

# Reducing images to sets of colored segments: A real-time algorithm with applications in tracking and interpretation in structured environments

## Abstract

*We describe an algorithm to compute a set of segments characterized with color information from an input image. The method is reversible, in the sense that an approximate reconstruction of the original image from the segments is possible. Therefore, it can be considered as a lossy compression method, useful for scenes where segments are the main primitives, such as interior of buildings, man-made objects, and so on. The compression factor is interesting by its own –approximately two orders of magnitude with respect to the original image, depending on the structure and detail of the scene–, but perhaps the main advantages of the algorithm are other: First, the compressed data can be very useful for further processing, as they are based on a simple though powerful geometric primitive, the segment with color information along each of its two sides. Second, the algorithm works in real time with low-cost hardware, such as an off-the-self image acquisition card and a standard PC. We show examples of segments extraction and reconstruction on some input images, discussing compression factors and execution times. Finally, we propose possible applications of the algorithm.*

## 1. Introduction

Segments are powerful medium level primitives commonly used in machine vision systems where subsequent postprocessing consists in 3D reconstruction of the scene [8], or the identification of patterns in it, perhaps using a database of known objects [1]. These problems have been active research topics in the last years, mostly due to significant advances in the application of projective geometry to computer vision [4].

Projective geometry uses points and lines as its main primitives. Segments and interesting points (also known as *features*) are easily found in structured environments. Furthermore, although there exist many methods to find corners in an image, they can also be easily localized by crossing two segments whose extremes are near enough.

Color information, on the other hand, has also been successfully used in several vision problems, such as segmentation [2], tracking [6], or learning and classification [3]. But these approximations usually discard any geometric information, often working only with unstructured sets of spatially related pixels.

In this paper, we propose an efficient algorithm to reduce images to sets of colored segments, trying to take advantage of both the color information and geometric structure of the scene. We justify that the obtained output data effectively summarize the contents of the input image, as using them we can recover a very approximate reconstruction of the original image. Finally, given the efficiency of the algorithm, we propose several applications of the algorithm in systems with real-time requirements, such as robot navigation and environment reconstruction and interpretation.

## 2. Segment Detection

Finding lines or segments in an image is a basic problem studied in many generic computer vision references. Most of the available methods are different variations of the well known Hough Transform [5]. Other authors, on the contrary, prefer to find segments by locally grouping edge pixels into contours (for example, using the Canny operator, and then approximating them by piecewise linear polygonal approximations [4]).

In fact, the procedure used to get the segments is irrelevant to our algorithm. The only requisite is that it must be fast, as it is intended to be used in real-time applications. In our case, we use a method based on an inexpensive high-pass filtering of the image, followed by an efficient sequential grouping stage, that takes into account the local orientation of the edges. The method cannot be explained here because of the space limitations. Suffice to say that it is implemented using the Intel Image Processing Library (optimized to run in Pentium processors), and that it works at 10-20 fps (depending on the complexity of the scene) for 288×384 indoor images running on a 533 MHz CPU. We show some examples of the results obtained with the algorithm in sections 5 and 6.

### 3. Characterizing Segments with Color

Once the segments of the scene have been found, the next step is to label them with robust color information. We use the median gray value (or RGB, in color images) of image pixels in both sides of the segment. Figure 1 shows an algorithm that generates a table to run orderly through the environment of a given input segment. It is based on Bresenham's algorithm, a standard computer graphics method for drawing lines on pixel grids. We extend this algorithm to "thicken" the segment, by replicating the pixels in both sides of it. Observe that the algorithm uses only integer operations, thus being very fast.

Tables computed with this algorithm have two advantages: First, they cover the segment environment in a *dense* way, that is, without leaving uncovered pixels. Second, they arrange the pixels in the environment in an ordered way, by means of a rectangular array of positions (Figure 2 shows the table generated for a given example segment).

Having computed these tables, it is easy to find the median value of gray for the left side of each segment:

$$I_{med}^{left} = \text{median}(\{I_{i,j}^y, t_{i,j}^x\}_{i=0 \dots \lfloor \frac{m}{2} \rfloor - 1, j=0 \dots n-1})$$

Where  $I_{y,x}$  is the gray value at pixel position  $(y, x)$ . The right side value can be obtained analogously using  $i = \lfloor \frac{m}{2} \rfloor + 1 \dots m - 1$ . The median gray value of each side can be efficiently computed, for example, by accumulating the pixel values in a histogram table, and then running through it from left to right, adding the histogram values until the cumulative sum exceeds  $(\lfloor \frac{m}{2} \rfloor n)/2$  (half the total number of pixels for each side). The value at which this number is reached is the expected median value of the distribution. Finally, extension of the procedure for RGB images is trivial: we only have to apply the former equations independently to the red, green and blue channels of the image.

