



HAL
open science

Robust and Accurate Vectorization of Line Drawings

Xavier Hilaire, Karl Tombre

► **To cite this version:**

Xavier Hilaire, Karl Tombre. Robust and Accurate Vectorization of Line Drawings. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2006, 28 (6), pp.890-904. inria-00000394v1

HAL Id: inria-00000394

<https://inria.hal.science/inria-00000394v1>

Submitted on 4 Oct 2005 (v1), last revised 15 Aug 2006 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Robust and Accurate Vectorization of Line Drawings

Xavier Hilaire, LORIA-UHP

Karl Tombre, LORIA-INPL

LORIA, B.P. 239, 54506 Vandœuvre-lès-Nancy CEDEX, France

Xavier Hilaire is currently research and teaching assistant at Université Henri Poincaré Nancy 1. Most of the work described in this article was undertaken while he was employed by FS2i, 8 impasse de Toulouse, BP 141, 78001 Versailles Cedex, France, under a CIFRE contract while preparing his PhD.

Karl Tombre is professor at École des Mines de Nancy, INPL, France.

Abstract

This paper presents a method for vectorizing the graphical parts of paper-based line drawings. The method consists in separating the input binary image into layers of homogeneous thickness, skeletonizing each layer, segmenting the skeleton by a method based on random sampling, and simplifying the result. The segmentation method is robust with a best bound of 50% noise reached for indefinitely long primitives. Accurate estimation of the recognized vector's parameters is enabled by explicitly computing their feasibility domains. Theoretical performance analysis and expression of the complexity of the segmentation method are derived. Experimental results and comparisons with other vectorization systems are also provided.

Index Terms

Document analysis, Graphics recognition and interpretation, vectorization, curve segmentation, performance evaluation, line drawings.

I. INTRODUCTION

The automated conversion of paper-based documents into a set of features which can be edited and stored by some document management software is a problem that has a long history, and which has received considerable attention during the last decades. In the case of line drawings, the aim is usually to convert the graphics represented by pixels into vectors and other simple geometrical features. This problem of *vectorization*, i.e. raster-to-vector conversion, has been at the center of research in graphics recognition for many years already. A number of techniques have been proposed to solve this problem, leading to commercial software packages as well as to research prototypes. All these systems provide quite acceptable results on simple, moderately noisy documents, but recognition rates for poor quality and huge documents such as technical drawings hardly exceed 80% in practice, leading to high post-vectorization editing costs [7], [32].

Several reasons may explain this. First, it is very unlikely that a “blind” segmentation process can provide satisfying results, not only because the optimal solution for the segmentation is usually not known, but also because it is generally not unique. Introducing ad hoc rules or contextual expert knowledge to obtain better quality solutions therefore appears mandatory, and several authors have explored this pathway [5], [22], [27], [39]. However, due to potential

combinatorial explosion, such approaches may only provide limited improvements, and still rely on a good quality low-level vectorization. Also, both the robustness and scale invariance of most of the cited methods are not known, and only little (if any) theoretical analysis of their performances have been established so far. At last, a crucial point in most existing methods is the requirement to tune an important set of parameters, often set on completely empirical bases.

The context in which the present work was initiated added another constraint on us, that of high accuracy. The raster-to-vector conversion method we designed had to be used for recovering the dimensions and positions of the different components of architectural drawings with high precision. It was therefore not enough to detect vectors; their position had to be very accurate with respect to the original drawing.

All this convinced us that a *robust* and *accurate* vectorization method, involving only a limited set of parameters [42], and including a theoretical analysis of its behavior, would bring a significant contribution to the field. This paper aims at presenting such a method. Like most raster-to-vector conversion systems, it is currently limited to the detection of straight line segments and circular arcs, the two mostly used shapes in technical documents. There is however no absolute limitation in the proposed approach which would prevent it from detecting other kinds of primitives, although we suspect it to become prohibitively costly in terms of computational complexity.

The method has been implemented and tested on a wide variety of line drawings. In particular, we have made use of the database set up for the international graphics recognition contests organized during the IAPR international workshops on graphics recognition [7], [32] to provide comparative performance evaluations with respect to four other systems. A probabilistic analysis of the method's behavior is also provided.

The remainder of this paper is organized as follows. In section II, we briefly review some other proposed approaches, and point out their main weaknesses. In section III, we provide the necessary material for skeleton-based vectorization. In section IV, we describe the details of our method and give a theoretical analysis of its behavior. Section V details the complexity of the method. Experimental results and performance comparisons are provided by section VI, and the paper closes with a discussion about lessons learned and further perspectives of this work in section VII.

II. RELATED WORK

Most vectorization methods consist of several processing steps, including *finding the lines* in the original raster image, *segmenting the lines* found into a set of vectors and arcs, and performing various post-processing treatments, to find better positions for the junction points, to merge some vectors and remove some others, etc.

As we mentioned in section I, post-processing mostly consists in adding some kind of contextual knowledge. This includes simple heuristics to correct the result, setting junctions straight, merging those which are close to each other, reconnecting lines split up by a missing pixel [5], or taking into account the nature of the drawing to progressively simplify it [39]. Although such approaches can significantly improve raw results, they tend to introduce a number of additional thresholds and parameters. Hence, they will not be considered here, because of our need for a robust and non-contextual method, with as few parameters and thresholds as possible. In this section, we will therefore concentrate our review of literature on the two first steps in vectorization.

A. *Finding the lines: raster-to-vector conversion*

Raster-to-vector methods may be roughly classified into three families, depending on the basic technique they rely on: parametric model fitting, tracking and matching opposite contours, and skeletonization.

Parametric model fitting consists in using a line model to detect the lines present in the image. The most general and best known technique for that is the Hough transform, which has nevertheless not been widely used for raster-to-vector conversion, except some attempts by Dori [10], Kultanen et al. [25], Yamada [51], Song and Lyu [37]. These methods have only been evaluated on simple or very specific cases. The Hough transform is a global and additive transformation of an image. It does not guarantee that two noisy discrete lines having very close respective slopes won't mix and map into a dummy peak in parameter space. It is time efficient, but memory inefficient. However, Song and Lyu have recently presented a vectorization making use of the Hough transform, but adapted to the analysis of line drawings. Their idea is to perform a subsampling of the image by finding interest points through a horizontal and a vertical scan of the image. Only these points are used in the Hough vote. In addition, the neighborhood of the points is analyzed to only vote for plausible angles. All this reduces the computation time and

memory needs of the method and yields promising results. Still, the methods cited have been designed to extract lines only, not lines and circular patterns.

Other parametric models have been applied to domain-specific applications such as finding staff lines [8] or regular lines in forms [52]. This is related to vectorization methods using higher-level contextual knowledge, and therefore not relevant for the present work, which aims at designing a general method, without a priori contextual knowledge about the positions or orientations of the lines and arcs.

Tracking and matching opposite contours is an alternative technique that was tested in the late 80's [1], [36]. It provides correct and accurate results if the junctions are simple. However, it is noise sensitive, and fails when it comes to match complicated structures, even at the price of an expensive computation time [42].

Skeletonization definitely remains the most widely used technique for raster to vector conversion [4], [20], [24], [30]. The basic idea is simply to compute some kind of medial axis of the shape to be vectorized, by thinning it or by looking for the local extrema in its distance transform. This relies on the assumption that the shapes to be vectorized are elongated shapes. Once the skeleton has been computed, the problem of vectorization is reduced to that of segmenting a 2D discrete curve into meaningful features.

Unfortunately, the problem is complicated by the crossings and junctions between thick shapes, as parabolic arcs appear in the continuous space for usual crossings involving thick segments [3], regardless of the type of skeletonization methodology used (for an overview of thinning methodologies, see [26], and for non-thinning based methods see [47]). Most of the existing methods have only limited ability to detect such parabolic arcs; they approximate them by vectors, which obviously weakens the quality of the solution and is confusing for higher level treatments (namely, recognition and analysis tasks). Early attempts to overcome this drawback were performed by Hori and Tanigawa [19], who proposed to combine skeletonization with contour matching. The critical point in their method resides in the limitations of matching at the junctions. Another approach is that of Janssen and Vossepoel [20], who propose an iterative algorithm based on maximum threshold morphology to adjust the vectors' extremities. Their method performs well on simple junctions, but has the main drawback of relying on local adjustments, which makes it sensitive to noise. Sensitivity to noise is also a critical point in SPV, proposed by Dori and Liu [11]. Their method acts as a simplified and fastened skeletonization

operator, as the orthogonal zig-zag walk only visits a few pixels on the shapes' boundaries. Junctions are detected by the presence of an abnormal border-to-border length with respect to the average length observed during the walk; as a result, a junction involving only small segments might be skipped. The method also requires the settings of more than ten thresholds.

All the above methods have only limited chances to properly process junctions, and are noise sensitive. Another more complicated problem is that of extracting circular arcs from the image [21], [48]. This is generally achieved by testing whether a subset of adjacent segments can constitute a circular arc or not; scalability and noise sensitivity of such methods have unfortunately not been examined.

B. Curve segmentation

With respect to the segmentation of the 2D curves resulting from the previous step, the literature offers a considerable number of methods, which can be classified in two categories. First, there are methods using curvature estimation to detect critical points on the curve [40], sometimes taking scale space into account [2], [29]. Such methods perform well on curves that are not too noisy, and whose scale variation is reasonable, but do not explicitly guarantee any upper bound on the most usual error metrics (such as the integral square error, Hausdorff distance, etc.). Very soft transitions between a pair of lines or lines and circles, with very large radii which are commonly found in technical drawings, make such methods unsuitable for this class of documents.

