [14] and [7] required around 0.3 sec and 80 sec respectively to obtain the presented results. Note that, compared to results in




Fig. 13. Palmprint CLS extraction results on centrally positioned regions of interest: in first two columns 500 × 500 pixels, and in the third— 750 × 750. From top to bottom: initial images (inverted, histogram stretched), extraction results weighted according to the energy $\exp(-E_n)$, identified long lines with colour coding, result of multi-scale CLS profile extraction [14] and contourlet shrinkage [7] result.

Section V-A, the appreciably lower computation time reported by the developed method comes from the rank-based optimization complexity reduction.

We finally validate the robustness of the developed framework to rotations and scale variations. Note that both of these can have big affects on CLS extraction but neither are usually known a priori. Specifically, rotation invariance is necessary due to a common lack of information on optimal image orientation/aliment, and a certain degree of scale tolerance within a reasonable range of values is crucial from the implementation point of view. In Fig. 14 we present extraction results obtained at the 30-by-30 pixels region scale at four orientations: (a) $\alpha = 30$, (b) $\alpha = 45$, (c) $\alpha = 75$, and (e) $\alpha = 0$ degrees. We observe that the detection results are very similar; numerical comparisons confirm over 90% of agreement between the results. Note that due to the stochastic nature of the approach some variations within 2 – 5% (depending on data complexity, scale, etc.) can be expected when applied




Fig. 14. Palmprint CLS extraction after rotation by α degrees (first row), and at different grid-region sizes (second row). (a) $\alpha = 30$, 30×30 . (b) $\alpha = 45$, 30×30 . (c) $\alpha = 75$, 30×30 . (d) $\alpha = 0$, 20×20 . (e) $\alpha = 0$, 30×30 . (f) $\alpha = 0$, 40×40 .

to the same data with the same parameter configuration. The results at three different scales in Fig. 14(d)-(f) show that less robustness is apparent with respect to the scale parameter. Three principal lines are extracted at all three considered scales, with slight variations in details; weaker and shorter CLS (wrinkles) are more affected. Such variations are to be expected due to the key role of the segment length parameter in the extraction procedure. These experiments confirm that the designed extraction framework demonstrates a high level of rotation invariance and a reasonable tolerance to scale.

VI. CONCLUSIONS

We have here proposed an automatic approach to elongated CLS extraction from images affected by noise, blur, and low contrast. The flexible framework, designed here, efficiently combines the fast deterministic Radon-based line segment detection with the stochastic Markov chain Monte Carlo optimization procedure. Furthermore, the conditional Random Field model successfully captures and incorporates local interactions. Since optimization is performed over a partitioned grid this enables a parallelized architecture which further improves computational performance. The developed method demonstrates good rotation invariance and scale tolerance for elongated and curved CLS extraction in the presence of noise.

Experiments demonstrate the efficiency of the approach in three different CLS extraction applications, namely: mammographic image analysis, road network extraction, and palm line detection. We observed that, with an appropriate choice of unary data term, the proposed automatic technique performs successfully on images with various geometrical properties and types of noise. Experimental validation on blurred, curved, and low contrast CLS confirm the robustness and flexibility of the proposed framework.

As a future extension of the work we envisage the development of a multi-scale approach that automatically selects the optimal grid-scale at each location using criteria such as the one discussed in, for example, in [53].

A MATLAB implementation of the developed approach is available at the authors' webpages <https://sites.google.com/site/vlkryl/> and <http://www.homepages.ucl.ac.uk/~ucakjdb/>.

