A Simple and Efficient Skew Detection Algorithm via Text Row Algorithm

Ray Smith
Personal Systems Laboratory
HP Laboratories Bristol
HPL-94-113
December, 1994

document image processing,
skew detection

An important part of any document recognition system is detection of skew in the image of a page. This paper presents a new, accurate and robust skew detection algorithm based on a method for finding rows of text in page images. Results of a test of the new algorithm and a comparison against Baird's well known algorithm on 400 pages show the new algorithm to be more accurate, robust and somewhat faster. In particular the new algorithm only breaks down at skew angles in excess of 15 degrees, compared to the almost uniform distribution of breakdowns of Baird's algorithm.

Internal Accession Date Only

A simple and efficient skew detection algorithm via text row accumulation

1 Introduction

An important part of any document recognition system is detection and correction of skew in the image of a page. Page layout analysis and preprocessing operations used for character recognition depend on an upright image or, at least, knowledge of the angle of skew.[1][3][5][9] One example of a process which is spoilt by skew is the use of horizontal and vertical projection profiles. Projection profiles have many applications in document image processing and they rely on horizontal and vertical lines being aligned to the axes.

This paper presents a new, accurate skew detection algorithm based on a simple and robust method for finding rows of text independently of the skew angle of the image. After finding the rows of text, it is possible to obtain an accurate estimate of the skew angle of each text line and thereby estimate the skew angle of the whole page to a high degree of accuracy.

2 Previous work

A common method of skew detection is to simplify the Hough transform[4] by applying it to only a subset of pixels in the image.

The basic concept behind methods based on the Hough transform is the same:

1. Select some subset of pixels of the image which are few in number and most likely to form straight lines parallel to the baselines of the text rows.

2. For as many different directions as is necessary to achieve the desired accuracy, project the points selected in step 1 parallel to each direction in turn.

3. The direction for which the highest spikes are achieved in the projection is the direction of the page skew.

Baird[2] and Nakano et. al.[5] use very similar methods by projecting the bottom of the bounding box of connected components parallel to a set of angles and choose the angle which maximizes a measure of the variance of the counts in the projection.

Hinds et. al.[3] use vertical run-length coding and project the bottoms of the runs and Spitz[9] achieves a similar effect using Group 4 fax coding. Both of these methods avoid doing the connected component analysis required by Baird.

The basic idea behind the algorithm described in this paper was first outlined by Pratt et. al.[6], but has been improved significantly to cope with a much wider range of types of page.

3 Finding text lines independently of skew

The new algorithm was initially developed to solve the problem of finding the rows of text on a page in the presence of broken and joined characters, speckle noise and page skew, even in excess of 20 degrees.

The algorithm operates as follows:

- Connected component analysis.

  Perform a connected component analysis[7] of the image. The connected components are (most likely) needed for the recognition processes anyway. For brevity, the connected components will be referred to as blobs. All future references to dimensions and location of a blob relate to the coordinates of its upright bounding box.

- Filter blobs.

  The purpose of this process is to select a subset of the blobs which have a good chance of representing the body text. Exactness is unimportant – the key is in filtering out drop-caps, underlines and isolated noise. The blobs left out at this stage will be put in their place in the text lines later.

  Blobs smaller than a fixed number of pixels in height are removed. The remaining blobs are then filtered by size. Blobs with height between the 20th and 95th percentile are retained, provided their width is not excessive.

- Sort blobs.

  Sort the blobs (into ascending order) using x-coordinate (of the left edge) as the sort key. This sort makes it possible to track the skew across the page.

- Make initial rows.