Secondly, some methods operate directly on the curve, either by minimizing a particular error measure, or by splitting or breaking the curve as long as the local or global error remains too high. Such methods are known as polygonal approximation methods [34], and are more suitable to the problem of vectorization because of the guarantees they provide. The most usual criterion used by these methods is that of the maximum distance between the curve and the segment [35]. This leads to recursive splitting of the curve at the maximum deviation points, until all segments are valid approximations. As the position of the extrema points of the segments tends to be constrained by the initial pixel positions, it makes sense for the method to be followed by a fitting phase, where each segment is displaced to best fit the original curve. A second possible criterion is that of the algebraic area of the surface between the curve and the segment [46]. As this area can be computed iteratively, very time-efficient iterative methods have been proposed.

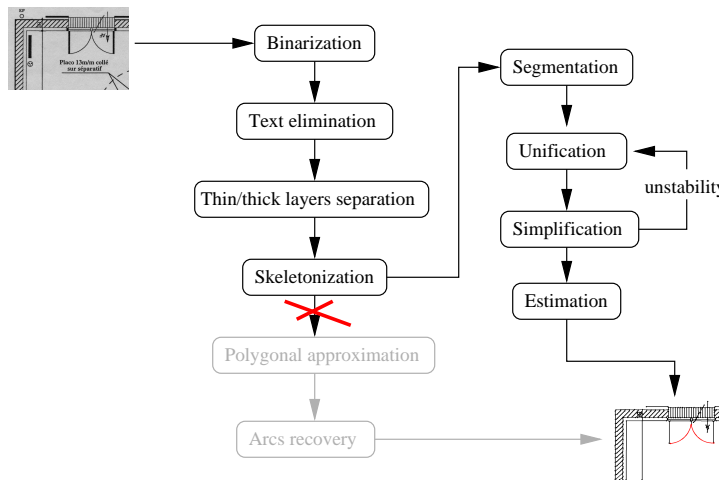


Fig. 1. Overview of our vectorization method with changes w.r.t [13] emphasized.

A problem with the approach is that it tends to displace the angular points, as the method only detects a change of general direction after having followed several points past the true angle.

However, whatever criterion is chosen, these approximation methods still have lacks with respect to the requirements of the vectorization problem, as a good vectorization method should simultaneously extract circular arcs and segments, while being scale invariant and robust to digitization noise as well as to text potentially attached to the graphics, and should not represent any of the parabolic arcs that appear in a crossing of primitives as a result of the skeletonization process.

III. BACKGROUND

The vectorization method we propose improves that of Dosch et al. [13], and is fully summarized in figure 1. The reader must note that the first four stages of this method (binarization, text elimination, thin/thick layers separation, and skeletonization) should therefore not be considered as a contribution of the present work. They are recalled here for the sole sake of clarity.

A. Binarization and filtering

We regard binarization as an optional stage, as most of the source images to be analyzed are directly available in binary format. In our system, the binarization stage is followed by an elementary filtering procedure that fills the holes that the graphics might carry, removes blobs

that may have appeared outside of them, and performs a morphological closure. It is possible to obtain a bound on the surface of spurious holes by resorting to Kanungo’s document degradation model [23], and to the following result, proved in [17] (keeping Kanungo’s notations):

Lemma 1: The expectation of the size of a fake hole (due to degradation) is lower than $\frac{\alpha_0 e^{-\alpha+\eta}}{(\alpha_0 e^{-\alpha+\eta}-1)^2}$

The proof is trivial and is not reproduced here. A similar result may be obtained for blobs.

B. Text elimination

The kind of documents we have to process often contains a text layer which is not to be vectorized. In addition, this layer may interfere with the segmentation of the graphics layer into curves when text touches the graphics. We therefore use the text/graphics separation method of Tombre et al. [44], which is based on a method first proposed by Fletcher and Kasturi [16], with a number of improvements. One of these is the ability to retrieve in the text layer characters connected to graphics, provided they belong to a string for which at least some characters are disconnected and have thus been detected as a “string seed”.

As we will see in section VI-C, under certain conditions, our method is actually capable of retrieving the correct segments and arcs even when the assumption of string seed presence does not hold.

C. Thin/thick layer segmentation

Line drawings can be made by the superimposition of several layers with different thicknesses, each layer carrying some specific information. In that case, our system offers the option of working independently on each thickness layer. In our implementation, we use mathematical morphology operations to perform thin-thick segmentation: an erosion is followed by a partial geodesic reconstruction [13]. Alternatively, this could also be done through segmentation of the distance transform.

Our method actually does not need this and is able to correctly segment a line drawing with different line thicknesses.

D. Skeletonization

There are many skeletonization algorithms [26], either based on the iterative deletion of the most outer pixels, or on a distance transform and medial axis computation. In our vectorization

system, we use Sanniti di Baja's skeletonization algorithm [9] for three simple reasons. First, the metric used in [9] is the (3,4)-chamfer distance, which effectively induces a metric over \mathbb{Z}^2 [41]. This means that given a discrete thick curve in \mathbb{Z}^2 , we can expect its skeleton to coincide with a thin curve of the same nature within a gap of one pixel. Second, the skeleton is reversible, and the computation of the whole distance transform will allow us to estimate the thickness of each extracted shape, as explained further, in section IV. Third, the skeleton is robust. The detection of saddle points, the computation of the medial axis from the distance map and its reduction to unit width using positive gradients as described in [9], ensures that skeletal points will always be chosen in the same manner whenever a choice is possible. In particular, adding noise points to the boundaries of a shape do not generally lead to changes in the resulting skeleton [17].

IV. PROPOSED VECTORIZATION METHOD

A. Segmenting the skeleton

1) *Hypothesis, definitions, and notations:* We assume the documents to be preprocessed as described in section III, until skeletonization. If the option of segmentation into different thickness layers has been chosen, the following discussion applies to one such layer at a time. The skeleton computed with Sanniti di Baja's algorithm is connected. For that reason, we will only have to deal with connected discrete arcs and skeletal branches, whose definitions are recalled above.

Definition 1: A discrete simple arc (DSA) is a sequence $((x_1, y_1), \dots, (x_n, y_n))$ of points of \mathbb{Z}^2 such as $1 \leq |x_{i+1} - x_i| + |y_{i+1} - y_i| \leq 2$ for $i = 1, \dots, n-1$. Furthermore, if $|x_n - x_1| + |y_n - y_1| > 2$, the DSA is said to be open.

Definition 2: A discrete skeletal branch (or skeletal branch, for short) is a DSA whose points (x_i, y_i) are skeletal points, and such that (x_1, y_1) and (x_n, y_n) are skeletal endpoints. A skeletal endpoint is a skeletal point having either exactly one or at least three 8-connected neighbors.

In the sequel, we will deal with segmentation of skeletal branches only. It is assumed that an upper bound f of the line thickness is known and that the skeleton has been pruned using this value.

Let \mathcal{C} be a skeletal branch, we denote by $\mathcal{C}[k]$ its k th pixel, and by $\mathcal{C}[i..j]$ the subarc of \mathcal{C} containing pixels $\mathcal{C}[i]$ to $\mathcal{C}[j]$, $i \leq j$. The Euclidean distance is denoted by $d(., .)$, and $|\cdot|$ denotes either the length (in number of pixels) or cardinality. For any *continuous* shape S of the plane,




Fig. 2. Definition of fuzzy primitives: (a), segment; (b), circular arc. White pixels, that do not belong to $D(S)$, are discarded and replaced by the proper pixels once the primitive is identified.

we denote by $D(S)$ its digitization according to the OBQ¹ scheme (the shape is assumed to be a part of the frontier of a closed, planar object). As stated in section I, we are motivated by a robust segmentation of the skeleton. We therefore propose to segment any skeletal branch \mathcal{C} using *fuzzy segments* and *fuzzy circular arcs*, whose definitions follow.

Definition 3: A DSA \mathcal{C} is a fuzzy segment (FS) if there exists a straight line S of the real plane such that

- (i) $|D(S) \cap \mathcal{C}| \geq |\mathcal{C}|/2$
- (ii) $\forall X \in \mathcal{C}, \quad d(X, S) \leq m$, where m is a parameter to specify.

Definition 4: A DSA \mathcal{C} is a fuzzy circular arc (FCA) if there exists a circular arc S of the real plane with radius ρ and opening angle θ such that

- (i) $|D(S) \cap \mathcal{C}| \geq |\mathcal{C}|/2$
- (ii) $\forall X \in \mathcal{C}, \quad d(X, S) \leq m$, where m is a parameter to specify
- (iii) $\theta \geq \theta_{\min}$, and $\rho \in [\rho_{\min}, \rho_{\max}]$, where θ_{\min} , ρ_{\min} and ρ_{\max} are constants to specify.

The underlying idea in both definitions is that we accept a given chain of pixels as a segment or as an arc, provided we can find a continuous shape S around which pixels stand closer than a certain distance m (see Fig. 2). For circular arcs, we furthermore require that S satisfies the validity criteria imposed by θ_{\min} , ρ_{\min} and ρ_{\max} . At last, we require that at least half of the pixels of a candidate arc coincide with the digitization of S , which is the largest amount of invalid data we can allow without possible confusion (in the Cramer-Rao sense).

¹Object Boundary Quantization.

