

# Optimal Line and Arc Detection on Run-Length Representations

Daniel Keysers and Thomas M. Breuel

Image Understanding and Pattern Recognition Research Group  
German Research Center for Artificial Intelligence (DFKI GmbH)  
and University of Kaiserslautern  
D-67663 Kaiserslautern, Germany  
E-mail: {daniel.keysers, thomas.breuel}@dfki.de

**Abstract.** The robust detection of lines and arcs in scanned documents or technical drawings is an important problem in document image understanding. We present a new solution to this problem that works directly on run-length encoded data. The method finds globally optimal solutions to parameterized thick line and arc models. Line thickness is part of the model and directly used during the matching process. Unlike previous approaches, it does not require any thinning or other preprocessing steps, no computation of the line adjacency graphs, and no heuristics. Furthermore, the only search-related parameter that needs to be specified is the desired numerical accuracy of the solution. The method is based on a branch-and-bound approach for the globally optimal detection of these geometric primitives using runs of black pixels in a bi-level image. We present qualitative and quantitative results of the algorithm on images used in the 2003 and 2005 GREC arc segmentation contests.

*Keywords:* Graphics Recognition, Line Drawings, Technical Drawings, Branch-and-Bound algorithms.

## 1 Introduction

A large amount of technical and engineering information is still available only in paper form. The goal of graphics recognition is to make that information available in electronic format. After initial scanning, a key problem in the conversion of drawings into electronic form is the reliable detection of the geometric primitives from which those drawings are constructed [1].

This problem has been the subject of intensive research, and it has even been the subject of several graphics recognition contests (held in conjunction with the GREC workshops; please refer to [2,3] and the references therein for an overview; a report of the 2005 contest is also included in this proceedings volume).

A variety of approaches to line detection have been described in the literature, many of them based on the Hough transform [4]. A survey of various methods is given in [5]. Many approaches also first perform a vectorization, i.e. they determine a representation by poly-lines and then use this intermediate step for

line or arc detection, e.g. [6,7]. One method along this line that performed very well in the 2005 GREC arc segmentation contest is an extension of the method described in [8].

While these existing approaches have found numerous practical applications, they have some difficulties and limitations. Hough transforms, for example, require careful choice of parameters like bucket size in order to reduce the risk of losing solutions. Preprocessing steps like vectorization and polygonalization, are themselves complex processing steps, and, since they generally only represent approximations to the original bitmap, also potentially reduce the accuracy and reliability of the overall system.

This paper introduces an algorithm for line and circle detection that operates directly on a pixel-accurate representation of the original binary image. This avoids the complexity and potential for errors introduced by separate preprocessing steps. Furthermore, it yields optimal solutions for a broad class of well-defined geometric and statistical models of thick lines.

The basis of the algorithm is the RAST algorithm (Recognition by Adaptive Subdivisions of Transformation Space [9]), which is guaranteed to find globally optimal solutions under given error models and requires no parameters other than desired numerical accuracy to be specified by the user. RAST algorithms are closely related to (hierarchical) Hough transforms but have more desirable combinatorial properties for object recognition. Recently, effective and simple to implement variants of RAST algorithms that are based on interval arithmetic have been introduced [10]. The RAST algorithm avoids the drawback of fixed parameter discretization by subdividing regions of the parameter space adaptively only when it is possible that a desired solution exists within that region of the parameter space. The closely related adaptive Hough transform uses the same strategy, but cannot guarantee to find the global optimum because it does not keep track of intermediate scores and thus cannot recover from a decision to drop a portion of the parameter space from the search.

The work described in this paper is directly based on the principles presented in [10,11], where these principles have been applied to settings in which not pixels but the outlines of connected components were used. The application of branch-and-bound algorithms to matching problems has been described in the computer vision literature before [9,12,13,14,15,16].

In our approach, we use run length coding as the representation of the image content. Run-length coding for images is commonly used for in-memory representation of large binary images, thereby making it easy to apply the algorithm to potentially large images in memory. Run-length coding is also familiar as an important intermediate step in some binary image compression methods.

Run-length coding is not only a convenient representation in terms of memory, it also serves a semantic function: in a binary image, each line becomes a collection of adjacent runs in the run-length encoded representation. Previous approaches to line detection in images have attempted to take advantage of this representation by considering the line adjacency graph and, for example, applying Hough transform methods to the midpoints of each run [17]. We use

run-length coding in a slightly different way. Rather than using the midpoints of each run, we match the endpoints of each run against the two boundaries of a thick line model.