  Set the running average y shift to zero.
  For each blob in sorted order:
  >   Find the existing row which has most vertical overlap with the blob
  >   If there is no overlapping row
  >   Then
  >>   Make a new row and put the blob in it.
  >>   Record the top and bottom coordinates of the blob as the
  >>   top and bottom of the row.
  >   Else
  >>   Add the blob to the row.
  >>   Expand the top and bottom limits of the row with the top and
  >>   bottom of the blob, clipping the row height to a limit.
  >>   Update the running average y shift with the bottom of the blob.
  >   Endif
  Endfor

  The sort ensures that blobs are processed in an order which gives some degree of overlap between adjacent blobs on a line. The sort also enables the use of a running average to measure the gradual vertical shift across the page. The running average y shift is used to vertically shift blobs when measuring their overlap with the rows. It is updated for each blob using:

  $$Shift_{n+1} = \alpha Shift_n + (1 - \alpha) Newshift$$

  where $0.5 < \alpha < 0.7$ and increasing to the larger end with an increasing number of rows on the page. The running average ensures that it is possible to track even large skew angles while being immune to descenders.

Where there is a significant horizontal gap in the available blobs, the y shift is extrapolated across the gap.

As the rows are accumulated, the top and bottom limits of the row are expanded to fit the inverse shifted y limits of the new blob. This helps the rows collect all the blobs which belong to them, rather than starting an extra row because the first blob on a line was small. The size of each row is restricted to a limit proportional to the upper quartile of heights from the filtering process. This prevents a row from expanding due to the different y positions of blobs on a skewed page, and growing so much that it collects blobs from more than one row.

- Fit baselines.

  At the end of the above process, most of the blobs on the page are associated with some row. The arrangement is already good enough to obtain an accurate baseline for most of the rows on the page. This is done using a least median of squares fit.[8] The least median of squares fit is ideal for this application, as the baselines are close to straight, and there are outliers, (descenders, punctuation and other special characters) which would cause a least mean squares fit to exhibit significant error.

The process of finding the global page skew stops here. The median gradient of the baselines provides a highly accurate estimate for the page.

4 Testing and results

The algorithm was tested by running it on a database of 400 A-size and A4 page images scanned at 300 dpi. These images are a mixture of real office documents varying in quality from original business letters and magazine pages to badly degraded photocopies and faxes. A significant proportion of the pages contain tables, pictures or figures and it is important that the skew detection is not distracted by the presence of these components.

The algorithm was tested against that of Baird,[2] since it is one of the most widely referenced skew estimation algorithms and Baird has also claimed it to be one of the fastest and most accurate reported.[1]

Two versions of Baird's algorithm were implemented from the description by Spitz,[9] and also run on the same images. One version (Simple_Baird) applies a simple search of every angle in $[-\pi/4, \pi/4]$ with steps of $\pi/20000$. Another version (Fast_Baird) applies a hierarchical search by an initial coarse search of every angle in $[-\pi/4, \pi/4]$ with steps of $\pi/200$. This coarse search is improved by choosing the best angle and narrowing the search by a factor of 10 with a factor of 10 finer resolution. A third iteration is used to give a final effective search equivalent to Simple_Baird.

The majority of images in this test set have very little skew, but at the accuracy of these algorithms, the skew is non-zero for most of them. To thoroughly test the accuracy of each algorithm, each was run on the original image and three rotations

of the original, with the rotations chosen randomly from (-0.5,0.5) radians. The images were rotated in greyscale by bilinear interpolation[7] and then thresholded. Artifacts caused by the rotation were therefore negligible.

