Techniques for the Interpretation of Thermal Paint Coated Samples*

Andrew Griffin, Josef Kittler, Terry Windeatt, George Matas
Dept. of Electronic and Electrical Engineering,
University of Surrey,
Guildford, Surrey.
GU2 5XH
U.K.

Abstract

Thermal paint, paint that changes colour with maximum temperature, has a variety of uses within the automotive industry. The analysis of samples after exposure to heat is a task currently performed by humans. We present the current findings of our attempts to automate this procedure by means of image analysis. A framework has been developed within which colour is mapped to an RGB space curve which is a function of temperature. This models the context imparted by the physical process, as well as spatial context, and gives us better results than those obtained by using conventional image segmentation techniques.

1. Introduction

The specific task to which this paper relates is the segmentation of images of objects that have been coated with thermochromic paint and then subjected to heat. Thermochromic paint, as its name suggests, is paint that changes colour with temperature. This change is permanent and allows a rough temperature measurement to be made by observation of the colours on the sample; a task currently performed by a human operator.

After heating, the colour profile is such that at certain temperatures, the colour changes quite prominently as perceived by the eye. At these points the temperature has been calibrated empirically; it is by locating these colour boundaries, and then classifying, that the human operator performs his analysis. The aim is to provide a means for the automatic analysis of images of this type, therefore schemes for both identifying the boundaries, and then classifying pixels according to these boundaries need to be developed.

The location of edges, or intensity boundaries in monochromatic images has received much attention [4], [6], [7]. When trying to locate colour boundaries, the problem is one of defining a colour edge, or more specifically, how to combine the gradient measures from the three monochromatic gradients into a salient measure of colour gradient [1], [2], [3], [5]. In our experience colour boundary detection methods failed to give satisfactory results. The main reason is that for some temperature band transitions the contrast is very weak. Similarly, pixel based classification methods produced unsatisfactory results due to feature space ambiguity. We have the benefit of a very specialised problem. This allows us to overcome some of the problems inherent in generalised systems by incorporating high level knowledge about the physical phenomena we are dealing with. This leads us to develop a framework, within which we can obtain results that are an improvement over those obtained using traditional techniques.

The paper is organised as follows. In section 2 we develop a novel model based classification scheme. In the model employed by the scheme temperature variation is mapped on a colour space curve representing the paint. The classification of image pixels is based on the distance to this curve. In section 3 we describe a number of experiments carried out to validate the proposed method. Finally section 4 summarises the main accomplishments and draws conclusions on the feasibility of the proposed approach.

2. Methodology

The paints currently under analysis have been designed with human interpretation in mind. Consequently the colour changes are such that at certain calibrated temperatures the colour change is perceptually 'larger' than at other temperatures where the colour is intended to appear more uniform. The task is therefore one of identifying the colour boundaries which when combined with classification of the uniform regions indicate the calibrated temperatures.

Temperature is related to colour at each image point by the mapping

*The authors would like to acknowledge the support of Rolls Royce plc.

1015-4651/96 $5.00 © 1996 IEEE
Proceedings of ICPR '96

Proceedings of the 1996 International Conference on Pattern Recognition (ICPR '96)
1051-4651/96 $10.00 © 1996 IEEE

$$t(x,y) \longmapsto P(\lambda, x, y)$$

where $t(x,y)$ is the temperature to which the position (x,y) on the object was heated, and $P(\lambda, x, y)$ is the spectral reflectance function of the paint at that point after heating. After image capture we have the situation where each RGB triplet is a representation of $P(\lambda, x, y)$ containing an error due to noise, from which we would ideally like to map back to temperature. However the original mapping is assumed to be not known except in terms of the calibrated colour changes, and the mapping is not necessarily injective. The solution is to map each RGB triplet to some parameter that reflects temperature, ie. the parameter is monotonic with temperature.

$$t \overset{\substack{Chemical\\Process}}{\longmapsto} P(\lambda) \overset{Image}{\longmapsto} \begin{bmatrix} R+\delta R \\ G+\delta G \\ B+\delta B \end{bmatrix} \overset{Interpretation}{\longmapsto} r$$

By observing the distribution of pixels from the image in the RGB cube it can be seen that they form, for each different paint, a curve. The samples used in this initial work were such that the true temperature profile of the image was known. It was a simple matter to verify that the position on the curve varied continuously with temperature. Taking $r(x,y)$ as the arc length from one of the endpoints of the curve to the position on the curve corresponding to $[R(x,y), G(x,y), B(x,y)]^T$ it can be demonstrated that we have satisfied the requirement for monotonicity.

