Colour Profiling Using Multiple Colour Spaces

Nicola Duffy and Gerard Lacey
Computer Vision and Robotics Group,
Trinity College, Dublin.Ireland
`duffynn@cs.tcd.ie`

Abstract

This paper presents an original method using colour histograms for flaw detection in automated industrial inspection. The colour histogram of an image is constructed by mapping the pixels into a colour space composed of discrete 3-D colour intervals called colour bins. Swain [9] demonstrated the use of colour histograms to locate a target object in an image. The colour histograms of the target object and of the image are used to create a ratio histogram. The ratio histogram is back projected onto the image by replacing each pixel in the image by the value of the ratio histogram bin into which it falls. The resulting image is smoothed to reduce the effects of noise. The peak is found and taken to be the location of the target object. The method presented in this paper also uses histogram back projection; its novelty is in the construction of the colour profile, which is back projected in the place of the ratio histogram. The colour profile is the result of presenting the system with a set of training images; some known to contain flaws some without flaws. The colour profile is constructed by comparing the colour histograms of the images of flawed components to those of non-flawed components. This results in an array of weights, proportional to the frequency of each discrete colour occurring in flawed versus non-flawed images. The correct choice of colour space and the quantisation of the colour space are very important. These issues are thoroughly explored in this paper and a solution is presented which maximises both accuracy and speed, by using the RGB and HSV colour spaces in the training phase, and combining them into a single RGB based profile for real-time flaw detection. Using colour profiles has the following advantages: they are accurate, real time, and very importantly can be trained by the end user. Colour profiling could be applied successfully to other automated inspection problems.

1 Introduction

This paper describes the use of colour profiling to perform industrial inspection on moulded plastic filter casings. The filter casings are manufactured in a variety of different colours and structural designs. They are transparent and have a highly complex structure,

causing many internal shadows. The casings need to be inspected for small burn marks and discolourations. The flaws vary not only in size and location but also in colour. The colours present in a burn cover a range that includes yellow, green, brown, and black to blue. The parts are small (approx. 3cm in diameter) and complicated and because of this manual inspection by humans requires a high level of concentration. In fact human inspection has proven to be very inaccurate; only during the first half-hour are humans able to perform significantly better than random. Even though 6-8 different humans inspect each filter casing many flawed casings are not discovered which has created the need for an automated inspection system.

Vent in the
moulding - not
a flaw

Flaw
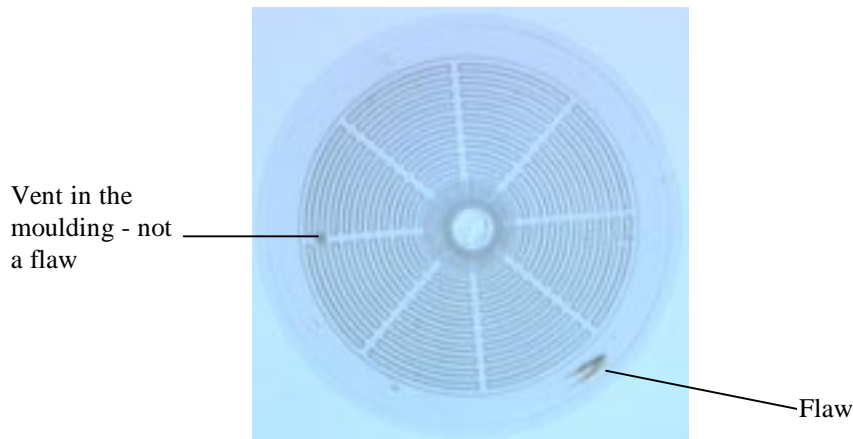
The casings are manufactured at the rate of 1.4 per second. Due to the complexity and real time requirements established vision techniques which search for structural defects are not applicable. Instead of focusing on the more difficult question of what makes a good component, the approach taken in this paper is to search for the presence of a flaw, based on colour, and treat the rest of the component as background. Although not all of the flaws have uniform colour, colour profiling has been successful in detecting the range of flaws. This approach does not take into account of the structure of the casings and therefore can easily be trained to inspect objects with colour separable flaws. The accuracy at which flaws can be found in the training set is recorded and used as an estimate of the ability to detect flaws in unknown images.

2   Colour Histograms

A colour histogram of an image is constructed by mapping each pixel onto a discrete colour space containing n colour bins. The j^{th} colour bin B_j is defined in the following way:

$$B_j = \left\{ (x, y, z) \quad \middle| x_{j_x} \leq x < x_{j_x+1}, y_{j_y} \leq y < y_{j_y+1}, z_{j_z} \leq z < z_{j_z+1} \right\}$$

where (x, y, z) are the three colour channels. A pixel, which falls into a particular bin, is said to occupy that bin. The colour histogram $H(M)$ is a vector $(m_{c_1}, m_{c_2}, \ldots, m_{c_n})$, where m_{c_j} is the number of pixels that fall into the j^{th} colour bin.