### 4. Image reconstruction

The procedure that recovers the original image from the set of colored segments also uses the tables described in section 3. The complete algorithm is summarized in figure 3 (for the most general case of RGB images). The pixels located immediately to the left and right of each segment are initialized and fixed with their respective RGB median values. The rest of pixels are uninitialized. Then, a simple iterative diffusion procedure is performed, in which in each iteration a RGB pixel value is recomputed as the mean of its initialized neighbors in a  $3 \times 3$  window. The algorithm keeps on running until a desired convergence level is attained (typically 200-400 iterations, depending on the number of segments and the desired precision).

```

Input: (Segment extremes and desired width)
(y1, x1), (y2, x2), W.
Output: (m × n table of image positions)
(ti,jy, ti,jx), for i = 0, ..., m - 1, j = 0, ..., n - 1.

Algorithm:
Initialization: (Rounding, increments and table dimensions)
y1 := [y1]; x1 := [x1]; y2 := [y2]; x2 := [x2];
dy := |y2 - y1|; dx := |x2 - x1|;
m := 2[ $W \frac{\sqrt{dy^2 + dx^2}}{2(dx+1)}$ ] + 1; n := dx + 1;
Main segment: (Bresenham's algorithm)
e := 2dy - dx; y := y1; x = x1;
for j := 0 to dx - 1 do begin
  t[ $\frac{m}{2}$ ],jy = y; t[ $\frac{m}{2}$ ],jx = x;
  if e ≥ 0 then begin
    y = y + 1; e = e - 2dx;
  end
  x = x + 1; e = e + 2dy;
end
t[ $\frac{m}{2}$ ],dxy = y2; t[ $\frac{m}{2}$ ],dxx = x2;
Segment environment: (Segment replication)
for i := [ $\frac{m}{2}$ ] + 1 to m - 1 do begin
  p1 := (y2 - y1, x2 - x1) · (ti-1,0y - y1 + 1, ti-1,0x - x1);
  p2 := (y2 - y1, x2 - x1) · (ti-1,dx-1y - y2 + 1, ti-1,dx-1x - x2);
  if |p1| ≤ |p2| then begin /* Complete segment replication: */
    for j := 0 to dx do begin
      ti,jy = ti-1,jy + 1; ti,jx = ti-1,jx; /* Right: */
      tm-i-1,jy = tm-i,jy - 1; tm-i-1,jx = tm-i,jx; /* Left: */
    end
  end else begin
    /* Partial segment replication, plus additional point: */
    for j := 1 to dx do begin /* Right: */
      ti-1,j-1y = ti-1,j-1y + 1; ti,jx = ti-1,j-1x;
    end
    ti,0y = ti-1,dx-1y - y2 + y1 + 1; ti,0x = ti-1,dx-1x - x2 + x1;
    for j := 0 to dx - 1 do begin /* Left: */
      tm-i-1,jy = tm-i,j+1y - 1; tm-i-1,jx = tm-i,j+1x;
    end
    tm-i-1,dxy = tm-i-1,0y + y2 - y1;
    tm-i-1,dxx = tm-i-1,0x + x2 - x1;
  end
end

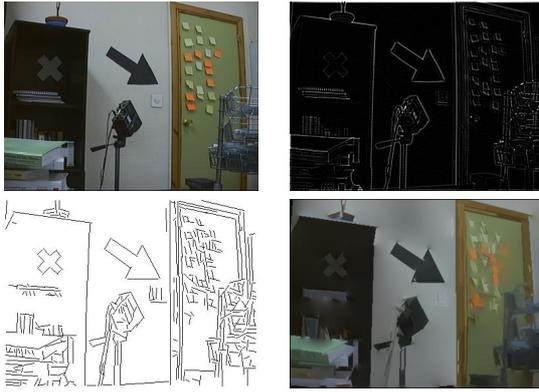
```

**Figure 1.** Algorithm that generates a table to run through the environment of a line. The algorithm works only in the first octant ( $X$  dominates  $Y$ ,  $dx > 0$  and  $dy > 0$ ). Extensions for the rest of octants should be straightforward.