The task can now be split into three distinct subsections.

- Curve approximation
- Mapping pixel values to the curve approximation
- Locating colour boundaries to obtain calibrated temperatures

2.1. Curve Approximation

As mentioned earlier the samples currently under study are such that the true temperature profile is known. By sampling the image along a temperature axis we are left with a set of RGB values ordered by increasing temperature. Each of these is assigned a parameter r, and a polynomial expression may be fitted to the data to give the three equations $R(r), G(r), B(r)$.

In practice however, to reduce computational expense, rather than a polynomial representation, a piecewise linear approximation was used; replacing intervals in the three polynomials by linear sections wherever the curvature in each of the three colour expressions was not significant. The number of points chosen to specify the curve is not critical,

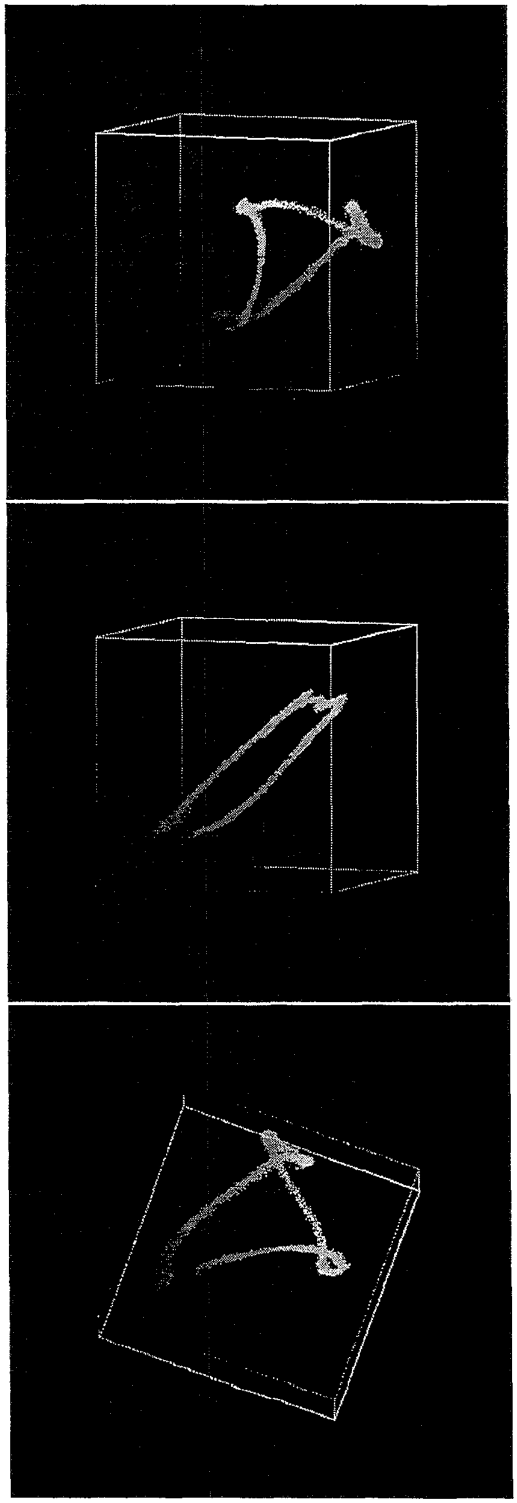

Figure 1. An example of a temperature curve in RGB space with orthogonal views

Proceedings of the 1996 International Conference on Pattern Recognition (ICPR '96)
1051-4651/96 $10.00 © 1996 IEEE