2) *Main procedure*: The pseudo-code for the main procedure of our algorithm is available in Alg. 1, in which f is the thickness of the processed layer and L is initially an empty list. The algorithm is based on random sampling [15] and operates as follows (see Fig. 3):

- 1) Two indexes i and j are randomly chosen along \mathcal{C} .
- 2) The subarc $\mathcal{C}[i..j]$ is then tested (lines 6–9), first as a FS, next as a FCA, in that order. If it is found to be one or another, then the corresponding **extract_** function is called, which finds the lowest and largest integer indexes i_0 and j_0 such as $i_0 \leq i$, $j_0 \geq j$, and $\mathcal{C}[i_0..j_0]$ is still a fuzzy primitive of the right type. The i_b and j_b indexes, which denote indexes of the longest primitive found, are updated accordingly.
- 3) If there is no primitive instantiable from i, j , or if the longest primitive found does not exceed a threshold value T , then we jump to step 1 to make a new attempt, unless the number of attempts has reached n_{max} .
- 4) If a primitive could be found by the following steps (regardless of it being longer than T or not), then it is added to L , and the segmentation is reconducted on the remaining parts of \mathcal{C} . Otherwise, we terminate in failure and apply an alternative method, namely the split and merge segmentation algorithm [33].

In this procedure, T is a confidence threshold on length upon which we accept any found primitive, and break the repeat loop in order to accelerate the process. In other words: if we can find any primitive longer than T , then we retain it, otherwise we loop at most n_{max} times to retain the longest primitive we found. Setting proper values for T and n_{max} is therefore highly desirable; we detail how to do this in section V-C.

3) *Fuzzy primitive test and extension*: The **test_** and **extract_** functions used in Alg. 1 both require to test whether a given discrete arc \mathcal{C} is a FS or a FCA. According to definitions 3 and 4, this question may only be answered by building and updating the corresponding feasibility domains in the dual space, which requires $O(n^d)$ time with $d = 2$ for FS and $d = 3$ for FCA.

Rather than using such a costly procedure, we slightly weaken the test by using *linear regression*, that is, we obtain S in definitions 3 and 4 by fitting a line or a circle to \mathcal{C} using the algebraic least squares method. This approach is justified by the two following lemmas, both proved in [17]:

Lemma 2: Let \mathcal{C} be a FS or FCA, \hat{T} be the estimate shape of \mathcal{C} using the least square method, and $\hat{\mathcal{C}}$ be the discrete arc obtained by rounding each point of \hat{T} to the nearest integer couple in




Fig. 3. (a) Illustration of the random sampling procedure used in Alg. 1. (b) The gray pixels may either belong to the lines or to their closest parabolic arc in the case of an L junction

\mathbb{Z}^2 . Then it holds that

$$\frac{|\mathcal{C} \cap \hat{\mathcal{C}}|}{|\mathcal{C}|} = 1 - O(1/|\mathcal{C}|)$$

Lemma 3: Addition of bounded and centered noise to a FS or FCA \mathcal{C} does neither change the conditions, nor the values towards which their respective least squares estimates converge as $|\mathcal{C}| \rightarrow \infty$.

Both lemmas ensure that recovering all the valid pixels of a fuzzy primitive is always possible when its length becomes arbitrary large. We conducted and compared both approaches and found only little differences for sufficiently large primitives (length ≥ 10 pixels). Therefore, the results reported in section VI follow the linear regression approach.

B. Optimization

The result of the segmentation stage is a set of primitives that approximate the skeleton of a particular layer of the document. This set, however, may not yet be considered as the solution to vectorization for two obvious reasons: (1) we segmented the skeleton of the image, not the image itself; as a result, we still have to get rid of the skeletal arcs that appear in crossings of primitives, and (2) we considered only skeletal branches, which implies that any primitive of the ground truth has been detected as fragmented if it is hit by any other at a crossing.

Algorithm 1 Pseudo-code for the main procedure of our segmentation method.

```

1: Function partition( $\mathcal{C}, m$ )
2: Let  $l \leftarrow |\mathcal{C}|$ ,  $n \leftarrow 0$ ,  $(i_b, j_b) \leftarrow (0, 0)$ 
3: repeat
4:   Randomly choose  $i$  and  $j$  such as  $1 \leq i < j \leq l$ 
5:   stop  $\leftarrow$  false,  $(i_0, j_0) \leftarrow (0, 0)$ 
6:   if test_segment( $\mathcal{C}, i, j, m$ ) then
7:      $(i_0, j_0) \leftarrow$  extract_segment( $\mathcal{C}, i, j, m$ )
8:   else if test_arc( $\mathcal{C}, i, j, m$ ) then
9:      $(i_0, j_0) \leftarrow$  extract_arc( $\mathcal{C}, i, j, m$ )
10:  else
11:     $n \leftarrow n + 1$ 
12:  end if
13:  if  $j_0 - i_0 > j_b - i_b$  then
14:     $(i_b, j_b) \leftarrow (i_0, j_0)$ 
15:  end if
16: until  $n \geq n_{max}$  or  $j_b - i_b > T$ 
17: if  $j_b - i_b \neq 0$  then
18:   Add the found primitive  $\mathcal{C}[i_b..j_b]$  to  $L$ 
19:   if  $i_b \geq 2$  then Segment( $\mathcal{C}[1..i_b], m$ ) end if
20:   if  $j_b \geq l - 1$  then Segment( $\mathcal{C}[j_b..l], m$ ) end if
21: else
22:   Segment  $\mathcal{C}$  using the split and merge algorithm, and add the result to  $L$ .
23: end if
24: End

```

The optimization stage addresses these problems. It consists of a *simplification* and a *unification* procedure that are applied in turn, independently. Both procedures share a connectivity graph G built immediately after the segmentation stage; this graph has the following properties:

- its nodes are the primitives found upon completion of the segmentation stage.

- any of its edges denotes a connection between two primitives. The edges are unoriented.

1) *Simplification of the skeleton*: The *simplification procedure* eliminates the spurious arcs that appeared inside the crossings of primitives during skeletonization. It uses the basic underlying assumption that any discrete arc with length lower than an acceptance threshold f may be removed provided that the remaining primitives are sufficient to explain and reconstruct the image. To detail this procedure, we first give a few definitions and results from [17].

Definition 5: The preimage of a DSA \mathcal{X} is the set of all the primitives of the plane whose digitization fully includes \mathcal{X} . It is denoted $Pre(\mathcal{X})$.

Dorst and Smeulders [12] gave a characterization of the feasibility domain of discrete segments in the dual space: it consists of a quadrilateral with at most 4 vertices (Fig. 4). From their result, it immediately follows that the preimage of a discrete segment can be represented using at most two segments and four half-lines. For discrete circular arcs, we proved in [17] that the preimage could only have three possible shapes, using a lemma from O'Rourke et al. [31]. These shapes are presented in Fig. 5.




Fig. 4. The domain (a) and the preimage (b) of a discrete segment

We now consider the intersection of two discrete primitives, that we define in a particular manner:

Definition 6: Let \mathcal{P} and \mathcal{Q} be two discrete primitives, let L and R denote the integer endpoints of \mathcal{P} , and let \bar{I} denote the complement of the source image. We say that \mathcal{P} joins \mathcal{Q} by L (resp. by R), and denote $\mathcal{P} \vdash_L \mathcal{Q}$ (resp. $\mathcal{P} \vdash_R \mathcal{Q}$), if and only if the set $Pre(\mathcal{P}) \cap Pre(\mathcal{Q})$ is made of at least one closed connected component Δ_L (resp. Δ_R) for which the following condition holds: there exists one integer point $K \in \Delta_L$ and one discrete primitive \mathcal{X} of the same nature




Fig. 5. The three different ways to represent the part of the preimage of a discrete circular arc falling outside its convex hull: (a) degenerated case (the circular arc is a discrete segment), (b) full circle, and (c) regular case.

as \mathcal{Q} passing through K such that $\mathcal{Q} \subseteq \mathcal{X}$, $R \in \mathcal{X}$ (resp. $L \in \mathcal{X}$), and $\mathcal{X} \cap \bar{I} = \emptyset$. Δ is called the L -endpoint domain (resp. L -endpoint domain) of \mathcal{P} .

An illustration of this definition is given in Fig. 6. In each case, the idea is the same: \mathcal{P} and \mathcal{Q} being given, if we can find a piece of discrete shape that completely overlaps \mathcal{P} and permits to reach \mathcal{Q} while remaining inside the image, then we have $\mathcal{P} \vdash_L \mathcal{Q}$ or $\mathcal{P} \vdash_R \mathcal{Q}$, or possibly both, depending on the nature of \mathcal{P} and \mathcal{Q} and on which points the linking piece of primitive stems from. In such a case, Δ_R (resp. Δ_L) simply defines the feasible region of the plane for the endpoint R (resp. L). In the following, we will always assume that any discrete primitive has two endpoints L and R : the full discrete circle can be viewed as a circular arc whose endpoints are identical, but it can be processed exactly the same way as segments and open circular arcs.

