

# Detecting a Document's Skew: A Simple Stochastic Approach

A. Hennig, G. Raza, N. Sherkat, R. J. Whitrow

*Department of Computing, The Nottingham Trent University,*

*Burton Street, Nottingham NG1 4BU, UK*

*{amr,ghr,ns,rjw}@doc.ntu.ac.uk*

## Abstract

*A simple stochastic approach to the detection of the skew angle of a scanned document is proposed in this paper. The method first estimates the undirected skew within the range of  $-90^\circ$  to  $+90^\circ$ . I-dots and full-stops are then used to determine whether the document has been scanned upside down, yielding the directed skew. Only a fraction of the information available in the document is exploited, without assumptions as to the general layout of the page. The method has been shown to be robust and accurate for a variety of documents containing both hand-written and printed text. The average error of the directed skew angle was observed to be  $0.05^\circ$  for a set of synthetic documents. For facsimile images, however, the upside-down detection failed in 9.6% of the documents examined.*

## 1 Introduction

The detection of the overall skew of an electronic document is a vital early step in its analysis and recognition. In the majority of documents, lines of text are oriented horizontally. In a rotated document, skew angle is defined as the difference between the dominant orientation of the text lines and the horizontal. When a document is scanned, skew might be introduced in several ways. A relatively small skew might be added by imprecise feeding of the paper into the scanner or fax machine. A document in landscape orientation might have to be scanned in portrait orientation if the size of the scanner is insufficient. The resulting image is then rotated by  $+90^\circ$  or  $-90^\circ$ , depending on the preferences of the person that feeds the document. The document might even be inserted upside down, resulting in an additional skew of  $180^\circ$ . If a paper document is photocopied before the scanning, these effects might accumulate further.

In this paper, the angle of the '*directed skew*' (denoted by  $\varphi, \varphi \in (-180^\circ, +180^\circ]$ ) describes the total rotation that has been applied to the original (right way up) document. The '*undirected skew*' (denoted with  $\bar{\varphi}, \bar{\varphi} \in (-90^\circ, +90^\circ]$ ),

however, assumes that the document has been scanned with a normal orientation. Most documents have a uniform skew. If they are moved during scanning or copying, or if the original contains curved or non-parallel lines of text, the skew angle is not constant throughout the document. This problem has been addressed in work described in [6][3] and is beyond the scope of this paper.

Various methods have been applied to the detection of a uniform skew. Projection profiles of connected components have been used by [2]. The Hough transform has been widely used, e.g. by [4], [1] and [9], even though it is computationally expensive and usually requires an assumption about the interval that contains the skew. A multi-layer perceptron is used for cursive handwriting in [5], whereas [8] applies a least square linear regression to reference points found on the page. Most of these methods, however, assume a right-way-up document, or even require that lines of text can be identified correctly.

The method presented in this paper uses a simple stochastic method to estimate the overall undirected skew of the document. A second step aims to detect whether the document has been scanned upside down, thus obtaining the directed skew. These steps are described in the following sections, followed by experimental results and concluding remarks.

## 2 Detection of the undirected skew angle

The undirected skew detection algorithm is based on the fact that elements of one straight line of text correlate with each other. The directions between elements of the same line of text distribute around the skew, whereas the connection to other lines can appear in almost any other direction. The histogram obtained from the angles of all possible connections hence shows a maximum at the desired skew angle. Elements might be black pixels or the centres of regions of connected black pixels. The example in Fig. 1a shows the connections from the centre of the region that forms the letter 'n' to the other regions of the page. Fig. 1b shows the resulting histogram with a clear maximum at the document's skew.

