Rapid and Robust Algorithms for Detecting Colour Targets

J. Coughlan, R. Manduchi*, M. Mutsuzaki* and H. Shen
Smith-Kettlewell Eye Research Institute, San Francisco, CA
**Dept. of Computer Engineering, Univ. California Santa Cruz, Santa Cruz, CA*
Corresponding author: J. Coughlan (coughlan@ski.org)

ABSTRACT

We introduce a computer vision algorithm for the rapid and robust detection and localization of colour targets. It is intended to be run in real-time on an inexpensive handheld PDA/cell phone with colour camera, which has limited computational power. Thus, the targets are specially designed with distinctive colour patterns, so that they are easy to find in typical indoor scenes. In order to handle varying lighting conditions, the detection algorithm draws on certain properties of colour gradients that are invariant to a wide range of illumination conditions. A 9 cm diameter target may be detected up to 10 meters from the camera and takes 0.12-0.5 sec per image in our prototype PDA system, demonstrating its feasibility for real-time use.

1. INTRODUCTION

Colour is a distinctive feature that can vastly simplify visual search for objects. However, in order to exploit the use of colour features, some form of colour constancy is needed to cope with unknown lighting conditions. Unfortunately, current computer vision colour constancy algorithms[1,2,3] generally make restrictive assumptions about illumination conditions and/or require significant computational resources. We propose a computer vision algorithm for the rapid and robust detection and localization of specially designed colour targets. The algorithm draws on properties of colour gradients that are invariant to a wide range of illumination conditions, without needing to implement full colour constancy. It is intended to be used in an environmental labelling system for helping visually impaired people navigate their environment. Since the algorithm needs to run in real-time on an inexpensive handheld PDA/cell phone with colour camera (see Figure 1a), which has limited computational power, algorithmic efficiency is of paramount importance.

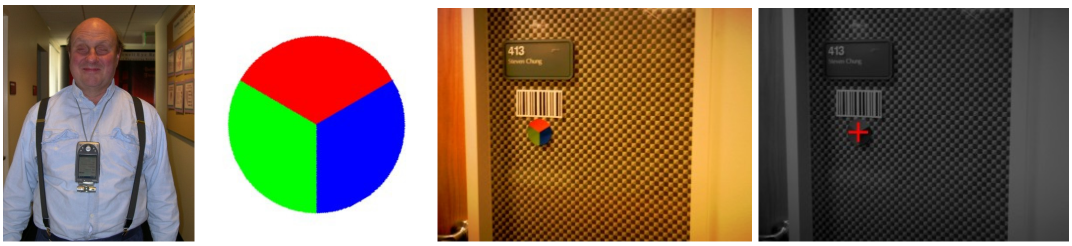

Figure 1: (a) PDA with camera worn by visually impaired user. (b) Colour target consists of three colour patches. (c) Target shown in typical environment (e.g. below barcode). (d) Detection result from (c) indicated by red cross.

Our algorithm for detecting colour targets exploits the fact that, although the apparent target colours depend strongly on lighting conditions, colour *gradients* within them vary much less. Moreover, the spatial and spectral properties of the targets make them easy to distinguish from typical background clutter. Colour targets are located in an image through a series of simple colour gradient tests. Targets may be placed adjacent to signs or barcodes, which can be rapidly located and read given the location of the target (and conveyed to the user by synthetic speech). The targets are small and can be detected from a distance; the maximum range is proportional to the size of the target and is up to 10 meters from the camera using a 9 cm diameter target. Detection takes 0.12-0.5 sec per image in our prototype system, demonstrating its feasibility for real-time use.

2. METHOD

The colour targets were designed with a distinctive combination of colours arranged in a particular configuration (Figure 1b). A series of simple and very rapid tests performed on an input image (Figure 1c) will quickly detect and localize the colour target(s) (Figure 1d). The tests exploit invariants based on colour gradients that we have derived empirically under a variety of indoor and outdoor lighting conditions for our colour pattern. While the precise colour gradients among the three colour patches vary depending on illumination and noise, some aspects of the gradient are highly predictable and create a nearly unique signature of the target.

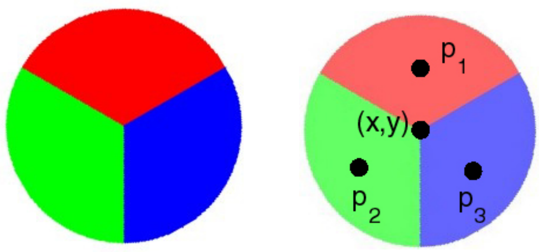

Figure 2: (a) Colour target pattern. (b) Sample point locations p_1, p_2 and p_3 in red, green and blue regions are defined by fixed offsets relative to the center (x,y).