With the above material, we are now ready to explain the **simplify** procedure detailed in Alg. 3. This procedure operates in two stages. First, for each *long* primitive ($length \geq f$), we seek for similar neighbors primitives that can be reached through a chain of *short* primitives – those *might* be removed. This is done by the **search** function, whose pseudo-code is available in Alg. 2, and which returns the list of all possible paths stemming from a given node n and leading either to a long primitive or to an end of chain of short primitives. To each node that represents a long primitive, the list of these paths is being attached. For any node n , we process all the paths P we find, and look at p , the last element of the path P . Several cases may then occur:

- p is a short primitive: the edges defined by the path P are simply marked as deletable;
- p is a long primitive and $p \vdash_L n$ (resp. $p \vdash_R n$): we redefine Δ_L (resp. Δ_R) as the proper component of $Pre(n) \cap Pre(p)$ following Def. 6. The endpoint L (resp. R) can then be

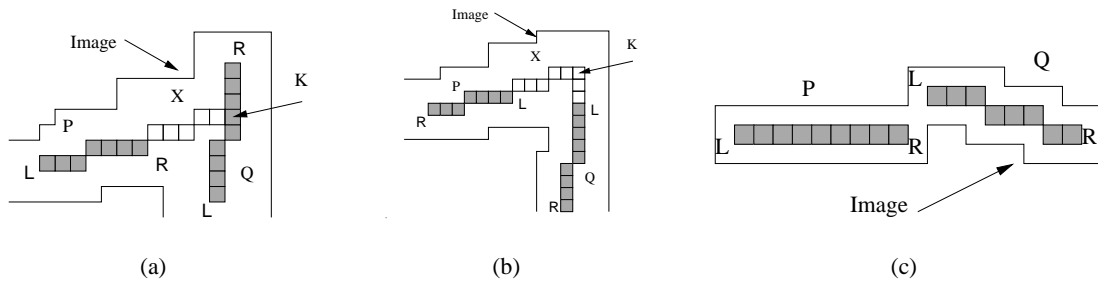


Fig. 6. Illustration of Def. 6. (a): $P \vdash_R Q$, but $P \not\vdash_L Q$, and $Q \not\vdash_R P$, $Q \not\vdash_L P$; (b): $P \vdash_L Q$ and $Q \vdash_L P$, but $P \not\vdash_R Q$ and $Q \not\vdash_R P$; (c) no joint relation at all.

chosen as any integer point of Δ_L (resp. Δ_R) in subsequent calls;

- p is a long primitive, but $p \not\vdash_L n$ and $p \not\vdash_R n$: nothing happens.

As the two first cases lead to a valid intersection, the edges formed by all the found paths between the two primitives are marked for deletion. Finally, all marked edges and all unconnected nodes of G are removed. Note that Alg. 3 is provided in its simplest form for the sake of clarity. It can be practically optimized in many ways, in particular:

- 1) the computations of $Pre(n) \cap Pre(g)$ for various n and p , and the information "does $n \vdash_L last(P)$ or $n \vdash_R last(P)$?", can be cached;
- 2) the redundancies of the deletions can also be avoided by resorting to some proper data structures;
- 3) most importantly: if the required endpoints have to be output as *integers*, which is actually the case for the VEC format [6] used in section VI, then *all the geometric operations* may be handled in \mathbb{Z}^2 rather than \mathbb{R}^2 .

Obviously, the latter point greatly simplifies the implementation of the method, as handling nonlinear pieces of curves for circles, as well as irrational representation and precision issues, becomes irrelevant. Our implementation does a little bit better: we allow a rational representation of all points by subsampling the various domain boundaries. All geometric operations are then carried out using standard algorithms on rational polygons. It is however noticeable that the method we describe is still valid in \mathbb{R}^2 , but at the price of much costlier algorithms due to precision issues.

Algorithm 2 Pseudo-code of the auxiliary **search** function used in the simplification stage.

```

1: Function search( $n, G$ )
2:  $L \leftarrow \emptyset$ ; mark  $n$ 
3: for  $x \in neighbors(n, G)$  such that  $x$  is not marked do
4:   if  $length(x) \geq f$  or  $(T \leftarrow search(x, G)) = \emptyset$  then
5:      $L \leftarrow L \cup \{x\}$ 
6:   else
7:     for  $t \in T$  do
8:        $L \leftarrow L \cup \{xt\}$ 
9:     end for
10:  end if
11: end for
12: unmark  $n$ ; return  $L$ 
13: End search

```

To illustrate how Alg. 3 works, consider the case of the X junction in Fig. 7. The branches in this junction have unequal thicknesses, and we may not expect them to meet in a single point (Fig. 7-(a)). We assume that A, B, C, D are long segments that have been properly recognized during the segmentation stage, whereas a, b, c, d, i, j, k are spurious arcs with length smaller than f (Fig. 7-(b)). We obtain the connectivity graph G (Fig. 7-(c)) and, node by node, the following paths: $A \rightarrow \{bB\}$, $B \rightarrow \{cdC\}$, $C \rightarrow \{dcB, dkD, dijD\}$, $D \rightarrow \{kdcB, kdC, jidC\}$. Starting from A , we obtain that A leads to B and that $A \vdash_R B$, so the graph is modified accordingly, as in Fig. 7-(d). From B , we obtain two candidates C and D , but only $B \vdash_R C$, so the edges Bc, cd, dC are deleted, and the arc BC is added. From C , we find three paths, among which one leads to B which has already been processed, and the two others lead to D . Since $C \vdash_R D$, we delete Cd, dk, kD, di, ij, jD and insert CD . Processing from D is the symmetrical case of processing from C and doesn't change anything. Finally, the remaining, unconnected nodes are simply deleted. We therefore obtain the connectivity graph of Fig. 7-(h), and the X junction is reconstructed as illustrated by Fig. 7-(i). It is worth noting that reconstructing a given junction generally leads to more than a single solution. It may be interesting to group primitives that

Algorithm 3 Pseudo-code for the **simplify** procedure of the optimization stage.

```

1: Procedure simplify( $G$ )
2: for  $n \in G$  such that  $length(n) \geq f$  do
3:   for  $P \in search(n, G)$  do
4:     let  $p$  be the last element of path  $P$ ;  $remove \leftarrow (length(p) < f)$ 
5:     if not  $remove$  then
6:       if  $p \vdash_L n$  then
7:         redefine  $\Delta_L(p)$  according to Def. 6;  $remove \leftarrow true$ 
8:       end if
9:       if  $p \vdash_R n$  then
10:        redefine  $\Delta_R(p)$  according to Def. 6;  $remove \leftarrow true$ 
11:       end if
12:     end if
13:     if  $remove$  then
14:       add edge  $(n, last(P))$  to  $G$  if it doesn't exist
15:       mark all edges defined by  $P$  for deletion
16:     end if
17:   end for
18: end for
19: remove all edges from  $G$  marked for deletion
20: remove all unconnected nodes from  $G$ 
21: End simplify

```

overlap each other as much as possible, as we suggested in [18]. It is even sometimes desirable to determine the concurrency of lines through the calculation of the maximal clique of their intersection graph [45], which is an NP-complete problem. The solution we suggest is therefore somewhat arbitrary, but has the advantage to preserve connectivity, and is able to find the "carrier" primitives in a junction – primitives can only be grown, never reduced.

2) *Unifying primitives*: Due to the crossing of primitives, we also need to cope with the problem of recovering from primitive fragmentation whenever it occurred. The unification test

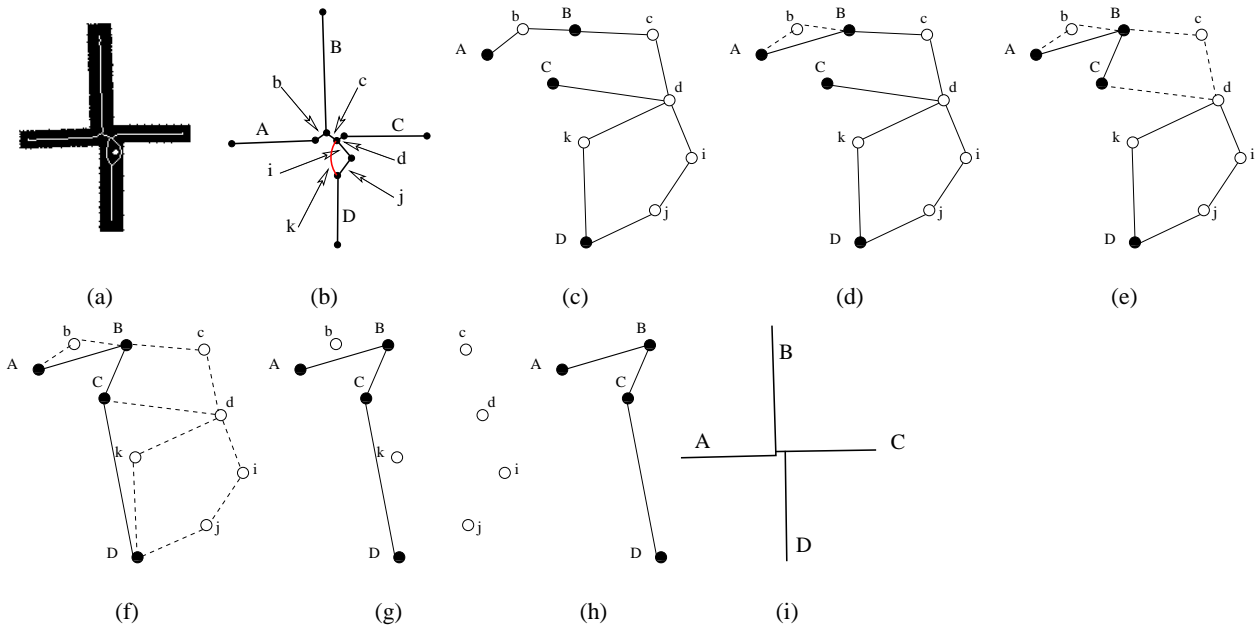


Fig. 7. Processing of an X junction by Alg. 3. (a) source image, (b) resulting segmentation, (c) connectivity graph G , (d) – (h) action of Alg. 3, (i) the reconstructed junction.