The presented approach separates the search algorithm from the quality function used. This modularity makes it easy to also incorporate and test other quality functions or use other representations of the image content. For example, the algorithm can also be used on a quality function using segments of a polygonal approximation of the edges present in the image.

In the rest of the paper, we first describe the framework for optimal matching and its use of interval arithmetic, followed by a description of the specific quality functions used for the detection of the thick lines and arcs. We then present results on example images of scanned technical drawings, briefly discuss the results of the GREC 2005 arc segmentation contest, and end with concluding remarks.

## 2 Detection Framework

Assume we are given a set  $R = r_1, \dots, r_N$  of pixel runs in the image plane, where each run  $r = (x_0, x_1, y)$  is defined by its starting point  $(x_0, y)$  and its end point  $(x_1, y)$ . The computation of this set from a given image is straight-forward. Note that the decision to choose runs in horizontal direction is arbitrary and only based on the most common pixel order found in the representation of digital images. The algorithms presented here would work equally well with runs of pixels in a diagonal direction.

Now we are interested in finding the best matching geometric primitives with respect to the set  $R$ . The geometric primitives are characterized by a set of parameters  $\vartheta \in T$ , e.g. center, radius, and thickness for a circle (cp. Section 4). Note that we avoid adding additional parameters for the start and end points of lines and arcs here, because doing so would enlarge the search space dimensionality. Instead, we evaluate all runs associated with a primitive after the search and then determine the parameters of the line or arc based on these.

We perform the detection by finding the maximizing set of parameters

$$\hat{\vartheta}(R) := \arg \max_{\vartheta \in T} Q(\vartheta, R) \quad (1)$$

where the total quality  $Q(\vartheta, R)$  of a parameter set is defined as the sum of local qualities

$$Q(\vartheta, R) := \sum_{n=1}^N q(\vartheta, r_n) \quad (2)$$

where  $q(\vartheta, r)$  is a local quality function that evaluates the goodness of fit for a given run of pixels  $r$  and a set of parameters  $\vartheta$ .

This maximization will be a complex task for most functional forms of  $Q$ . In many applications, such fits of parameters are carried out iteratively and

heuristically, which involves the risk that the results found are only locally optimal solutions. Other methods include randomized approaches like e.g. random sample consensus [18].

We propose to employ a branch-and-bound technique [9,16,19] to perform this task that guarantees to find the globally optimal parameter set by recursively subdividing the parameter space and processing the resulting parameter hyper-rectangles in the order given by the upper bound on the total quality. Moreover, the algorithm allows us to efficiently determine the  $k$  best matches, not only the best match.

The computation of an upper bound on the quality of any geometric primitive with parameters in a hyper-rectangular region  $T$  is given by

$$\max_{\vartheta \in T} Q(\vartheta, R) \leq \sum_{n=1}^N \max_{\vartheta \in T} q(\vartheta, r_n), \quad (3)$$

where  $\max_{\vartheta \in T} q(\vartheta, r_n)$  is straightforward to compute. Note that the computation of the upper bound is implicitly handled using interval arithmetic as described in Section 3. We can now organize the search as follows (for details see [9,19]):

1. Pick an initial region of parameter values  $T$  containing all the parameters that we are interested in.<sup>1</sup>
2. Maintain a priority queue of regions  $T_i$ , where we use as the priority the upper bound on the possible values of the global quality function  $Q$  for parameters  $\vartheta \in T_i$ .
3. Remove a region  $T_i$  from the priority queue; if the upper bound of the quality function associated with the region is too small to be of interest, terminate.<sup>2</sup>
4. If the region is small enough to satisfy our accuracy requirements, accept it as a solution; remove all runs involved in that solution from further consideration and continue at Step 3.
5. Otherwise, split the region  $T_i$  along the dimension furthest from satisfying our accuracy constraints and insert the subregions into the priority queue; continue at Step 3.

This algorithm will return all maximum quality matches greedily in decreasing order of total quality. Note that the optimality is only guaranteed for the best match, then for the second match excluding the already matched components, and so on. We cannot guarantee optimality for the combination of all returned matches, which is a problem of much higher complexity. To make this sequential approach practical and avoid duplicate computations, we use a matchlist

<sup>1</sup> The initial region can always be picked large enough to include all possible arcs and lines within a given image and thus it is not an additional parameter of the algorithm. If we decide to choose a smaller initial region, this is because we want to exclude certain solutions.