2.1 Back Projection

Swain [9] developed a method of comparing colour images using their colour histograms and a similarity measure. They used this to index images in a database. Database image retieval using colour indexing was a very sucessful and much research into improving this idea followed [4, 7, 3, 8, 5 and 6]. In Swain's article on colour indexing a second method, histogram back projection, was also presented. Histogram back projection is a method in which a target object can be located in an image. A ratio histogram, R, is constructed using the colour histogram of the target object, $H(T)$, and the colour histogram of the image, $H(I)$.

$$r_i = \min\left(\frac{t_i}{i_i}, 1 \right)$$

where r_i is the i^{th} bin of R, t_i the i^{th} bin of $H(T)$, i_i the i^{th} bin of $H(I)$.

If the value is greater than one it means that, this colour appeared more often in the target than in the image. The pixel values in the image are then replaced by the value of the colour bin of R into which they fall. The back projected image is then smoothed, to reduce noise. The peak is found and taken to be the location of the target object.

3 Single Colour Space Colour Profiling

The back-projection algorithm described above locates a known object in an image. The problem described here is to test for the presence of a not yet seen flaw, whose exact colour composition, shape, and size, can not be predicted. Building a ratio histogram with a typical flaw as the target object and a typical non-flawed filter as the background requires that that the flaws have a consistence colouration, this is not the case. A new type of histogram, in which all the colours that appear in any of a set of flaws receive a high weight, colours that appear in non-flawed images a low weight and colour never seen before a special flag weight. The array in which these weights are stored is called the colour profile.

3.1   The Training Phase

In the training phase a set of known images are presented. Each image in the training set is already classified as containing a flaw or not containing a flaw, although the location and size of the flaw is not known.   To record the information in the training set two histograms, called the training histograms, are constructed according to the following rules:

1.   Initialise all bins of both training histograms to zero. Use one for flawed images and one for non-flawed images.
2.   For each image in the training set construct a colour histogram. Then convert this to a binary colour histogram, B, by replacing the values of all non-zero bins with the value one. B is an array with ones in bins that are occupied by at least one pixel and zeros in bins that are never occupied.
3.   If the training image is flawed add B to the flawed training histogram and if the training image is non-flawed add B to the non-flawed training histogram.

To clarify: after all the training images have been added, a bin in the flaw training histogram contains the frequency with which that bin was occupied in all the flawed images and similarly, a bin in the non-flaw training histogram contains the frequency with which that bin was occupied in all the non-flawed images.

A colour histogram is not only a measure which colours are present in an image, but also the area each colour covers. This is very useful when comparing two particular images, but this is not the objective here; all that is wanted is a list of colours present in each image. This is exactly the information contained in a binary colour histogram.

3.2   Constructing the Colour profile

The training histograms are combined to construct the colour profile, which is then used to test unknown images for the presence of a flaw.   The two training histograms are normalised, then each bin in the colour profile is assigned a weight proportional to how much more often that bin was occupied in flawed versus non-flawed images. The weights are constructed in such a way that the higher the weight assigned to a colour profile bin the more likely it is that a pixel falling into that bin came specifically from a flaw, not just a flawed image. The profile bin weight zero indicates that pixels falling into this bin are not flaw pixels and is assigned in two instances. Firstly, if the bin was occupied in at least 1/5 of the non-flawed images, regardless of how often it was seen in the flawed images. Secondly, if the bin was occupied in more of the non-flawed images than the flawed ones. The first of these two cases reduces the false positive rate. The weight of 100 indicates that pixels falling into this bin certainly came from flaws. If a bin has only been seen in flawed images, that is, if it's non-flawed training histogram is zero but it's flawed training histogram is non-zero, then it is assigned the value of 100. Bins, where the flawed training histogram bin is greater than the non-flawed training histogram bin, are assigned values

between 0 and 100. If a bin has been occupied neither in the flawed images or the non-flawed images it is assigned the value 101, indicating it has never been seen.

To initialise the profile an array is created and each bin is assigned a value 101. The profile is then constructed by combining these histograms using the following rules:

$$P(i) = \begin{cases} 0 & if \quad NF[i] > (1/5)TNF \\ 0 & if \quad NF(i) > F(i) \\ 100 & if \quad NF(i) = 0 \quad and \quad F(i) > 0 \\ 101 & if \quad F(i) = NF(i) = 0 \\ f(i) & if \quad F(i) > NF(i) \end{cases}$$

where $f(i) = 100 * (F(i) - NF(i)) / (F(i) + NF(i))$,

TNF is the total number of non-flawed images, $P(i)$ is the i^{th} bin in the profile histogram, $F(i)$ is the i^{th} bin in the flaw training histogram and $NF(i)$ the i^{th} bin in the non-flaw training histogram.