we use involves lemmas 2 and 3, and the following domain construction procedure: let \mathcal{Q} and \mathcal{P} be two adjacent fuzzy primitives, we unify them into a new primitive iff we can find a continuous segment or circle S such as

$$|D(\mathcal{Q}) \cap D(S)| \geq |D(\mathcal{Q})|/2 \text{ and } |D(\mathcal{P}) \cap D(S)| \geq |D(\mathcal{P})|/2 \quad (1)$$

Most of the time, unifying two primitives \mathcal{Q} and \mathcal{P} will result in a primitive S for which $Pre(S) \subset Pre(\mathcal{P}) \cup Pre(\mathcal{Q})$. As a result, some concurrency relations established during the simplification stage *may* have to be reconsidered. Indeed, two cases of figure may occur:

- we find that the endpoint domains of \mathcal{Q} and \mathcal{P} still have a non-empty intersection with $Pre(R)$: in this case, this intersection defines the new endpoint domains;
- we find that the previous intersection is empty: in this case, we simply disconnect the new primitive from those previously attached to its endpoints.

Disconnecting primitives in the latter case results in the processing of two independent primitives and is most of the time desirable (as an example, consider a full circle interrupted by a line). However, it may also become an arbitrary decision in certain cases (shall an X junction




Fig. 8. Use of the domains to estimate the parameters of (a) segments, and (b) circular arcs.

be reconstructed with four segments and a point, or only with two segments?). Also note that there exists an arbitrary order within which primitives eligible for unification can be processed. In our method, we first process those that are the most likely; that is, we maintain a table of all couples of adjacent primitives that satisfy Eq. 1, and always process first the pair that maximizes the quantity $|D(S) \cap D(Q)| + |D(S) \cap D(P)|$.

C. Final parameters estimation

This last stage aims at estimating the parameters of the found shapes and their thickness. It is done once, after all other stages.

1) *Estimating the shapes parameters:* Estimation of the shape parameters for the case of segments is straightforward from their domains. First, recall from section IV-B that, M being L or R , the fact that $\mathcal{P} \vdash_M \mathcal{Q}$ results in a non empty domain Δ_M (Fig. 8), inside which the endpoint M will have to lie. The optimal coordinates of M are the respective expectations of x and y over Δ_M if we assume a uniform distribution of the feasible points in \mathbb{R}^2 . In other words:

$$P = \frac{1}{\text{area}(\Delta_M)} \left(\int_{D \in \Delta_M} x(D) dD, \int_{D \in \Delta_M} y(D) dD \right)$$

Second, observe from the left of Fig. 8 that a similar domain may also be built in the case of a “free” endpoint M by considering the set of points in the neighborhood of M that don’t increase the digitization of the segment (or circular arc).

For circles, Δ_M is bounded by two circles with centers Ω_i and Ω_e (not necessarily identical) and radii ρ_i and ρ_e . The optimal circle is centered in the middle of Ω_i and Ω_e , and has radius $\frac{1}{2}(\rho_i + \rho_e)$. For circular arcs, the estimation is slightly more complicated, because the problem

is overconstrained. As for circles, we obtain an optimal estimated circle ξ with radius ρ (see Fig. 8(b)), but the problem is that the optimal endpoints L and R may not necessarily lie on it. We have therefore two possibilities: either we fix points L and R , and seek for a non optimal radius, or we determine ξ with an optimal radius, and seek for non optimal L and R points on it. We adopted the latter.

2) *Thickness estimation*: The thickness of the primitives remains a quite unreliable information in practice as it is known most of the time with only limited accuracy. Three simple reasons explain this: (1) there is an obvious dependency to the digitization and binarization schemes; (2) the shapes are noisy; (3) even for non noisy shapes and a fixed threshold binarization scheme, the theoretical precision is rather poor (for example, any real half-plane may be moved along the Ox or Oy axes of a distance in the interval $]\sqrt{2}/2, 1[$ depending of the plane orientation, without affecting its digitization).

The $(3, 4)$ -distance transform computed during skeletonization provides a straightforward solution to obtaining a lower bound to the thickness of a shape. Let I be the digital image and T be the $(3, 4)$ -distance distance of I . A point $p \in T$ with maximal label v ensures that there exists a discrete ball centered on p with radius v that is fully included in I . Two cases may then occur, depending on whether p is the sole point having the maximal label in its 8-neighborhood or not. If it is the case, then v is the value of the thickness with the $(3, 4)$ -distance. Otherwise, the local description of the shape is a union of balls, in which case we can deduce the thickness from the diameter of this union, and add a half of the current label's value to correct the estimate. At last, we also have to cope with noise and dummy characters or primitives touching the graphics, which may severely alter the values of the thickness. We therefore estimate the thickness of each segmented shape with skeleton \mathcal{C} as follows:

- 1) For any pixel $p = (i, j) \in \mathcal{C}$, let $E_p = \{(x, y) \in T : |x - i| + |y - j| \leq 2, T(x, y) = T(i, j)\}$, and let $v_p = \max_{(x, y) \in E_p} \{|x - i| + |y - j|\}$
- 2) Collect a set L of values by adding the following values for each $p \in \mathcal{C}$ to it:

$$\begin{cases} T(p) & \text{if } v_p = 0 \\ T(p) + 3/2 & \text{if } v_p = 1 \\ T(p) + 2 & \text{if } v_p = 2 \end{cases}$$

- 3) Retain the mean value of L as the thickness.

V. COMPLEXITY AND PERFORMANCE ANALYSIS

This section evaluates the complexity of the algorithms involved in our method. We also provide clues to set the T , and n_{\max} parameters for the segmentation step (Alg. 1). We assume an $h \times w$ pixels wide image, containing p connected components. The lowest maximal surface area of these components is denoted by s . The segmentation stage is assumed to process n skeletal branch, each formed of l pixels in the worst case, and N denotes the highest number of primitives carried by at least one of these branches.

A. Preprocessing steps

The first steps, which we “inherited” from previous work (see section III), do not present particular high complexity. Noise filtering is essentially linear. Text elimination is based on classical connected component analysis which can also be performed in linear time.

Thin-thick separation is performed by iterating elementary morphological operations a limited number of times. More precisely, let S be the source image, from which we wish to extract parts with thickness at least equal to some $e > 0$. If $r = \lfloor (e - 1)/2 \rfloor$, the operation consists in iterating r elementary erosions, followed by $r + 1$ dilations and intersections with the original image. The computational complexity for separating into two layers is therefore in $O(rhw)$.

B. Skeletonization

The time necessary to complete the skeletonization step is bounded by that of the computation of the (3,4)-chamfer distance. It turns out to be $O(hw)$ in a sequential scheme [41].

C. Skeleton segmentation

We provide a more detailed analysis of Alg. 1 in order to establish its complexity and assess its performances. Consider a fuzzy primitive \mathcal{C} with length l . Assuming stochastic noise, let α be the probability that a given point of \mathcal{C} is incorrect. Then the probability to obtain at least $\lceil l/2 \rceil$ incorrect points in \mathcal{C} is

$$P(n \geq \lceil l/2 \rceil | \mathcal{C}) = \sum_{k=\lceil l/2 \rceil}^l \binom{l}{k} \alpha^k (1 - \alpha)^{l-k} \quad (2)$$

$\alpha \setminus l$	5	10	15	20	25	30
0.45	0.4069	0.4956	0.3465	0.4086	0.3063	0.3552
0.35	0.2352	0.2485	0.1132	0.1218	0.06044	0.06519
0.25	0.1035	0.07813	0.0173	0.01386	0.00337	0.00275
0.15	0.02661	0.009874	0.0006096	0.0002484	1.689e-5	7.078e-6
0.05	0.001158	6.369e-5	1.83e-7	1.134e-8	3.591e-11	2.306e-12

TABLE I

A FEW VALUES OF $P(n \geq \lceil l/2 \rceil | \mathcal{C})$ GIVEN l AND α .

An elementary study of the right expression in equation 2 shows a convergence towards 0 as $l \rightarrow \infty$ whenever $\alpha < 1/2$. As a result, there always exists a value of T sufficiently large so that $l > T \Rightarrow P(n \geq l | \mathcal{C}) < \varepsilon$ holds for any $\varepsilon > 0$ holds. Moreover, if α is small, then the convergence of $P(n \geq l | \mathcal{C})$ to zero with l is exponential, namely

$$P(n \geq \lceil l/2 \rceil | \mathcal{C}) = O(n^{\lfloor \frac{l}{2} \rfloor})$$

Table I provides a few values of $P(n \geq \lceil l/2 \rceil | \mathcal{C})$ given l and α . The reader may observe that these values depend on the parity of l , which is no surprise given equation 2. Of course, the problem of estimating α remains, as we are not supposed to know it before runtime. However, we may see from Tab. I that setting $T \geq 25$ leads to a very acceptable test for all values of $\alpha \leq 0.4$.