Four sequential colour gradient tests suffice to rule out all but a small fraction of the image pixels that do not lie on a colour target. These tests are based on the following gradient components: the red channel gradient component across the red-green and red-blue boundaries and the green channel gradient component across the red-green and green-blue boundaries. The gradients are estimated by computing differences in RGB channels between nearby pixels. Colour boundaries between adjacent regions are rarely sharp in real images -- because of effects such as colour bleeding, motion blur and pixel interpolation -- so we calculate these gradient differences across a distance of several pixels rather than between neighboring pixels. We choose this distance to be as large as possible, consistent with the requirement that the samples used to compute the gradients all fit within the target region. (The minimum scale of the target in the image -- which is inversely proportional to the maximum distance it can be resolved from the camera -- thus determines the maximum allowable distance between samples.)

The colour gradient tests are performed as follows. Each pixel in the image is considered in turn as the possible location of the center of a target. Given any candidate target center with pixel coordinates (x,y), three points p_1, p_2 and p_3 are computed with fixed offsets relative to (x,y), corresponding to nearby, equidistant locations in the red, green and blue patches, respectively (Figure 2b). The four tests are the following inequalities (in sequence):

$$R_1 > R_2 + T,\ G_2 > G_1 + T,\ G_2 > G_3 + T \text{ and } R_1 > R_3 + T \quad (1)$$

Here the subscripts refer to the RGB intensity channels (ranging from 0-255) at points p_1, p_2 and p_3, and T is a threshold (empirically chosen to be 20). The tests are applied in a loop that scans the entire image to test each possible candidate center location. Testing of a location ends as soon as any of the inequalities fail. The overwhelming majority of pixel locations fail the first or second tests, which means most of the image requires very little processing. Locations that pass all four tests are classified as on-target. Typically, a cluster of many pixels at or near the center of a target are classified as on-target; in some cases pixels closer to the edge of the target are also classified in this way.

Additional processing is needed to rule out the few false positives that survive the filter cascade. In our preliminary experiments we have implemented a grouping procedure which classifies sufficiently large clusters of on-target pixels as belonging to a single target (by contrast, most false positive pixels are scattered thinly throughout the image). Since relatively few pixels are classified as on-target in most images, the grouping procedure adds negligible cost to the overall target detection algorithm. In the future we plan to implement more discriminating tests to verify the existence of candidate targets in the image by verifying the presence of the entire three-colour target.

3. RESULTS

We implemented a preliminary version of the algorithm on a Dell Axim X3 (624 MHz CPU) with Veo Photo Traveler 130S 1.3 megapixel camera. The algorithm detects multiple targets in about 0.12 sec - 0.5 sec (depending on camera resolution). From our preliminary tests we estimate a detection rate of 95%. The current false positive rate is non-negligible (about one false positive target per every 4 images), but we plan to add more stringent target verification tests in the future to reject the false positives and simultaneously allow us to increase the detection rate. (False positives typically occur in parts of the image with several colours in proximity to each other; some of the colours may be artifacts of colour bleeding or motion blur.) The detection is invariant over a range of scales (from about 0.5 m to as far as 10 m), and accommodates significant rotations (up to about 30 degrees in the camera plane) and slant. See Figures 3 and 4 for detection results.