### 5. Results

Figure 4 shows the results of the segment extraction and posterior reconstruction procedure in a typical indoor scene. For this  $288 \times 384$  RGB image, the computing time needed to extract the segments and get the color information was only 0.07 seconds in a Pentium III 533 MHz processor, with a Matrox frame-grabber. That means around 15 fps. The original, uncompressed image size is  $288 \times 384 \times 3 = 324$  KB (using one byte per pixel for each red, green and blue channel). The compressed information is formed by 350 segments, each labeled with two RGB values. Each segment extreme can be coded using  $\lceil \log_2(288 \times 384) \rceil = 17$  bits, and each RGB value using 24 bits. Therefore, the whole set of segments will use  $350 \times ((2 \times 17) + 2 \times 24) = 28700$  bits



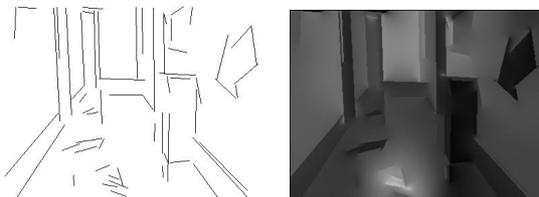


**Figure 4.** (Top-down, left-right) Original image, high-pass filter, segments, reconstructed image.

robot navigation. We propose to characterize each segment only with six additional scalar values, the median RGB values of the pixels along each of its two sides. These values can be used, for example, to compute a *similarity measure* between candidate pairs of segments, that takes into account both the geometric distance (as in [1]) and the color information. Using this measure, we could make the tracking more robust, without sacrificing the speed of the system.

## 6.2. Segment classification

The obtained median color values could also be useful in interpretation. For example, if we search for objects with a characteristic, *a priori* known color (signals for navigation, doors, the floor, etc.), we can use the color information of each segment in order to directly associate it with real objects. Figure 6.2 shows a typical indoor scene taken by an autonomous robot with a camera mounted on it. The robot can easily classify and group the segments by looking at their color information. Currently, this system is being used by our research group in a robot that navigates avoiding obstacles, obeying signals and interpreting the local situation (corridor, corner, room, and so on) in real time.



**Figure 5.** Indoor image taken by our robot: (Left) Segments (Right) Reconstructed image.

## 6.3. Image compression

The compressing capability of the algorithm can also be successfully exploited. The robot mentioned above, for example, is equipped with an onboard PC that is communicated with the rest of the world through a Radio Ethernet interface. These kind of interfaces are usually very slow (1 Mb/s bandwidth in our case). This is not enough to send images of moderate sizes at video rates of 25 fps. But, if we execute the segment extraction algorithm in the robot PC, and then send the compressed information through the network, the only bottleneck will be then the local CPU speed, not the radio interface. The rest of the processing (interpretation, color learning and classification, closed contour interpretation, calibration and 3D reconstruction) is performed in an external, unloaded computer. We have also used this technique in the application described in section 6.2.

## 7 Conclusions

We have described a very efficient algorithm which is able to compress images down to a 1-2% of the original size, while capturing the essential geometric structure of the scene. The algorithm works by detecting relevant segments that are then labeled with color information. The original image can be easily recovered by means of a simple diffusion algorithm. Thus, the obtained data structure contains all the relevant information of the input. This justifies the fact that the original image is no longer needed for any posterior processing. Finally, we have proposed a number of possible applications in real-time computer vision tasks, such as robot navigation and environment reconstruction and interpretation in indoor scenarios.

## References

- [1] J. Beveridge and E. Riseman. How easy is matching 2D line models using local search. *IEEE Trans. on PAMI*, 19(6), 1997.
- [2] J. Bruce, T. Balch, and M. Veloso. Fast and inexpensive color vision for interactive robots. In *Proceedings of the IRS*, 2000.
- [3] S. Buluswar and B. Draper. Color machine vision for autonomous vehicles. *International Journal for Engineering Applications of Artificial Intelligence*, 11(2), 1998.
- [4] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.
- [5] V. Leavers. Survey: Which hough transform? *CVGIP: Image Understanding*, 58:250–264, 1993.
- [6] S. McKenna, Y. Raja, and S. Gong. Tracking colour objects using adaptive mixture models. *Image and Vision Computing*, 17:223–229, 1999.
- [7] C. Schmid and A. Zisserman. Automatic line matching across views. In *Proceedings of the IEEE Conf. on CVPR*, 1997.
- [8] Z. Zhang and O. Faugeras. A 3D world model builder with a mobile robot. *Robotics Research*, 11(4):269–285, 1992.