Phil The Box: Estimating the Distribution Range of Vienna's free of cost Newspapers.

Andreas Zweng[1,2], Thomas Poenitz[1] and Martin Kampel[1]

[1]Cogvis GmbH
Wiedner Hauptstrasse 17/3a, A-1040 Vienna
[2]Vienna University of Technology
Computer Vision Lab, Favoritenstr. 9/183-2, A-1040 Vienna
zweng@caa.tuwien.ac.at
[kampel;poenitz]@cogvis.at

1   Introduction

A free daily newspaper is available for the citizen in Vienna on approximately 1000 public places, predominantly in the area of subway stations. In order to control the proper filling of the newspaper boxes, pictures are taken from the boxes before and after the filling from responsible employees. The estimated 3200 daily pictures are manually annotated afterwards to determine the print run. The goal of the project (Fill ("Phil") the Box) is to detect the paper stack automatically, which should save time and lead to a more objective result compared to manual analysis.

Fig. 1. Detection of markers and the box id (green circles) and the estimation of the paper stack limits (blue circles).

2   Methodology

The print run estimation is a framework based on algorithms to detect markers, text and the paper stacks. The pipeline is as follows. The algorithm detects two markers in the image which are located on the right side of the box, then the box number is localized and the stacks in between are segmented. The sizes of the

Fig. 2. Template matching using SIFT and RANSAC.

stacks are then estimated using the markers as bottom and top limits and the box number is identified using OCR. The top info of the box is localized at the end using feature matching and the additional information of the box position. As markers we use concentric contrasting circles to detect the markers and the rotation of the image [2]. The top info of the box has to be detected in order to detect if special advertisement is used or if it was removed. We use a template image and compute SIFT features [3] to find a match in the image using the RANSAC algorithm [1] (see Fig. 2).

Optical character recognition is used to identify the unique box number. We detect areas with a high gray value and use adaptive local threshold values to segment the character areas in the image. Candidate regions are then used to compute confidence values measuring if a specific area is a character and also if there is another character in the neighborhood. The confidence values are computed using character templates which are slided over the image in the region of the segmented characters to find the best matching position of the characters. To detect the limits of the paper stacks, we use seed points on bright pixels and a modified region growing algorithm which grows rectangular in order to support the shape of the paper stacks. The result of the algorithm are rectangles in the paper stack area which are combined to one big area. The resulting area includes the top and the bottom of the paper stack which is used to compute the fill level of the box.

First field tests have shown an overall precision of 76,3% of correctly analyzed pictures. After some training of the employees how to take the pictures, the system is expected to reach up to 90% accuracy. Final results will be presented at the conference, when the system is already up and running.

References

1. Fischler, M.A., Bolles, R.C.: Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Tech. Rep. 213, AI Center, SRI International, 333 Ravenswood Ave., Menlo Park, CA 94025 (Mar 1980)
2. Gatrell, L., Hoff, W.: Robust image features: concentric contrasting circles and their image extraction. SPIE 1612, 235–244 (1992)
3. Lowe, D.G.: Object recognition from local scale-invariant features. In: Proceedings of the International Conference on Computer Vision-Volume 2 - Volume 2. pp. 1150–1157. ICCV '99, IEEE Computer Society, Washington, DC, USA (1999)