REFERENCES

- [1] M. Amo, F. Martínez, and M. Torre, "Road extraction from aerial images using a region competition algorithm," *IEEE Trans. Image Process.*, vol. 15, no. 5, pp. 1192–1201, May 2006.
- [2] M. Negri, P. Gamba, G. Lisini, and F. Tupin, "Junction-aware extraction and regularization of urban road networks in high-resolution SAR images," *IEEE Trans. Geosci. Remote Sens.*, vol. 44, no. 10, pp. 2962–2971, Oct. 2006.
- [3] D. Rivest-Henault and M. Cheriet, "3-D curvilinear structure detection filter via structure-ball analysis," *IEEE Trans. Image Process.*, vol. 22, no. 7, pp. 2849–2863, Jul. 2013.
- [4] M. P. Sampat, G. J. Whitman, A. C. Bovik, and M. K. Markey, "Comparison of algorithms to enhance spicules of spiculated masses on mammography," *J. Digit. Imag.*, vol. 21, no. 1, pp. 9–17, 2008.
- [5] M. Jacob and M. Unser, "Design of steerable filters for feature detection using Canny-like criteria," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 8, pp. 1007–1019, Aug. 2004.
- [6] M. N. Do and M. Vetterli, "The contourlet transform: An efficient directional multiresolution image representation," *IEEE Trans. Image Process.*, vol. 14, no. 12, pp. 2091–2106, Dec. 2005.
- [7] A. L. da Cunha, J. Zhou, and M. N. Do, "The nonsubsampling contourlet transform: Theory, design, and applications," *IEEE Trans. Image Process.*, vol. 15, no. 10, pp. 3089–3101, Oct. 2006.
- [8] V. Velisavljevic, B. Beferull-Lozano, M. Vetterli, and P.-L. Dragotti, "Directionlets: Anisotropic multidirectional representation with separable filtering," *IEEE Trans. Image Process.*, vol. 15, no. 7, pp. 1916–1933, Jul. 2006.
- [9] D. L. Donoho and X. Huo, "Beamlets and multiscale image analysis," in *Multiscale and Multiresolution Methods* (Lecture Notes in Computational Science and Engineering), vol. 20, T. J. Barth, T. Chan, and R. Haimes, Eds. Berlin, Germany: Springer-Verlag, 2002, pp. 149–196.
- [10] S. Berlemont and J.-C. Olivo-Marin, "Combining local filtering and multiscale analysis for edge, ridge, and curvilinear objects detection," *IEEE Trans. Image Process.*, vol. 19, no. 1, pp. 74–84, Jan. 2010.
- [11] M. N. Do and M. Vetterli, "The finite ridgelet transform for image representation," *IEEE Trans. Image Process.*, vol. 12, no. 1, pp. 16–28, Jan. 2003.
- [12] J. Ma and G. Plonka, "The curvelet transform," *IEEE Signal Process. Mag.*, vol. 27, no. 2, pp. 118–133, Mar. 2010.
- [13] L. Liu, D. Zhang, and J. You, "Detecting wide lines using isotropic nonlinear filtering," *IEEE Trans. Image Process.*, vol. 16, no. 6, pp. 1584–1595, Jun. 2007.
- [14] C. Steger, "An unbiased detector of curvilinear structures," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 20, no. 2, pp. 113–125, Feb. 1998.
- [15] J.-H. Jang and K.-S. Hong, "Detection of curvilinear structures and reconstruction of their regions in gray-scale images," *Pattern Recognit.*, vol. 35, no. 4, pp. 807–824, 2002.
- [16] R. Gao and W. F. Bischof, "Detection of linear structures in remote-sensed images," in *Image Analysis and Recognition* (Lecture Notes in Computer Science), vol. 5627, M. Kamel and A. Campilho, Eds. Berlin, Germany: Springer-Verlag, 2009, pp. 896–905.
- [17] A. C. Copeland, G. Ravichandran, and M. M. Trivedi, "Localized Radon transform-based detection of ship wakes in SAR images," *IEEE Trans. Geosci. Remote Sens.*, vol. 33, no. 1, pp. 35–45, Jan. 1995.
- [18] C. Lacoste, X. Descombes, and J. Zerubia, "Point processes for unsupervised line network extraction in remote sensing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 10, pp. 1568–1579, Oct. 2005.
- [19] Y. Verdié and F. Lafarge, "Efficient Monte Carlo sampler for detecting parametric objects in large scenes," in *Proc. 12th Eur. Conf. Comput. Vis.*, 2012, pp. 539–552.
- [20] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 6, no. 6, pp. 721–741, Nov. 1984.
- [21] W. K. Hastings, "Monte Carlo sampling method using Markov chains and their applications," *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970.
- [22] L. Ding and A. Goshtasby, "On the Canny edge detector," *Pattern Recognit.*, vol. 34, no. 3, pp. 721–725, 2001.
- [23] R. Bracewell, *Two-Dimensional Imaging*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1995.