<sup>2</sup> The choice of the parameter for termination does not influence the search itself. It is rather a convenience parameter that lets the algorithm terminate before all runs have been processed, if it is known that we require a certain minimum quality.

representation [9]. That is, with each region kept in the priority queue in the algorithm, we maintain a list (the matchlist) of all and only those runs that have the possibility to contribute with a positive local quality to the global quality. These matchlists will be large for the initially large regions of the parameter space but shrink with decreasing size of the regions  $T_i$ . The justification for the use of matchlists is that runs for which the distance from the geometric primitive is surely greater than half the maximum line thickness  $d$  that is contained within  $T_i$  do not contribute to the computation of the quality function at all. Furthermore, it is easy to see that the upper bound of a parameter space region  $T_i$  is also an upper bound for all subsets of  $T_i$ . When we split a region in Step 5, we therefore never have to reconsider runs in the children that have already failed to contribute to the quality computation in the parent.

The use of matchlists enables us to add another feature to the algorithm: Before the execution of Step 5 above, we introduce another conditional processing of the region  $T_i$ . If the geometric primitive described by  $T_i$  is associated with a matchlist of runs that can be split into two parts that are sufficiently separated (say, by 10 pixels), split the matchlist into these two parts and re-insert the region  $T_i$  into the queue twice, once with each part of the matchlist associated. Doing so avoids the unification of primitives that belong to separated regions in the image. It furthermore separates noise components from the matchlists, which then can be pruned in the further search.

After running the algorithm, the bounding boxes for the lines are determined by considering all associated runs in the matchlists and for circular arcs the smallest arc that encompasses all runs in the matchlist.

### 3 Use of Interval Arithmetic

One of the key points in the RAST algorithm is the calculation of the upper bound of the quality function for a subset of the parameter set. This upper bound can be determined using interval arithmetic both efficiently from a computational point of view and from the point of view of coding the actual function.

Interval arithmetic was originally proposed as a means of bounding the error of numerical computations. For detailed information about interval arithmetic and its implementation, the reader is referred to the references [20,21].

The extension of the basic numerical operators to intervals is straight-forward. Instead of a rigorous definition, let us give two examples of such extensions to illustrate the principle. The operators  $+$  and  $\times$  are defined on intervals in the following way:

$$\begin{aligned}
 [x_0, x_1] + [y_0, y_1] &:= [x_0 + y_0, x_1 + y_1] \\
 [x_0, x_1] \times [y_0, y_1] &:= [\min(x_0y_0, x_0y_1, x_1y_0, x_1y_1), \max(x_0y_0, x_0y_1, x_1y_0, x_1y_1)]
 \end{aligned}$$

Thus, any addition of numbers from  $[x_0, x_1]$  and  $[y_0, y_1]$  will result in a number from  $[x_0, x_1] + [y_0, y_1]$ , and analogously for the multiplication operation. Due to monotonicity the extensions for  $\log()$  and  $\exp()$  are simple and for other

functions like e.g.  $\sin()$  a slightly more elaborate analysis leads to the appropriate formulation.

If the quality function used in the RAST algorithm is a composition of simpler functions like multiplication, addition, sine, minimum, square, etc., the function can be directly implemented using the basic interval arithmetic functions. In other cases, it might be more straight-forward to base parts of the determination of the upper bound on the interval arithmetic, like e.g. the calculation of the distance of a pixel from a line, and separately code other parts.

The use of interval arithmetic instead of an explicit derivation of bounding functions, as used in prior work on branch-and-bound matching algorithms, has several advantages [10]: In many geometric matching problems of practical interest, deriving bounding functions and implementing them correctly can be difficult and error prone. This problem is substantially alleviated by using interval arithmetic, which allows us to simply replace calculations on scalars by the corresponding calculations on intervals in many cases. Furthermore, the use of interval arithmetic also makes geometric matching numerically reliable, i.e. we can guarantee that no solutions are lost due to roundoff errors. This last statement is based on the use of interval functions that guarantee to include the real-valued result within an interval bounded by finite-precision floating-point numbers as representable on a computer [20].