The difference between the skew observed on the rotated image and the skew observed on the original should equal the rotation applied to the original. The random skew applied and the change in observed skew can be plotted on a scatter diagram to visually check the accuracy of the algorithms. The scatter diagrams in Figure 1 show the results of applying the new algorithm and Simple_Baird to the 400 pages, each with three random rotations as described above.
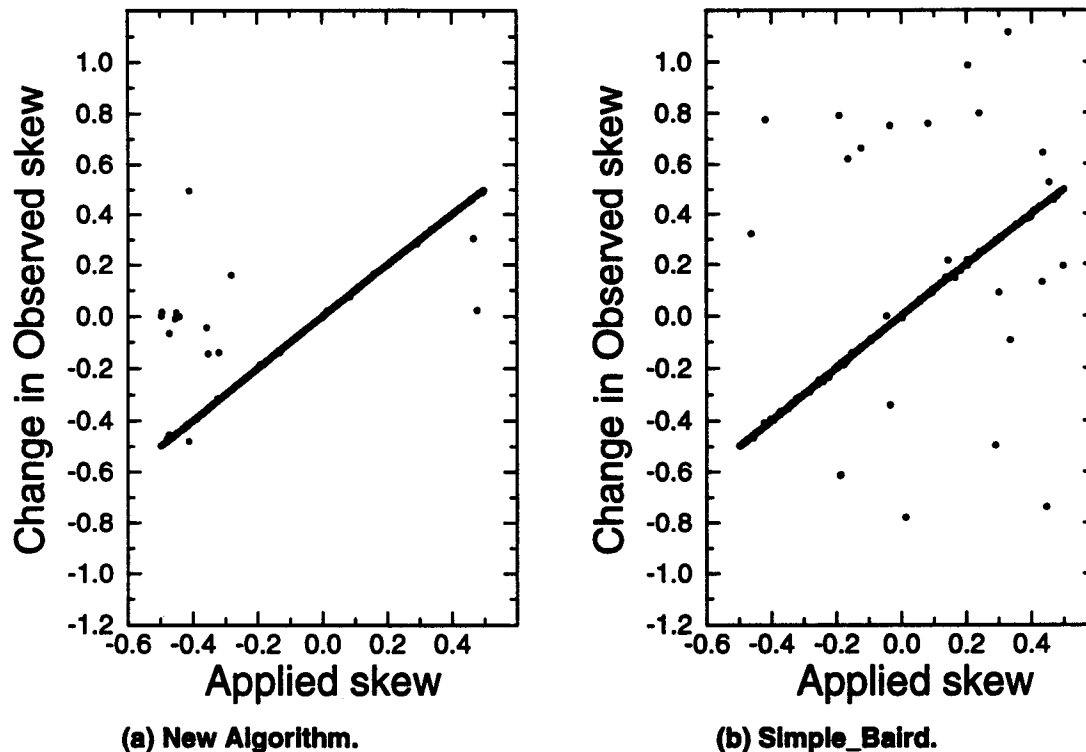

(a) New Algorithm.   (b) Simple_Baird.

Figure 1 Scatter diagrams of change in observed skew against applied skew.

It is clear that both algorithms basically measure the skew with a high degree of accuracy, but that Simple_Baird has a greater tendency to fail dramatically, even on documents with almost no skew. It can also be seen from the y=x line formed by the majority of the points that Simple_Baird also has a greater level of noise than the new algorithm and this is borne out by the numeric data in Table 2.

The scatter diagram for the Fast_Baird algorithm is shown in Figure 2. As might be expected, the price paid for the higher speed is an increasing tendency to fail with a very large error. These additional failures are due to a local maximum in the variance of the projection profile at an incorrect angle, with the global maximum falling between angles in the coarse search.

Investigation of the failures of Simple_Baird showed several factors:

4

Figure 2 Scatter diagram of change in observed skew against applied skew for Fast_Baird.

- Noise. Noise from pictures and the background seem to contribute to failure. This might be fixable by using a blob filter similar to the new algorithm.

- Tall, narrow columns. It is possible to hallucinate long runs of blobs at angles close to 90 degrees away from the correct angle when the text column is significantly taller than it is wide.

- Fixed pitch fonts. It is much easier to produce the problem above when the font is fixed pitch since at certain angles it is easy to fit lines to the text.