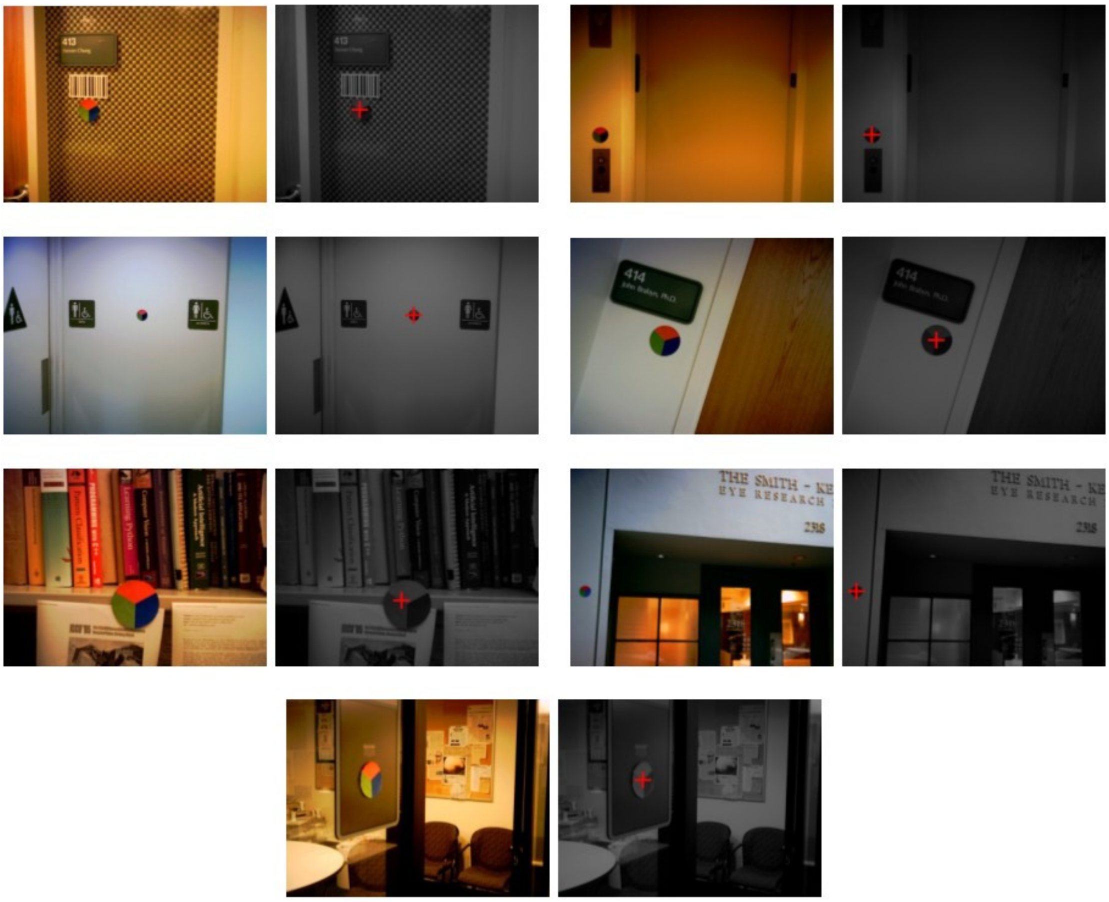

Figure 3: Image pairs show original image on left and detection result on right marked by red cross. Image descriptions (left to right, top to bottom): office door with barcode; elevator; restroom; rotated target; target against highly cluttered background; building entrance; target viewed from oblique angle.

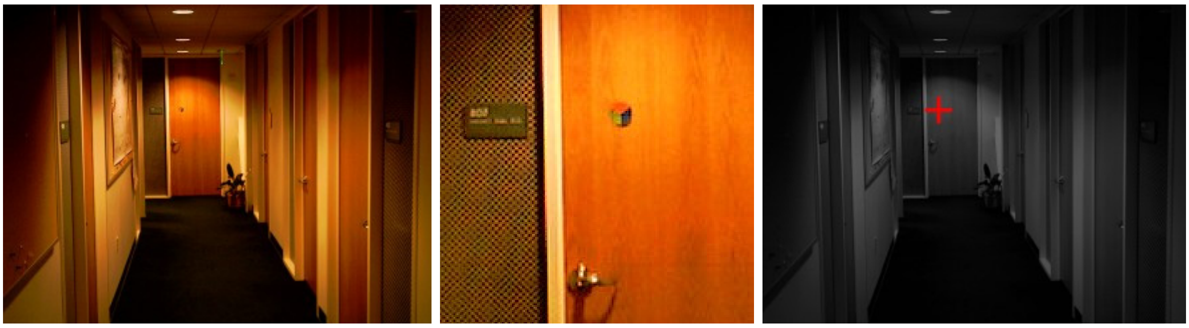

Figure 4: Long-range target detection. Target is 9cm in diameter and at range 10m. (a) Original image. (b) Zoomed-in view of target. (c) Detection result.

4. CONCLUSIONS

Our experiments demonstrate the feasibility of our proposed algorithm, which allows us to detect and locate colour targets in real time on a PDA.

Basing the search algorithm on the identification of specific colour gradient features, rather than less distinctive grayscale features such as edges, is crucial to the efficiency of our algorithm. Colour and colour gradients have been used in computer vision and image processing primarily for such purposes as skin detection[4], contour fitting[5] and edge detection[6]; a variety of work[7,8] has also been undertaken on the design and use of specially designed, easily localized landmarks (i.e. targets), some of which has exploited the use of colour. However, to the best of our knowledge our work is the first application of colour gradients to target detection that is fast enough to be implemented on a PDA.

References

1. B.V. Funt, K. Barnard, K and L. Martin. "Is Machine Colour Constancy Good Enough?", European Conference on Computer Vision, 1998.
2. D. Forsyth. "A Novel Approach for Color Constancy", International Journal of Computer Vision, 5:5--36, 1990.
3. D.H. Brainard and W.T. Freeman. "Bayesian Color Constancy", J. Opt. Soc. Amer.-A, 14:1393--1411, July 1997.
4. M. Jones and J. Rehg. "Statistical Color Models with Application to Skin Detection." International Journal of Computer Vision. Vol. 46 , Issue 1, pp 81 - 96. January 2002.
5. T. Gevers, S. Ghebreab and A. Smeulders. "Colour Invariant Snakes," British Machine Vision Conference. 1998.
6. J. Geusebroek , A. Dev , R. Boomgaard , A. Smeulders , F. Cornelissen , H. Geerts, "Color Invariant Edge Detection." Proceedings of the Second International Conference on Scale-Space Theories in Computer Vision, p.459-464, September 26-27, 1999.
7. Y. Cho and U. Neumann. (1998) "Multi-ring color fiducial systems for scalable fiducial tracking augmented reality". In: Proc. of IEEE VRAIS.
8. D. Claus and A.W. Fitzgibbon. (2004) "Reliable Fiducial Detection in Natural Scenes", in Proc. ECCV, 2004.