## 4 Quality Functions for the Detection of Lines and Circular Arcs

For the detection of the geometric primitives considered in this work we use the following parameterizations. The line model is parameterized by its angle  $\varphi$ , its distance  $s$  from the origin, and its thickness  $d$ . The circle (and circular arc) model is parameterized by the coordinate of its center point  $(x, y)$ , its radius  $s$  and its thickness  $d$ .

During the search algorithm, each partition of the search space is described by a Cartesian product of intervals for the parameters, i.e. for a line by a set of the form  $T = [\varphi_0, \varphi_1] \times [s_0, s_1] \times [d_0, d_1]$  and for a circle by a set of the form  $T = [x_0, x_1] \times [y_0, y_1] \times [s_0, s_1] \times [d_0, d_1]$ .

Given a run  $r = (x_0, x_1, y)$  we consider the starting point of the run  $(x_0, y)$  and the end point  $(x_1, y)$  and determine the interval of possible distances from the center of the thick line or circle, respectively. We then compute an upper and lower bound for  $q$  in the straight-forward way. The total upper bound of the quality  $Q$  is then used in the branch-and-bound algorithm.

The local quality function for a run and a geometric primitive (line or circle) is then defined to be

$$q(\vartheta, (x_0, x_1, y)) = \max\left(0, 1 - \frac{\left(\frac{d}{2} - d_\vartheta(x_0, y)\right)^2}{\sigma^2}\right) * \max\left(0, 1 - \frac{\left(\frac{d}{2} - d_\vartheta(x_1, y)\right)^2}{\sigma^2}\right) \quad (4)$$

where  $d_\vartheta(x, y)$  denotes the (absolute) distance of the point  $(x, y)$  from the center-line of the thick geometric primitive defined by the parameters  $\vartheta$ . This quality

function reaches its maximum value of one if the endpoints of the run coincide with the predicted endpoints given the parameters of the primitive. Deviations of the positions from their predicted values are tolerated up to a value of  $\sigma$  with decreasing quality values. The parameter  $\sigma$  was set to 0.5 pixels in the experiments.

It is of course possible that we may find quality functions that are more suitable for the task discussed here and one can think of a variety of other functions measuring the goodness of match between a run and a line with finite thickness. For example, the relative overlap of runs and lines is better suited for the detection of horizontal lines based on runs than using the end points of the runs. However, this drawback can be avoided by running the algorithm once for lines  $\pm\pi/4$  from the vertical axis, then rotating the remaining image elements by  $\pi/2$  and repeating the detection. We experimented with a few other possibilities for the quality function like e.g. the relative overlap of a run  $(x_0, x_1, y)$  and the expected  $y$ -run of the line. Among these, the informal experiments suggested that the local quality function (4) produces the best results.

## 5 Results

We present experimental results for the images of the GREC arc segmentation contests that took place in the years 2003 and 2005.

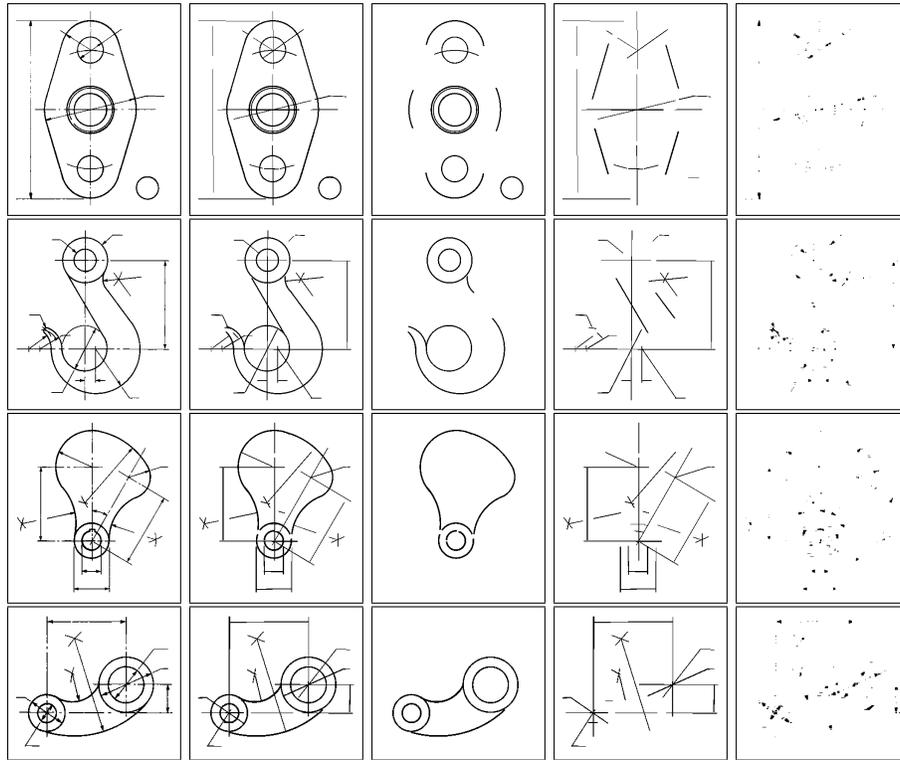
### 5.1 Images from the 2003 Arc Segmentation Contest