Suppose now that \mathcal{C} is made of N primitives of integer length $l/N < T$. We denote by E_1 the event “ i and j belong to the same ground truth primitive”, by E_2 the event “ $\mathcal{C}[i..j]$ has more correct than incorrect points”, and by E_3 the event “Alg. 1 finds a primitive within n_{max} trials”. The **test** and **extract** functions cannot fail if both E_1 and E_2 occur, so we have

$$P(E_3 | \mathcal{C}) \geq 1 - (1 - P(E_1 | \mathcal{C})P(E_2 | \mathcal{C}))^{n_{max}}$$

which immediately leads to

$$n_{max} \leq \frac{\log(1 - P(E_3 | \mathcal{C}))}{\log(1 - P(E_1 | \mathcal{C})P(E_2 | \mathcal{C}))} \quad (3)$$

Obviously, $P(E_1) = 1/N$, and the expression of $P(E_2|\mathcal{C})$ given by Eq. 2 may be assumed to be close to 1. If we neglect it, we can rewrite Eq. 3 as

$$n_{max} \approx \frac{\log(1 - \sigma)}{\log(1 - 1/N)} \quad (4)$$

in which $\sigma = P(E_3|\mathcal{C})$ represents the desired probability that Alg. 1 succeeds after n_{max} trials under the hypothesis made on \mathcal{C} . In our system, we use Eq. 4 to determine n_{max} . The worst case corresponds to a systematic failure of Alg. 1 after n_{max} iterations, followed by the application of the chord algorithm. Each iteration of Alg. 1 runs in linear time with the number of pixels carried by the candidate fuzzy primitive, and finally, the complexity of the chord algorithm in the worst case is $O(l \log l)$. It therefore follows that the time complexity of the segmentation stage in the worst case is

$$O\left(nl \left[\log l + \frac{\log(1 - \sigma)}{\log(1 - 1/N)} \right] \right)$$

A noticeable fact is that because of the log terms, the above expression is exponential w.r.t both σ and N . This means that the method we suggest is not suitable for large values of N , and that we will always have to tolerate a few chances for failure if we want to avoid irrelevant computational times.

D. Optimization

The worst case for the simplification substage occurs when the ground truth is made n skeletal branches all cutting each other. In that case, the execution time of Alg. 3 is bounded by $O(n^2)$ (for the path search, the deletion of edges, but also the computation of the domains and their intersections).

In the unification stage, the equation of the potential, unifying primitive is known in constant time, but the fuzzy primitive test requires linear time w.r.t the number of pixels to test. The worst case occurs when all the primitives can be unified in a single one, with the additional constraint that only a single segment is integrated at a time (for example with n adjacent, colinear segments). In that case, the resulting complexity is $O(n^2l)$.

VI. EXPERIMENTAL RESULTS

We evaluate the capabilities of our method over the set of images used during the third, fourth and fifth IAPR contests on graphics recognition, held respectively at the GREC'99, GREC'01 and GREC'03 workshops [7]. Whereas a minimal tuning of the different systems' parameters was allowed to during the third contest, no preparation and human intervention was permitted at all during the fourth and the fifth. In order to make a fair comparison of our results with those obtained by the contestants, *we strictly followed the respective contests' policies during our experiments*. In particular, the default settings we used to obtain the results reported hereafter are the following:

- thickness evaluation: f automatically set to twice the highest evaluated thickness from the (3,4)-distance transform image;
- bounds for circular parameters: $(\theta_{min}, \rho_{min}, \rho_{max}) = (0.2rd, 15, \max(w, h))$, where w, h are the image's dimensions;
- other parameters: $m = 1$, $\sigma = 0.9999$ (Eq. 4), filtering turned on.

To the exception of the images from the first contest, these default parameters are *systematically applied to all images* without any change. Our method has been recently reimplemented as a 64-bit application running on PowerPC architecture. The results and times hereafter reported have been obtained on an Apple PowerMac dual-G5, 2.7GHz computer with 2.5GB memory. All the results, as well as an evaluation version of our method are available for download at <http://www.loria.fr/~hilaire/ranvec²>.

A. General performances

To give an overall idea of the performances our method can achieve, we first considered the set of images from the third contest. This set consists of 10 real, large images (typical resolution: 4800×4600 pixels) containing architectural maps and technical drawings; illustrations are available in Fig. 9. As a minimal tuning of parameters was allowed during this contest, we slightly modified the setup described previously in the following way: for each class of images (architectural maps, technical drawings), the user is asked for a crude estimation of ρ_{min} and ρ_{max} ; other parameters were left unchanged.

²If the paper is accepted, the evaluation data will be provided as complementary web resources for archival purposes.




Fig. 9. Two parts of images used during the third contest.

Image	Mica.RW	Mica.WD	Scan2CAD	Tractrix	Vecory	VrLiu	Our method
2032_02.tif	0.699	0.644	0.449	0.504	0.586	0.637	0.480
2212_a.tif	0.963	0.955	0.896	0.863	0.865	0.852	0.817
2212_c.tif	0.897	0.891	0.616	0.682	0.483	0.738	0.626
2212_e.tif	0.870	0.859	0.620	0.703	0.612	0.713	0.693
2226_02.tif	0.703	0.681	0.474	0.621	0.465	0.561	0.426
pal151.tif	0.951	0.947	0.771	0.779	0.870	0.881	0.726
pal169.tif	0.960	0.958	0.805	0.758	0.763	0.881	0.729
pal311.tif	0.981	0.978	0.925	0.912	0.933	0.943	0.902
pal343.tif	0.956	0.953	0.829	0.809	0.871	0.886	0.749
pal349.tif	0.944	0.944	0.803	0.811	0.766	0.894	0.686
Averages	0.892	0.881	0.719	0.744	0.721	0.799	0.683

TABLE II

EDIT COST INDEX (ECI) VALUES FOR THE PARTICIPATING SYSTEMS AT THE THIRD IAPR CONTEST ON GRAPHICS RECOGNITION AND OUR METHOD, OBTAINED WITH TEXT EVALUATION DISABLED. THE BEST RESULT FOR EACH IMAGE IS EMPHASIZED IN BOLD.

Results are presented in Table II, obtained with the official data of the contest. This table adheres to Phillips and Chhabra's evaluation methodology [32], and has been obtained thanks to their evaluation software set with default parameters, and text evaluation turned off. The processing time varies from 2mn,14s to 4mn,42s per image.

B. Arc detection capabilities

We evaluate our method's ability to extract circular arcs by following the rules of the fourth and fifth contests. The fourth contest contained three real, scanned images (p1.tif to p3.tif), and four synthetic images with Gaussian and high-frequency noise added on one side (g05.tif,f03.tif);

	syn_freq_03.tif	syn_gauss_05.tif	syn_geo_03.tif	syn_pen_03.tif	p1.tif	p2.tif	p3.tif
Elliman [14]	0.853	0.904	0.896	0.927	0.547	0.482	0.371
Average: 0.895				Average: 0.467			
Final score: 0.681							
Hilaire [18]	0.889	0.891	0.944	0.958	0.707	0.311	0
Average: 0.921				Average: 0.339			
Final score: 0.630							
Liu & Dori [50]	0.850	0.892	0.808	0.918	0.836	0.692	0.420
Average: 0.867				Average: 0.649			
Final score: 0.758							
Song et al. [38]	0.870	0.951	0.954	0.945	0.817	0.747	0.822
Average: 0.930				Average: 0.795			
Final score: 0.863							
Our VRI	0.879	0.935	0.935	0.914	0.845	0.644	0.708
Average: 0.916				Average: 0.732			
Final score: 0.824							

Image	Elliman	Song	Our method
1.tif	0.567	0.641	0.756
1_230.tif	0.589	0.640	0.752
1_n4.tif	0.664	0.509	0.762
2.tif	0.439	0.753	0.653
2_100.tif	0.513	0.786	0.707
2_n4.tif	0.612	0.703	0.791
3.tif	0.272	0.532	0.724
3_100.tif	0.519	0.301	0.740
3_n4.tif	0.451	0.224	0.697
4.tif	0.500	0.735	0.663
4_230.tif	0.323	0.790	0.795
4_n4.tif	0.399	0.688	0.782
Average	0.487	0.609	0.735

TABLE III

VRI SCORES OBTAINED FOR DIFFERENTS SYSTEMS ON THE IMAGES FROM THE FOURTH (LEFT) AND FIFTH (RIGHT) IAPR CONTEST ON GRAPHICS RECOGNITION.

hardpencil, and wrapping effects on the other side (h03.tif, w03.tif). Samples are given in Fig. 11.

Complete results are available on the left of Table III, in which we also include results from [38] for the sake of comparison. The reader, however, may observe two important points. First, the VRI scores presented in Table III are different from the official results of the contest; our explanation is the following: during the contest, the circular bounds $(\theta_{min}, \rho_{min}, \rho_{max})$ for *all images* were estimated from a test image given prior to the contest; and the content of this image was also very close to that of the four synthetic images. This not only explains the very good scores obtained by all contestants for the synthetic images, but also why the average performance drops when it comes to processing real images, as re-tuning parameters was not allowed. Second, the conditions in which results have been obtained in [38] are not stated by the authors. Regarding execution times, whereas synthetical images syn_*.tif were all of very similar complexity and could be processed at around 8s each, real images p1.tif, p2.tif and p3.tif were of unequal complexites and led to respective times 7s, 12s, and 1mn23s.