- [24] G. Beylkin, "Discrete radon transform," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 35, no. 2, pp. 162–172, Feb. 1987.
- [25] S. Kumar and M. Hebert, "Discriminative random fields," *Int. J. Comput. Vis.*, vol. 68, no. 2, pp. 179–201, 2006.
- [26] J. D. Lafferty, A. McCallum, and F. C. N. Pereira, "Conditional random fields: Probabilistic models for segmenting and labeling sequence data," in *Proc. 8th Int. Conf. Mach. Learn.*, 2001, pp. 282–289.
- [27] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 3rd ed. New York, NY, USA: McGraw-Hill, 1991.
- [28] R. Szeliski *et al.*, "A comparative study of energy minimization methods for Markov random fields with smoothness-based priors," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 6, pp. 1068–1080, Jun. 2008.
- [29] C. Rother, V. Kolmogorov, V. Lempitsky, and M. Szummer, "Optimizing binary MRFs via extended roof duality," in *Proc. IEEE Conf. CVPR*, Jun. 2007, pp. 1–8.
- [30] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient belief propagation for early vision," *Int. J. Comput. Vis.*, vol. 70, no. 1, pp. 41–54, 2006.
- [31] M. N. M. van Lieshout and R. S. Stoica, "The Candy model revisited: Markov properties and inference," CWI, Amsterdam, The Netherlands, Res. Rep. PNA-R0115, 2001.
- [32] M. Rochery, I. H. Jermyn, and J. Zerubia, "Higher order active contours," *Int. J. Comput. Vis.*, vol. 69, no. 1, pp. 27–42, 2006.
- [33] L. C. C. Wai and M. Brady, "Curvilinear structure based mammographic registration," in *Computer Vision for Biomedical Image Applications* (Lecture Notes in Computer Science), Y. Liu, T. Jiang, and C. Zhang, Eds. Berlin, Germany: Springer-Verlag, 2005, pp. 261–270.
- [34] M. G. Linguraru, K. Marias, R. English, and M. Brady, "A biologically inspired algorithm for microcalcification cluster detection," *Med. Image Anal.*, vol. 10, no. 6, pp. 850–862, 2006.
- [35] R. Zwiggelar, S. M. Astley, C. R. Boggis, and C. J. Taylor, "Linear structures in mammographic images: Detection and classification," *IEEE Trans. Med. Imag.*, vol. 23, no. 9, pp. 1077–1086, Sep. 2004.
- [36] M. Berks, Z. Chen, S. Astley, and C. Taylor, "Detecting and classifying linear structures in mammograms using random forests," in *Information Processing in Medical Imaging* (Lecture Notes in Computer Science), vol. 6801, G. Székely and H. K. Hahn, Eds. Berlin, Germany: Springer-Verlag, 2011, pp. 510–524.
- [37] F. Moayed, Z. Azimifar, R. Boostani, and S. Katebi, "Contourlet-based mammography mass classification," in *Image Analysis and Recognition* (Lecture Notes in Computer Science), vol. 4633, M. Kamel and A. Campilho, Eds. Berlin, Germany: Springer-Verlag, 2007, pp. 923–934.
- [38] G. S. Muralidhar *et al.*, "Snakules: A model-based active contour algorithm for the annotation of spicules on mammography," *IEEE Trans. Med. Imag.*, vol. 29, no. 10, pp. 1768–1780, Oct. 2010.
- [39] J. Tang, R. M. Rangayyan, J. Xu, I. El Naqa, and Y. Yang, "Computer-aided detection and diagnosis of breast cancer with mammography: Recent advances," *IEEE Trans. Inf. Technol. Biomed.*, vol. 13, no. 2, pp. 236–251, Mar. 2009.
- [40] M. Heath, K. Bowyer, D. Kopans, R. Moore, and W. P. Kegelmeyer, "The digital database for screening mammography," in *Proc. 5th Int. Workshop Digit. Mammograph.*, 2001, pp. 212–218. [Online]. Available: <http://marathon.csee.usf.edu/Mammography/Database.html>
- [41] V. A. Krylov, S. Taylor, and J. D. B. Nelson, "Stochastic extraction of elongated curvilinear structures in mammographic images," in *Image Analysis and Recognition* (Lecture Notes in Computer Science), vol. 7950, M. Kamel and A. Campilho, Eds. Berlin, Germany: Springer-Verlag, 2013, pp. 475–484.
- [42] N. Merlet and J. Zerubia, "New prospects in line detection by dynamic programming," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 4, pp. 426–431, Apr. 1996.
- [43] Q. Zhang and I. Couloigner, "Accurate centerline detection and line width estimation of thick lines using the radon transform," *IEEE Trans. Image Process.*, vol. 16, no. 2, pp. 310–316, Feb. 2007.
- [44] D. Geman and B. Jedynak, "An active testing model for tracking roads in satellite images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 1, pp. 1–14, Jan. 1996.
- [45] D. Chai, W. Forstner, and F. Lafarge, "Recovering line-networks in images by junction-point processes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2013, pp. 1894–1901.
- [46] F. Lafarge, G. Gimel'farb, and X. Descombes, "Geometric feature extraction by a multimarked point process," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 9, pp. 1597–1609, Sep. 2010.
- [47] D.-S. Huang, W. Jia, and D. Zhang, "Palmprint verification based on principal lines," *Pattern Recognit.*, vol. 41, no. 4, pp. 1316–1328, Apr. 2008.
- [48] X. Wu, D. Zhang, and K. Wang, "Palm line extraction and matching for personal authentication," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 36, no. 5, pp. 978–987, Sep. 2006.
- [49] D. Zhang and W. Shu, "Two novel characteristics in palmprint verification: Datum point invariance and line feature matching," *Pattern Recognit.*, vol. 32, no. 4, pp. 691–702, 1999.
- [50] N. Duta, A. K. Jain, and K. V. Mardia, "Matching of palmprints," *Pattern Recognit. Lett.*, vol. 23, no. 4, pp. 477–485, 2002.
- [51] L. Zhang and D. Zhang, "Characterization of palmprints by wavelet signatures via directional context modeling," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 34, no. 3, pp. 1335–1347, Jun. 2004.
- [52] COEP, Pune, India. *Autonomous Institute of Government of Maharashtra*. [Online]. Available: <http://www.coep.org.in/>, accessed on 15 Aug. 2014.
- [53] T. Lindeberg, "Feature detection with automatic scale selection," *Int. J. Comput. Vis.*, vol. 30, no. 2, pp. 79–116, 1998.