We applied the algorithm to unprocessed images that were used in the fifth contest on arc segmentation at GREC 2003 to show the applicability and the quality of the results. These images show scanned technical drawings.

Each of the rows in Figure 1 shows (in that order) the original image, the union of all detected primitives, the detected circular arcs, the detected lines on the remaining runs not covered by arcs, and the remaining runs in the image that were not explained by any arc or line.

In some cases errors occur where geometric primitives extend over a larger range than visually appropriate. In that case e.g. arcs extend too far and thus prevent the detection of complete lines. It is also interesting that the remaining runs in the image often correspond to arrow-heads as these are not included in our model of the geometric primitives.

In the 2003 contest, in total 12 images were supplied for the test runs. Each of the four images shown in Figure 1 was additionally used in two version with artificially added noise. The participants in the 2003 contest obtained average VRI scores of 0.609 and 0.487, respectively [2]. Using the proposed method and the evaluation software corresponding to the 2003 contest, we obtained an average VRI score of 0.757, i.e. better than the two participating methods in 2003. Note, however, that the winning method of 2005 did not participate in 2003. Note also that although we did not specifically tune the parameters to these data, the comparison of the scores is not entirely fair because we had access to the images before actually running the evaluation.



**Fig. 1.** Results for the detection of arcs and lines using the proposed method. Each row shows (left to right) input, detected arcs and lines overlaid, detected arcs, detected lines, remaining runs not explained.

## 5.2 The 2005 Arc Segmentation Contest

We participated with the presented method in the GREC 2005 arc segmentation contest. As detailed in the contest overview, which the reader can find in this proceedings volume, the method obtained the second place among the three participating systems with an average VRI score of 0.615. The scores of the competing systems were 0.801 for the winning method and 0.514 for the method that obtained the third place.

The contest used six images, again each in an original version and with two different types of noise added. Figure 2 shows each of the original images along with the detection result of our method and the ground truth arcs. The results for the noisy images are very similar<sup>3</sup>. The only preprocessing that was performed was the application of the `pbmclean` program from the NetPBM library, which removes single pixels that are not part of any connected component in the image.

<sup>3</sup> The complete set of results can be found at <http://www.iupr.org/~keyzers/grec2005/grec2005.html>.

Looking at the results, we can observe that the most crucial problem for the algorithm is the following: When two or more image objects touch or overlap, the pixel runs do not coincide with the expected pixel runs of the geometric primitive any more. In the present formulation of the quality function, this is not considered, which leads to the majority of the observed errors.

Note that we did not tune our method to the evaluation criterion used in the contest, which may be possible. For example, Hilaire [8] tuned the method specifically to the performance measure in the 2001 contest as he writes in the discussion of the method that “the penalty associated with *fragmentation* is probably the worst we may encounter” and that “our method should *first worry about thick, long arcs or full circles*”, which are consequences of the evaluation scheme used. Note further that it is possible to obtain improvements in the score if one also outputs line segments, which our method did not do. It is possible that the the GREC 2007 arc segmentation contest will use a different evaluation scheme which would then be similar to the one discussed in [22].