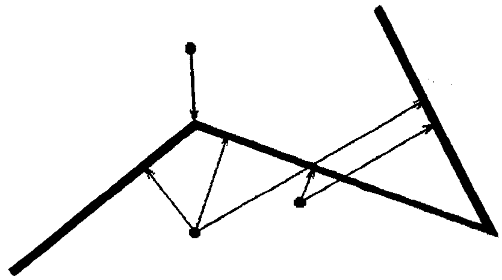

Figure 2. Possible mappings from pixel values to curve

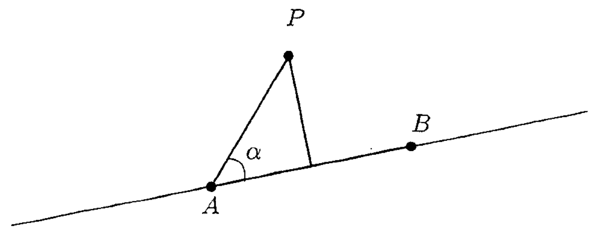

Figure 3. Projecting points onto a line segment

as long as each turning point in each of $R(r)$,$G(r)$ and $B(r)$ is included. However, it should be noted that whilst choosing say just the turning points, will lead to less computation in general, the removal of all curvature between turning points will lead to more computation in the removal of ambiguity, so some tradeoff is required.

All results shown subsequently were obtained using piecewise linear approximations to the polynomial representations of the curve.

2.2. Mapping pixel values to curve position

Given an RGB triplet, we need to find the value for r that gives $[R(r), G(r), B(r)]^T$, as the closest interpretation of that RGB triplet. Intuitively this is the arc length from an endpoint of the curve to the closest point on the curve to that RGB triplet.

In practice, this is achieved by projecting each image point to each line segment. Projections which cannot be made such that the projection is to a point within the endpoints of a line section are discarded unless they map to a corner (line intersection). These various possibilities are illustrated in Figure 2.

$$\cos\alpha = \frac{\vec{AP}.\vec{AB}}{|\vec{AP}||\vec{AB}|}$$

The perpendicular distance from the pixel to the line segment is given by $d_p = |\vec{AP}| \sin\alpha$ and the distance along the line segment to the normal of $\vec{AB}$ which passes through P by $d'_m = |\vec{AP}| \cos\alpha$

We shall normalise this distance by the length of $\vec{AB}$ to give

$$d_m = \frac{|\vec{AP}| \cos\alpha}{|\vec{AB}|}$$

Thus if the pixel projects to a point on the line segment between its endpoints d_m will lie in the interval $[0..1]$.

In this manner the pixel is projected onto each line segment within the curve approximation. Projections with $d_m < 0$ are discarded unless it is a projection onto the first line segment, and projections with $d_m > 1$ are discarded unless it is a projection onto the last line segment. Should a point lie outside a corner in the approximation, it is projected to that corner ($d_m = 0/1$) and d_p is simply the Euclidean distance from the point to that corner.

The value of d_p relates to the uncertainty of the classification. Ideally for one of the projections d_p will equal 0, but unfortunately the approximation to the curve is by no means exact, and the various signals are corrupted by noise. At this point a simple decision scheme would be to choose the projection that requires the shortest distance. The ambiguity introduced by noise and curvature removal however, requires that we consider any of the remaining projections as possibly being the correct projection. For simplicity we assume that the probability of a projection being the true projection diminishes with distance between the image point and the line segment. It is also assumed that the probability of two adjacent pixels having vastly different r values is zero.

A rigourous approach would be to consider all the probabilities derived from consideration of each projection for each image point, but this is considered too computationally expensive. We assume that if the distance for a projection is greater than some threshold, then the probability of that projection occuring is vanishingly small. This threshold depends on the amount of observation noise present, and on the modelling errors caused by the removal of curvature. From this, we may discount some possible projections and assume that we are left with some points for which there is only one possible projection and r can be recovered. From these points a region growing scheme is implemented with the assumption that the temperature gradient is sufficiently small to limit the magnitude of change in r. In this manner

Proceedings of the 1996 International Conference on Pattern Recognition (ICPR '96)
1051-4651/96 $10.00 © 1996 IEEE

many possible projections are excluded, and ambiguity has been reduced. Any points where ambiguity remains are assigned to the nearest projection. This should not prove too costly even if the wrong projection is chosen, provided these areas are not too large.