Vladimir A. Krylov received the "specialist" (M.Sc.) degree in applied mathematics and computer science, and the "candidate of physico-mathematical sciences" (Ph.D.) degree in statistics from Lomonosov Moscow State University, Moscow, Russia, in 2007 and 2011, respectively. From 2011 to 2012, he was a Post-Doctoral Fellow with the Ariana and Ayin Research Teams, National Institute for Research in Computer Science and Control (INRIA), Sophia Antipolis, France, and a Research Associate with the Department of Statistical Science, University College London, London, U.K., from 2012 to 2013. Since 2014, he has been with the Department of Electrical, Electronic, Telecommunications Engineering and Naval Architecture, University of Genoa, Genoa, Italy.

His research interests are in the field of statistical signal processing and pattern recognition applied to medical and remote sensing imagery.



James D. B. Nelson joined the Department of Statistical Science at University College London as a Lecturer (Assistant Professor) in 2010 and became Senior Lecturer (Associate Professor) in 2013. After a Ph.D. in applied harmonic analysis from the Mathematics Department at Anglia Polytechnic University (1998–2001), he held post-doc positions in: the Applied Mathematics and Computing Group at the University of Cranfield (2001–2004); the Information: Signals, Images, and Systems Research Group at the University of Southampton (2004–2006); and the Signal Processing and Communications Laboratory at the University of Cambridge (2006–2010).

Dr. Nelson's interests span aspects of harmonic analysis, signal processing, computational statistics, and machine learning.