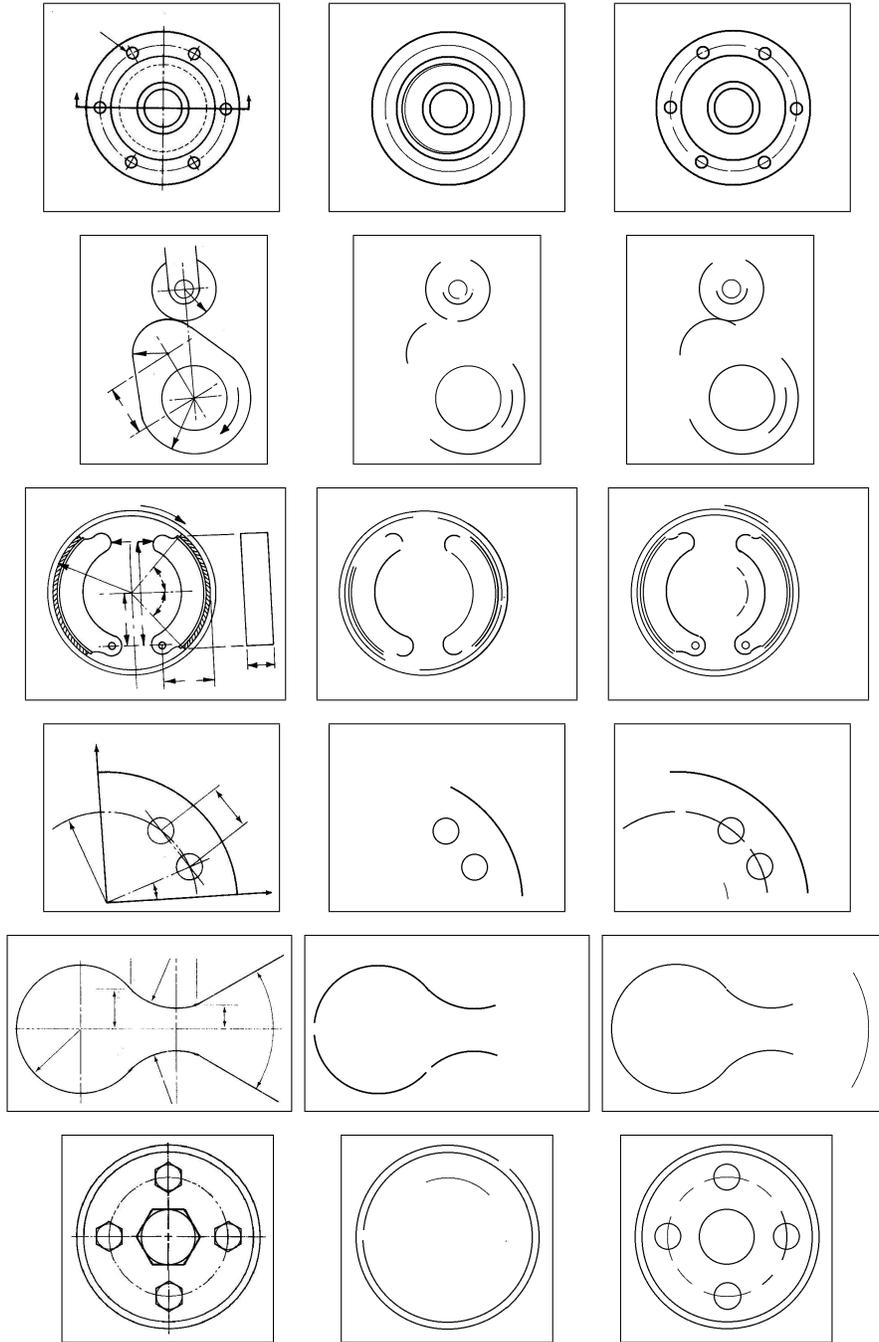
## 6 Conclusion

We have described a simple and effective method for locating geometric primitives like lines and circular arcs in bi-level images of scanned documents. Except for a library for interval arithmetic and image file I/O, the algorithm itself is easily implemented in a few hundred lines of C++ code. Another property of the algorithm that simplifies its implementation and use is that the only parameters that the user needs to specify are bounds for the thickness of the lines, the desired accuracy of the primitives, and the minimum quality desired. The experimental results show a promising behavior of the algorithm on images of technical drawings.

Note that the development of the described algorithm is work in progress, although it relies on two well-understood and tested methods, which are the RAST framework of branch-and-bound search and the use of interval arithmetic for the computation of bounds of functions defined on regions of parameter space.

To our knowledge, the algorithm presented in this paper is the first practical, globally optimal algorithm for the detection of thick lines and thick circles under well defined error models in technical drawings; other methods described in the literature generally rely on suboptimal matching methods (like the Hough transform) and approximations during preprocessing (like polygonal approximations).