We applied again the same test procedure to the data of the fifth contest. The results are available at the right of Table III and in Fig. 10. Finally, Fig. 11 illustrates the behavior of our method w.r.t the four types of noise deccribed above. These results corroborate what has

been proven in [17] for gaussian and high-frequency noise: there is a relatively good resistance for these kinds of noise, whereas the scores collapse for the two other types, either due to the disconnection of components (and therefore fragmentation) for hard pencil, or to the m parameter set too tight to tolerate large displacements for the wrapping effects.




Fig. 10. Vectorization of some of the images from the fifth contest. From left to right, in each group: source image, vectorized arcs, vectorized lines; (a) image 1.tif; (b) image 2_n4.tif; (c) image 3_100.tif; (d) image 4_230.tif

C. Robustness

A critical point for any vectorization system is its ability to distinguish text and graphical parts, as they are usually processed independently. The text/graphics separation method mentioned in section III-B is able to process characters attached to graphical shapes, provided that for every text block, there exists a string seed of at least two characters not attached to any graphical shape.

However, our method is still able to eliminate text when the previous condition does not hold, provided that the length acceptance threshold f used in section IV-B is set at a sufficiently high

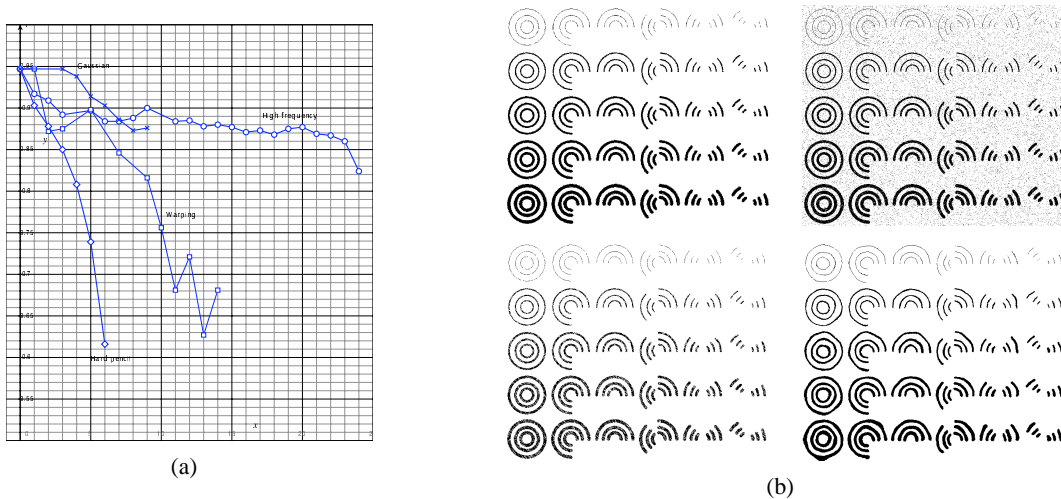


Fig. 11. (a) VRI scores [28] obtained with our method for various types of noise: abscissae: various noise levels as defined in [49]; ordinates: VRI scores; (b) sample test images with no noise, gaussian noise, hard pencil and wrapping effects added.

value. In short, we rely on the fact that a character, once skeletonized, can't be explained by Alg. 1 by any combination of graphical shapes with lengths all greater than f . Fig. 12-(a) gives a first example of such a situation: a sufficiently large value of f not only permits to eliminate the text, but also removes the dummy skeleton segment remained at the center of the X junction. A similar behavior may be observed on a complete image in Fig. 12-(b).

It is also noticeable that because of the criterions used during the simplification stage (Def. 6 for simplification, Eq. 1 for unification), it is impossible that the resulting vectorization lies outside the image. This prevents from precision and robustness issues; in particular, it ensures that a character remained attached to graphical part, once vectorized, creates aberrant results: at worst, its vectorization is not deleted, and remains as such in the output.

D. Other noticeable details

Figure 13 provides further details about the obtained results. Details (a) and (b) are two examples of correctly vectorized extracts of images from the third contest. Detail (c) reveals an interesting point: despite appearances, the lines in the staircase do not hit each other in a single point; indeed, either the staircase has been designed to fulfill this, or the drawing is simply not correct. At last, detail (d) shows an almost correct vectorization of the drawing: essential parts of the shapes have been properly extracted, but small circular arcs are sometimes recognized as




Fig. 12. Examples of text elimination. (a), from left to right; source image, direct vectorization (no text/graphics separation) with acceptance threshold $f = 4$, with $f = 15$. (b), from left to right: source image, result of text/graphics separation, result vectorization with $f = 20$.

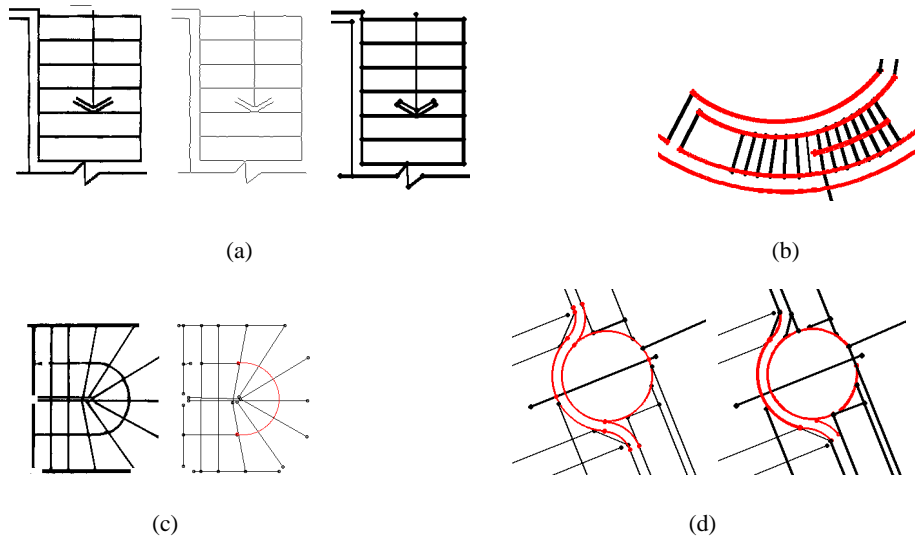


Fig. 13. A few noticeable elements fully vectorized with our vectorization method.

straight lines.

VII. CONCLUSION

In this paper, we have presented a robust and accurate vectorization method which does not include domain knowledge, and can thus be used on a variety of graphics-rich documents. The method is based on the random sampling paradigm and has the main following features:

- it segments the skeleton into the set of the most probable graphical primitives;
- it avoids fragmentation of the curve;
- it removes fake primitives without using ad-hoc rules;

- it yields excellent precision in the positioning of the junction points.

The method is presently limited to finding straight segments and circular arcs, although it could be extended to other primitives.

Obviously, the method has higher computational complexity than simple ad hoc rules applied to the skeleton, but it remains usable; in our opinion this is the price of the genericity and accuracy of the method.

Of course, the method still relies on a reasonably successful text-graphics segmentation and can be taken astray by characters or other noise touching the graphical parts; however, we have seen that under certain conditions, the method remains robust in the presence of some isolated text components. Being a non-contextual method, it also fails on some specific situations (such as X junctions with very small opening angles) for which simple contextual knowledge rules would probably lead to reliable and correct graphical interpretations.

Actually, we don't think that any form of contextual knowledge should be avoided, but we feel that it should be used in complement to accurate and robust generic methods. Therefore, the proposed method can very well be used as the basis for contextual vectorization methods, in which knowledge about the graphical configurations searched for might be added to the model.

More generally, we also believe that there is room for improvement in the method itself, by taking advantage of some kind of coarse-to-fine approach. The idea would be to use the present method to detect the most obvious primitives. These primitives would then be merged in a hierarchical way, starting with those having the highest confidence. The merging process can take advantage of locally available information and of contextual knowledge to perform some kind of "intelligent" reconstruction of the vector data.

Acknowledgments

The authors are grateful to the reviewers for valuable comments and suggestions on a first version of this manuscript. Special thanks also to Atul Chhabra, Ishin Phillips and Liu Wenyin, organizers of the third, fourth and fifth international contests on graphics recognition, for providing ground-truthed datasets and performance evaluation tools.

This work was to a large extent financed by a research collaboration with FS2i, 8 impasse de Toulouse, BP 141, 78001 Versailles Cedex, France, under a CIFRE scholarship.