The accuracy and time performance of the algorithms are shown in Table 1. The

Table 1. Accuracy and time performance of the skew detection algorithms

Algorithm	Mean angle err/radians	RMS angle err/degrees	RMS pixel err/pixels	Number of outliers	Mean CPU/sec	SD CPU/sec
CC analysis					10.11	9.44
New	3.77E-3	2.83	77.8	15	0.32	0.28
Simple_Baird	-2.59E-3	7.69	217	27	47.20	47.17
Fast_Baird	-4.95E-3	11.7	341	57	1.02	0.94

mean angle error really only confirms that there is no significant constant offset in the error of any of the algorithms. The Root-Mean-Square (RMS) angle error is the most statistically meaningful measure of angular accuracy. A more useful measure

however is in terms of vertical pixel error. This measure shows how close a baseline would be if fitted to the longest line on the page with the aid of the global skew alone. The RMS pixel error shows the vertical pixel error resulting from applying the angular error on each page to the widest column on that page. The CPU timings relate to the Hewlett Packard 9000/720 workstation on which the algorithms were run. The line labelled "CC analysis" is the connected component analysis used by both Baird's algorithm and the new algorithm.

The error figures are rather large, since they include all the results. Each algorithm failed drastically on some documents. Documents with pixel errors in excess of 100 are shown as outliers in Table 1. To obtain a better estimate of how the algorithms perform in cases where they do not fail drastically, the union of all the outliers was deleted from the statistics. The number of samples deleted was 73. The statistics resulting from the remaining common subset are shown in Table 2.

Table 2. Accuracy of the skew detection algorithms without outliers

Algorithm	Mean angle err/radians	RMS angle err/degrees	RMS pixel err/pixels
New	-5.93E-5	0.0718	1.87
Simple_Baird	3.96E-5	0.139	3.69
Fast_Baird	1.07E-4	0.132	3.86

5 Conclusions

An approach to skew detection has been described which is quite different to that usually employed. The method accumulates rows of characters independently of skew and then computes the skew from the rows of characters. Using a large test set, the new method has been shown to be more accurate and somewhat faster than Baird's algorithm. Perhaps the most important feature of the new algorithm is that it has greater robustness with respect to noise, causing it to produce fewer wild estimates of skew, especially for small angles.

6 Acknowledgements

The author would like to thank Chris Newton and Phil Cheatle for their useful comments.

7 References

[1] Baird H.S. Anatomy of a versatile page reader. *Proc. of the IEEE,* **vol. 80** no. 7, pp1059-1065 July 1992.

[2] Baird H.S. The skew angle of printed documents. *Proc. 1987 Conf. Society of Photographic Scientist and Engineers*, Rochester, N.Y, May 20-21 1987.

[3] Hinds S.C, Fisher J.L, D'Amato D.P. A document skew detection method using run-length encoding and the Hough transform. *Proc. 10th Int. Conf. on Pattern Recognition,* Atlantic City NJ 16-21 June 1990 **vol. I,** pp464-468.

[4] Hough, P. Method and means for recognizing complex pictures, U.S. Patent no. 3069654. 1962

[5] Nakano Y, Shima Y, Fujisawa H, Higashino J, Fujinawa M. An algorithm for the skew normalization of document image. *Proc. 10th Int. Conf. on Pattern Recognition,* Atlantic City NJ 16-21 June 1990 **vol. II,** pp8-13.

[6] Pratt W.K, Capitant P. J, Chen W, Hamilton E.R, Wallis R.H. Combined symbol matching facsimile data compression system. *Proc. of the IEEE,* **vol. 68** no. 7, pp786-796 July 1980.

[7] Rosenfeld A, Kak A.C. *Digital picture Processing,* **vol. 2,** 2nd ed. Academic Press 1982.

[8] Roth G, Levine M.D. Segmentation of geometric signals using robust fitting. *Proc. 10th Int. Conf. on Pattern Recognition,* Atlantic City NJ 16-21 June 1990 **vol. I,** pp826-831.

[9] Spitz A.L. Skew determination in CCITT Group 4 compressed document images. *Proc. [First] Symposium on Document Analysis and Information Retrieval.* Tropicana Hotel, Las Vegas Mar 16-18 1992 pp11-25.