Existing systems for the detection of geometric primitives often include complex, hand-tuned rules related to grouping, continuity, and plausible parameter ranges. If we identify any such cases or constraints, we can incorporate them into the quality function used by our algorithm. The optimality guarantees of the algorithm still hold (i.e., the algorithm will find the optimal results under the improved quality function), and so do the main advantages of the algorithm (few parameters, strict separation of what is being optimized from the search algorithm itself).



**Fig. 2.** GREC 2005 arc segmentation contest images (without noise), detected arcs and ground truth arcs.

## Acknowledgments

This work was partially funded by the BMBF (German Federal Ministry of Education and Research), project IPeT (01 IW D03).

## References

1. Elliman, D.: Tif2vec, an algorithm for arc segmentation in engineering drawings. In: Proc. 4th IAPR Workshop on Graphics Recognition, Springer LNCS 2390. (2001) 359–358
2. Wenyin, L.: Report of the arc segmentation contest. In: Proc. 5th IAPR Workshop on Graphics Recognition, Springer LNCS 3088, Barcelona, Spain (2003) 364–367
3. Wenyin, L., Zhai, J., Dori, D.: Extended summary of the arc segmentation contest. In: Proc. 4th IAPR Workshop on Graphics Recognition, Springer LNCS 2390. (2001) 343–349
4. Duda, R.O., Hart, P.E.: Use of the Hough transformation to detect lines and curves in pictures. *Communications of the ACM* **15** (1972) 11–15
5. Wenyin, L., Dori, D.: From raster to vectors: Extracting visual information from line drawings. *Pattern Analysis and Applications* **2** (1999) 10–21
6. Dori, D., Wenyin, L.: Sparse pixel vectorization: An algorithm and its performance evaluation. *IEEE Trans. Pattern Analysis Machine Intelligence* **21** (1999) 202–215
7. Dosch, P., Masini, G., Tombre, K.: Improving arc detection in graphics recognition. In: International Conference on Pattern Recognition. Volume 2., Barcelona, Spain (2000) 2243–2246
8. Hilaire, X.: RANVEC and the arc segmentation contest. In Blostein, D., Kwon, Y.B., eds.: *Graphics recognition—Algorithms and applications*. Volume 2390 of *Lecture Notes in Computer Science*. Springer-Verlag (2002) 359–364
9. Breuel, T.M.: Fast recognition using adaptive subdivisions of transformation space. In: *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition*. (1992) 445–451
10. Breuel, T.M.: On the use of interval arithmetic in geometric branch-and-bound algorithms. *Pattern Recognition Letters* **24** (2003) 1375–1384
11. Breuel, T.M.: Robust least square baseline finding using a branch and bound algorithm. In: *Proceedings of the SPIE - The International Society for Optical Engineering*. (2002) 20–27
12. Hagedoorn, M., Veltkamp, R.C.: Reliable and efficient pattern matching using an affine invariant metric. *International Journal of Computer Vision* **31** (1999) 203–225
13. Huttenlocher, D.P., Klanderman, G.A., Rucklidge, W.J.: Comparing images using the Hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **15** (1993) 850–63
14. Jurie, F.: Solution of the simultaneous pose and correspondence problem using Gaussian error model. *Computer Vision and Image Understanding* **73** (1999) 357–373
15. Mount, D.M., Netanyahu, N.S., Moigne, J.L.: Efficient algorithms for robust feature matching. *Pattern Recognition* **32** (1999) 17–38
16. Olson, C.F.: Locating geometric primitives by pruning the parameter space. *Pattern Recognition* **34** (2001) 1247–1256

17. Pavlidis, T.: Algorithms for Graphics and Image Processing. W. H. Freeman & Co., New York, NY (1983)
18. Forsyth, D.A., Ponce, J.: Computer Vision: A Modern Approach. Prentice Hall, Upper Saddle River, NJ (2003)
19. Breuel, T.M.: Finding lines under bounded error. *Pattern Recognition* **29** (1996) 167–178
20. Hickey, T.J., Ju, Q., van Emden, M.H.: Interval arithmetic: From principles to implementation. *JACM* **48** (2001) 1038–1068
21. Jaulin, L., Kieffer, M., Didrit, O., Walter, E.: Applied Interval Analysis. Springer Verlag, Berlin (2001)
22. Breuel, T.M.: Representations and metrics for off-line handwriting segmentation. In: 8th International Workshop on Frontiers in Handwriting Recognition. (2002)