The classification process is of course equivalent to finding approximate solutions of the equation

$$\begin{bmatrix} R(r) \\ G(r) \\ B(r) \end{bmatrix} = \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

Given the endpoints of each line segment, a solution may have occurred if R, G and B lie between two end points of a line segment. In this case an estimate of r can be recovered say as the root of $R(r) - R = 0$ in that interval using linear interpolation. If $|G(r) - G|$, and $|B(r) - B|$ are both within some threshold, again defined in terms of noise and curvature removal, then r is a solution to the above equation. If there is only one solution then the classification is unambiguous. From these unambiguous points a consistent interpretation of the data may be grown. Any points for which there are no solutions (or no possible projections) are deemed outliers and are not processed.

2.3. Location of Colour Boundaries

At this point we have converted the colour boundary location to a one-dimensional problem. The parameter r gives information about relative temperatures, but also retains some information about the distribution of colours. By taking colour gradient in an image as the Euclidean distance in RGB-space between two adjacent pixels, we can see that our arc-length parameter r can be used to calculate colour gradient. Pixels with large distance between them indicate a large colour gradient. This corresponds to a large change in arclength and therefore;

$$\frac{\partial RGB}{\partial x} = \frac{\partial r}{\partial x}$$

$$\frac{\partial RGB}{\partial y} = \frac{\partial r}{\partial y}$$

The method using r however allows us to threshold globally once a particular value for r at a boundary has been decided.

Another method of locating colour boundaries would be to use the property that colour is roughly uniform between the calibrated colour changes. This will give rise to distinct maxima and minima in the histogram of the image where RGB is replaced by r. The minima represent the values of r where there are fewest pixels, ie. the change points which are the calibrated temperatures.

3. Results

The aim of the above classification scheme is to provide a monotonic function of temperature given the RGB values. That is

$$t_{x,y} \longmapsto \begin{bmatrix} R_{x,y} \\ G_{x,y} \\ B_{x,y} \end{bmatrix} \longmapsto r_{x,y}$$

where $r_{x,y}$ is monotonic with $t_{x,y}$.

In the first instance the curve has been estimated by sampling the pixel values from the image along a temperature axis every few pixels giving rise to upwards of 30-40 points. In Figure 4 the curve has been estimated by examining the colour-space curve of the image and choosing roughly (in practice 6 or 7) the 'most significant' turning points, significantly reducing processing time. This plot of a scaled $r_{x,y}$ versus increasing temperature along an arbitrary temperature axis on the image demonstrates near monotonicity.

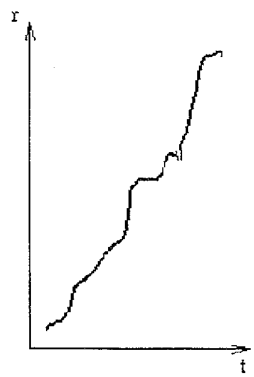

Figure 4. r vs temperature

3.1. Interpreting the results

Once each pixel RGB triplet has been mapped to an $r_{x,y}$ value the problem becomes one of selecting meaningful threshold values on $r_{x,y}$ for the colour boundaries.

Hopefully the mapping from RGB to r will be monotonic with temperature. Any selection of thresholds on the value of $r_{x,y}$ will give rise to temperature bands. As stated earlier, the images take the form of regions of homogeneous colour with diffuse transitions. It should be possible to locate the thresholds such that they lie on these transitions.

