A Comparison of Methods for Detection and Recognition of Playing Cards

João Pimentel and Alexandre Bernardino
Instituto de Sistemas e Robótica, Instituto Superior Técnico
joao.pimentel@ist.utl.pt, alex@isr.ist.utl.pt

Abstract

We present a playing card detection and recognition method for calibrated camera systems. The combination of a voting based scheme card detection method and a probabilistic recognition method based on edge distributions is able to successfully recognize rotated and occluded playing cards under wide illumination conditions.

1. Introduction

To prevent game losses due to mistakes or irregular strategies, or for statistical data acquisition, the casino game industry is interested in automatic monitoring of the games, recording their state for posterior analysis. We propose a method for detection and recognition of playing cards taking into account partial occlusion and rotation. This is carried out in two phases: a scene is analyzed to detect rectangles, fitting the card size, and then each detected rectangle is classified according to the figure present on the corner of the card. The cards are all assumed of a known size and scale, and the image to be analyzed has no perspective or optical deformations, though this can be taken into account in a previous calibration phase. We study and compare three methods for the card recognition stage, evaluating their performance in terms of robustness to brightness and contrast changes and computation time.

2. Card Detection

Card detection is carried by a voting scheme, based on the Generalized Hough Transform [1]. The voting accumulator is defined in three dimensions (x, y, θ): the card center localization and orientation. The outer contours in the image and the gradient directions are extracted. Each contour point is known to belong to rectangles in just two different orientations: the same as the gradient direction (corresponding to the short side), and the gradient direction plus 90º, corresponding to the large side. This way, the possible rectangles are the ones that contain this point in its contour, in just two known orientations. Each contour point votes on all these locations by incrementing the accumulator accordingly. The detected rectangles are represented by the local maxima of the accumulator. This detection model is simple to implement and, being a voting scheme, performs well in presence of partially occluded cards. It is tuned as to bring no false negatives at the cost of a possibly high number of false positives. The resulting detected rectangles are filtered as to remove the false positives as much as possible. Remaining false positives are rejected in the classification phase if no symbol is found at its top right corner.

3. Classification

Each detected card is rotated to its upright position and the image of its upper right corner is extracted to be classified. Classification was done using various methods for comparison. The rank of the card is first classified, and if considered a valid card (likelihood or other matching criteria above threshold) the suit is also classified. Otherwise, the card is rejected.

3.1. Probabilistic rigid method

A method based on [2] was implemented. It considers binary oriented edge features, in 8 orientations at increments of 45º, computed at each point x of a grid L. Let $X = \{X_e(x) | x \in L, e = 1, ..., 8\}$ be the set of features with orientation e. Having $X_e(x) = 1$ means a feature with orientation e is present at x. We define a probability map, defined in the grid L, as $P(X_e(x) = 1) = p_e(x)$. Assuming the edges present are conditionally independent with marginal probabilities given by $p_e(x)$, the likelihood of the edge data is given by

$$P(X) = \prod_{x \in L} \prod_{e} p(x)^{X_e(x)} [1 - p(x)]^{(1 - X_e(x))},$$

corresponding to a binomial distribution on the feature presence at point x. The classification is done choosing the class with maximum likelihood.

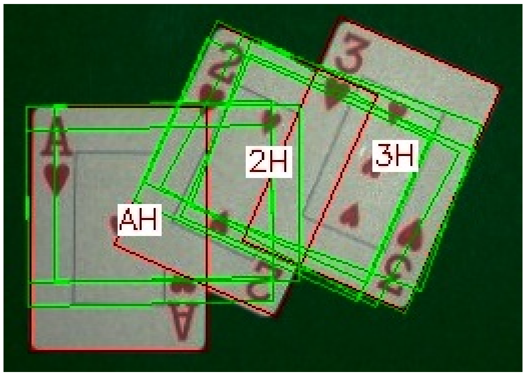

Figure 1. Detected cards in an image.

3.2 Probabilistic deformable method

Proposed in [2] is a *Patchwork of Parts* deformable model. This model is based in the rigid model from 3.1, but instead of a single probability map, it is divided into several overlapping small maps called *parts*. Each one is shifted independently as to adapt to the image to be analyzed. The parts are then combined by an averaging operation on the overlapping regions. The resulting combination provides the probability map to be used in the same likelihood computation as in the rigid model.

3.3. Template Matching

Having a fixed scale imposed on the card images and having the orientation of a detected card known, a natural choice for a classification method is the template matching. Each rank or suit of the detected card is compared with a template from each class, using a cross-correlation criteria. As before, the class chosen for the rank and suit is the class with maximum cross-correlation, if it is above threshold.

4 Results

Overall the method works successfully. The detection phase, when the card size is modeled carefully, delivers no false negatives. It takes a considerable amount of time (1 second for an image with about 13 cards). Partial occlusions pose no problems, except when there are several cards aligned vertically or horizontally. In these latter cases, the outer contours give not enough information about the number of cards present in the middle. These cards cannot be guaranteed to be detected. When at least a whole corner of a card is visible, the detection faces no problems. An example of a detection and recognition is shown on Figure 1. In red are the cards successfully recognized. In green, are the detected cards after rejection in the classification phase.

The classification phase is the most sensitive one. Having a possibly large number of false detection positives, the classifier should be able to correctly classify just the real detected rectangles, and reject any other. Both the rigid probabilistic model and the template matching perform this task. The first, being based on edge features, is relatively invariant to illumination conditions. The success of the latter can depend strongly on these conditions. This is shown in Figure 2, where the recognition rates with the probabilistic method behave well for large variations of brightness and contrast, whereas the template matching only works in small variations on the training images. The values are considered as 100 being the most positive variation possible, and -100 the most negative one. The probabilistic method is, however, a lot more computationally expensive when compared to the template matching, resulting in a performance as much as 200 times slower. The template matching can take as little as 0.02 seconds in an image with 13 cards. The deformable model works extremely slow. It correctly classifies a card rank when its image is deformed, e.g by an affine transform. However, in normal conditions it fails more frequently than the other methods.

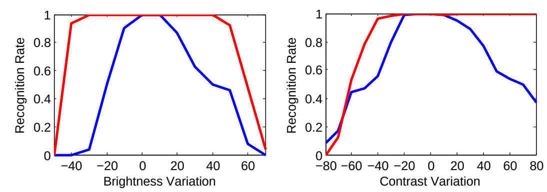

Figure 2. Recognition rates with brightness (left) and contrast variation (right) for the probabilistic rigid model (red) and template matching (blue).

Acknowledgments

This work was partially funded by FCT (ISR/IST pluri-anual funding) through the PIDDAC Program funds.

References

[1] D. Ballard, "Generalizing the Hough transform to detect arbitrary shapes," *Pattern recognition*, vol. 13, no. 2, pp. 111–122, 1981.

[2] Y. Amit and A. Trouvé, "Pop: Patchwork of parts models for object recognition," *International Journal of Computer Vision*, vol. 75, no. 2, pp. 267–282, 2007.

[3] D. H. Ballard and C. M. Brown, *Computer Vision*. Prentice Hall Professional Technical Reference, 1982.