REFERENCES

- [1] D. Antoine, S. Collin, and K. Tombre. Analysis of Technical Documents: The REDRAW System. In H. S. Baird, H. Bunke, and K. Yamamoto, editors, *Structured Document Image Analysis*, pages 385–402. Springer-Verlag, Berlin/Heidelberg, 1992.
- [2] H. Asada and M. Brady. The Curvature Primal Sketch. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):2–14, 1986.
- [3] H. Blum and R Nagel. Shape description using weighted symmetric axis features. *Pattern Recognition*, 10:167–180, 1978.
- [4] E. Bodansky and M. Pilouk. Using local deviations of vectorization to enhance the performance of raster-to-vector conversion systems. *International Journal on Document Analysis and Recognition*, 3(2):67–72, December 2000.
- [5] Y. Chen, N. A. Langrana, and A. K. Das. Perfecting Vectorized Mechanical Drawings. *Computer Vision and Image Understanding*, 63(2):273–286, March 1996.
- [6] A. K. Chhabra and I. T. Phillips. The Second International Graphics Recognition Contest—Raster to Vector Conversion: A Report. In Tombre and Chhabra [43], pages 390–410.
- [7] A. K. Chhabra and I. T. Phillips. Performance evaluation of line drawing recognition systems. In *Proceedings of the 15th International Conference on Pattern Recognition, Barcelona (Spain)*, volume 4, pages 4864–4869, 2000.
- [8] V. Poulain d’Andecy, J. Camillerapp, and I. Leplumey. Kalman Filtering for Segment Detection: Application to Music Scores Analysis. In *Proceedings of the 12th International Conference on Pattern Recognition, Jerusalem (Israel)*, volume 1, pages 301–305, 1994.
- [9] G. S. di Baja. Well-Shaped, Stable, and Reversible Skeletons from the (3,4)-Distance Transform. *Journal of Visual Communication and Image Representation*, 5(1):107–115, 1994.
- [10] D. Dori. Orthogonal Zig-Zag: an Algorithm for Vectorizing Engineering Drawings Compared with Hough Transform. *Advances in Engineering Software*, 28(1):11–24, 1997.
- [11] D. Dori and W. Liu. Sparse Pixel Vectorization: An Algorithm and Its Performance Evaluation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(3):202–215, March 1999.
- [12] L. Dorst and A. W. M. Smeulders. Discrete straight line segments: Parameters, primitives and properties. In R. Melter, P. Bhattacharya, and A. Rosenfeld, editors, *Vision Geometry, series Contemporary Mathematics*, volume 119, pages 45–62. American Mathematical Society, 1991.
- [13] Ph. Dosch, K. Tombre, C. Ah-Soon, and G. Masini. A complete system for analysis of architectural drawings. *International Journal on Document Analysis and Recognition*, 3(2):102–116, December 2000.
- [14] D. Elliman. TIF2VEC, An Algorithm for Arc Segmentation in Engineering Drawings. In D. Blostein and Y.-B. Kwon, editors, *Graphics Recognition – Algorithms and Applications*, volume 2390 of *Lecture Notes in Computer Science*, pages 350–358. Springer-Verlag, 2002.
- [15] M. A. Fischler and R. C. Bolles. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [16] L. A. Fletcher and R. Kasturi. A Robust Algorithm for Text String Separation from Mixed Text/Graphics Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(6):910–918, 1988.
- [17] X. Hilaire. *Segmentation robuste de courbes discrètes 2D et applications à la rétroconversion de documents techniques*. Thèse de doctorat, Institut National Polytechnique de Lorraine, 2004.
- [18] X. Hilaire and K. Tombre. Improving the Accuracy of Skeleton-Based Vectorization. In D. Blostein and Y.-B. Kwon, editors, *Graphics Recognition – Algorithms and Applications*, volume 2390 of *Lecture Notes in Computer Science*, pages 273–288. Springer-Verlag, 2002.

- [19] O. Hori and S. Tanigawa. Raster-to-Vector Conversion by Line Fitting Based on Contours and Skeletons. In *Proceedings of 2nd International Conference on Document Analysis and Recognition, Tsukuba (Japan)*, pages 353–358, 1993.
- [20] R. D. T. Janssen and A. M. Vossepoel. Adaptive Vectorization of Line Drawing Images. *Computer Vision and Image Understanding*, 65(1):38–56, January 1997.
- [21] S. H. Joseph. Unbiased Least Squares Fitting of Circular Arcs. *CVGIP: Graphical Models and Image Processing*, 56(5):424–432, September 1994.
- [22] S. H. Joseph and T. P. Pridmore. Knowledge-Directed Interpretation of Mechanical Engineering Drawings. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(9):928–940, September 1992.
- [23] T. Kanungo, R. M. Haralick, H. S. Baird, W. Stuezel, and D. Madigan. A Statistical, Nonparametric Methodology for Document Degradation Model Validation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(11):1209–1223, November 2000.
- [24] R. Kasturi, S. T. Bow, W. El-Masri, J. Shah, J. R. Gattiker, and U. B. Mokate. A System for Interpretation of Line Drawings. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(10):978–992, 1990.
- [25] P. Kultanen, E. Oja, and L. Xu. Randomized Hough Transform (RHT) in Engineering Drawing Vectorization System. In *Proceedings of IAPR Workshop on Machine Vision Applications, Tokyo (Japan)*, pages 173–176, 1990.
- [26] L. Lam, S.-W. Lee, and C. Y. Suen. Thinning Methodologies — A Comprehensive Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(9):869–885, September 1992.
- [27] K. H. Lee, Y. C. Choy, and S. B. Cho. Geometric Structure Analysis of Document Images: A Knowledge-Based Approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(12):1224–1240, 2000.
- [28] W.Y. Liu and D. Dori. A protocol for performance evaluation of line detection algorithms. *Machine Vision and Applications*, 9(5-6):240–250, 1997.
- [29] F. Mokhtarian and A. K. Mackworth. A theory of multiscale, curvature-based shape representation for planar curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(8):789–805, 1992.
- [30] V. Nagasamy and N. A. Langrana. Engineering Drawing Processing and Vectorization System. *Computer Vision, Graphics and Image Processing*, 49(3):379–397, 1990.
- [31] J. O’Rourke, S. R. Kosaraju, and N. Meggido. Computing circular separability. *Discrete and Computational Geometry*, 1:105–113, 1986.
- [32] I. T. Phillips and A. K. Chhabra. Empirical Performance Evaluation of Graphics Recognition Systems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(9):849–870, September 1999.
- [33] U. Ramer. An Iterative Procedure for the Polygonal Approximation of Plane Curves. *Computer Graphics and Image Processing*, 1:244–256, 1972.
- [34] P. L. Rosin. Techniques for Assessing Polygonal Approximation of Curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(6):659–666, June 1997.
- [35] P. L. Rosin and G. A. West. Segmentation of Edges into Lines and Arcs. *Image and Vision Computing*, 7(2):109–114, May 1989.
- [36] R. W. Smith. Computer Processing of Line Images: A Survey. *Pattern Recognition*, 20(1):7–15, 1987.
- [37] J. Song and M. R. Lyu. A hough transform based line recognition method utilizing both parameter space and image space. *Pattern Recognition*, 38(4):539–552, April 2005.
- [38] J. Song, M. R. Lyu, and S. Cai. Effective Multiresolution Arc Segmentation: Algorithms and Performance Evaluation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(11):1491–1506, November 2004.

- [39] J. Song, F. Su, C.-L. Tai, and S. Cai. An Object-Oriented Progressive-Simplification Based Vectorization System for Engineering Drawings: Model, Algorithm, and Performance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(8):1048–1060, August 2002.
- [40] C.-H. Teh and R. T. Chin. On the Detection of Dominant Points on Digital Curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(8):859–872, 1989.
- [41] E. Thiel. *Les distances de chanfrein en analyse d’image : fondements et applications*. Thèse de doctorat, Université Joseph Fourier - Grenoble I, 1994.
- [42] K. Tombre, C. Ah-Soon, Ph. Dosch, G. Masini, and S. Tabbone. Stable and Robust Vectorization: How to Make the Right Choices. In A. K. Chhabra and D. Dori, editors, *Graphics Recognition—Recent Advances*, volume 1941 of *Lecture Notes in Computer Science*, pages 3–18. Springer-Verlag, September 2000.
- [43] K. Tombre and A. K. Chhabra, editors. *Graphics Recognition—Algorithms and Systems*, volume 1389 of *Lecture Notes in Computer Science*. Springer-Verlag, April 1998.
- [44] K. Tombre, S. Tabbone, L. Pélissier, B. Lamiroy, and Ph. Dosch. Text/Graphics Separation Revisited. In D. Lopresti, J. Hu, and R. Kashi, editors, *Proceedings of 5th IAPR International Workshop on Document Analysis Systems, Princeton (NJ, USA)*, volume 2423 of *Lecture Notes in Computer Science*, pages 200–211. Springer-Verlag, August 2002.
- [45] Peter Veelaert. Concurrency of line segments in uncertain geometry. In *Proceedings of 10th International Conference on Discrete Geometry for Computer Imagery, Bordeaux (France)*, volume 2301 of *Lecture Notes in Computer Science*, pages 289–300. Springer-Verlag, 2002.
- [46] K. Wall and P. Danielsson. A Fast Sequential Method for Polygonal Approximation of Digitized Curves. *Computer Vision, Graphics and Image Processing*, 28:220–227, 1984.
- [47] L. Wenyin and D. Dori. A Survey of Non-Thinning Based Vectorization Methods. In A. Amin, D. Dori, P. Pudil, and H. Freeman, editors, *Advances in Pattern Recognition (Proceedings of Joint IAPR Workshops SSPR’98 and SPR’98, Sydney, Australia)*, volume 1451 of *Lecture Notes in Computer Science*, pages 230–241, August 1998.
- [48] L. Wenyin and D. Dori. Incremental Arc Segmentation Algorithm and its Evaluation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(4):424–431, April 1998.
- [49] L. Wenyin, J. Zhai, D. Dori, and T. Long. A system for performance evaluation of arc segmentation algorithms. In *CVPR Workshop on Empirical Evaluation in Computer Vision*, December 2001.
- [50] Liu Wenyin and Dov Dori. Genericity in Graphics Recognition Algorithms. In Tombre and Chhabra [43], pages 9–20.
- [51] H. Yamada. Paper-Based Map Processing. chapter 19, pages 503–528. World Scientific, 1997.
- [52] Y. Zheng, H. Li, and D. Doermann. A Parallel-Line Detection Algorithm Based on HMM Decoding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(5):777–792, May 2005.