3.1.1 1-d Clustering

The images take the form of roughly homogeneous regions of colour with diffuse transition regions. Taking the

Proceedings of the 1996 International Conference on Pattern Recognition (ICPR '96)
1051-4651/96 $10.00 © 1996 IEEE

cumulative histogram of pixels projected to each position on the curve should lead to significant peaks and troughs. The peaks then refer to the homogeneous regions, and the troughs to the transition regions. The threshold values may therefore be selected as say the minima of the 1-d histogram which will correspond to the greatest Euclidean distance between values occuring along each transition and should refer to the perceptual boundaries.

In practice depending on the accuracy of the curve approximation there may be some warping of the clusters as an artifact of the curvature. That is, if the true arc length between two of the vector endpoints is much greater than the length of that vector then the curvature will lead to a significantly non linear projection along the line.

3.1.2 Temperature gradient analysis

Another approach for locating the threshold values is to look at the gradient of the $r_{x,y}$ greyscale image. The position along the curve for the homogeneous regions of colour will be roughly constant, and the position across the transition regions will change more rapidly. If the clusters were distinct in RGB space, ie. there were no transition regions, the gradient of $r_{x,y}$ along any line in the image would take the form of roughly flat sections with steps between them referring to the colour boundaries. If there were no homogeneous regions of colour, the gradient would have a slope $\neq 0$ at each point. In practice the result lies somewhere between. Colour boundaries occur somewhere between each roughly flat section of the temperature gradient, and are initially located at the position of maximum gradient between flat regions. In the results, the thresholds have been chosen in this manner by examining the gradient of the $r_{x,y}$ image along an arbitrary temperature axis on the image.

In Figure 5 the results have been obtained by firstly selecting the thresholds via histogram analysis and secondly by selecting global thresholds along a single arbitrary temperature axis after projection. These are shown in that order, together with the original image.

4. Conclusions

We have presented a framework for the segmentation of images of thermal paint coated samples. In this framework, colour is mapped to an RGB space curve which is a function of temperature. This curve models the context imparted by the physical process, as well as spatial context. The pixel classification scheme based on this model gives better results than those obtained by conventional image segmentation techniques. This has been verified experimentally on real thermal paint image data. We believe that the results are very promising, and that the proposed scheme offers a framework for solving the thermal paint image analysis problem. Work to date has been performed on planar objects. In the future methods will need to be devised to deal with three dimensional objects and the associated problems of shading and possible occlusion.

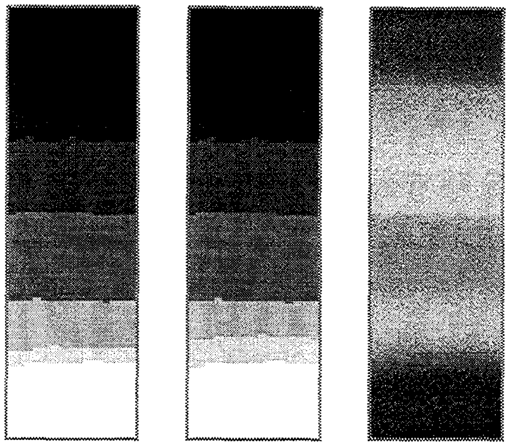

Figure 5. Final results

References

[1] M. Chapron. A new chromatic edge detector used for color image segmentation. In *International Conference on Pattern Recognition*, pages 311–314, 1992.

[2] A. Cumani. Edge detection in multispectral images. *CVGIP: Graphical Models and Image Processing*, 53:40–51, 1991.

[3] S. Di Zenzo. A note on the gradient of a multi-image. *Computer Vision, Graphics, and Image Processing*, 33, 1986.

[4] A. Low. *Introductory Computer Vision and Image Processing*. McGraw Hill, 1991.

[5] R. Nevatia. A color edge detector and its use in scene segmentation. *IEEE Trans. Sytems Man and Cybernetics*, SMC-7:820–826, 1977.

[6] M. Petrou. The differentiating filter approach to edge detection. *Advances in Electronics and Electron Physics*, 88:297–345, 1994.

[7] W. K. Pratt. *Digital Image Processing*. John Wiley and Sons, 1978.

Proceedings of the 1996 International Conference on Pattern Recognition (ICPR '96)
1051-4651/96 $10.00 © 1996 IEEE