3.3   Testing an Unknown Image

An unknown image is tested and classified as flawed or non-flawed by the following method:

1.   Replace each 3-D pixel value with the 1-D value of the colour profile bin into which the pixel falls.

2.   Smooth this transformed image and threshold on the smoothed value of each pixel and assign it as a flaw or not. The weight above which a pixel is classified as a flaw pixel is called the pixel threshold.

3.   If more than a threshold of flaw pixels are found in the image classify this image flawed. The number of pixels above which an image is classified as a flawed image is called the image threshold. In testing for flaws in the filter casings this was always set to be one, as any flaw detected, no matter how small was sufficient to reject a component. In a different application only bigger flaws might be sought and this threshold would be set appropriately.

This Image shows a flaw before and after detection with the flaw pixels found marked.

The colour profile had 16x16x16 sized bins.

Flaw Pixels

Before detection   After detection
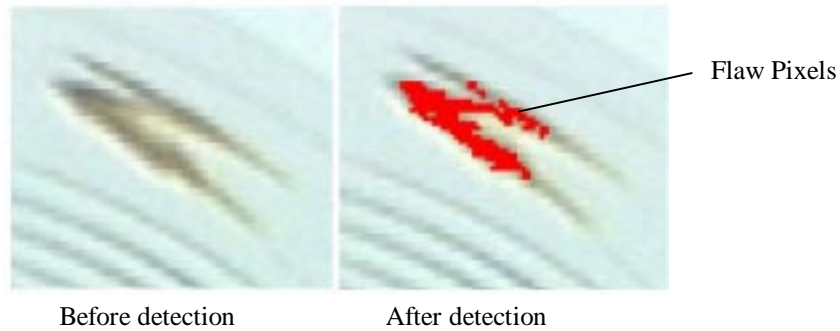
4   Discussion

4.1   Varying the Bin Size

The biggest decision to be made when implementing this method is deciding how to set up the bins in the colour profile. Which bin volume maximises accurate classification of pixels? To answer this question the bin size was constructed dynamically. In the training session the training histograms were constructed with small bin volume. Each colour channel was divided up into 64 intervals of length 4 each. The initial colour profile was constructed with this small bin size and then it was condensed. Condensing is a process in which sets of 8 adjacent bins merge and are thereby replaced by a bigger bin whose weight is an average of the 8 sub-bins. First the colour profile is scanned and the variance of each set of 8 adjacent bins is calculated. If any set have a variance that exceeds a threshold value the profile is not condensed, otherwise it is. This is repeated until it is no longer possible to condense without violating the variance-threshold. In the condensing phase the colour profile is replaced with a new colour profile 8 times smaller constructed using the following rules.

- If all the 8 sub-bins have the value 101 assign the new bin the value 101.
- Calculate the average of the sub-bins, leaving out any bins with the value 101.

This creates a colour profile with a bin size that maximises the finding of a flaw pixel without falsely classifying a non-flaw pixel as a flaw pixel. A trade off is always present: the bin size should be small enough that non-flaw pixels do not land in the same bin as flaw pixels, thus reducing the weights in those bins, but large enough that flaw pixels not yet seen, but similar to ones already seen, fall into the same bin. The highest accuracy occurred when the profile condensed down to the level that the bins were sized 16x16x16. Experiments showed that if the colour profile did not condense down to this size it gave very inaccurate results, including false positives and false negatives. If the training set was too small the colour profile could not condense to 16x16x16 without violating the variance-threshold, but as more images where added, it could achieve this. After the training set

contained enough images to achieve 16x16x16 adding more images did not allow it to condense more. The colour profile reached it's most stabile and accurate level at 16x16x16 and so this bin size is used.

4.2   The Colour Representation

Digital cameras use the RGB space so to use other colour spaces the RGB image must be transformed. This is computationally expensive and to justify this expense an improvement in results must be apparent. It has been observed that constructing colour histograms using RGB space "gives perceptually redundant bins and perceptual holes" [5].

The opponent colour space is defined through the following transformations [1]

$$rg = r - g$$
$$by = 2*b - r - g$$
$$wb = r + g + b$$

and was used by Swain [9] because the intensity information was contained in one axis, the wb-axis, making it possible to use a courser quantisation along this axis, reducing the sensitivity of colour histograms to changes in lighting. In the case of industrial inspection, lighting can be carefully controlled, so this is not an issue that needs to be addressed. The system was trained and tested using several different colour representations; RGB, Opponent, and HSV. In the initial testing [2], there was a casing with a very small flaw that RGB missed and HSV found and also one casing where the exact opposite occurred. This gave rise to the idea to create a colour profile that combines the information contained in both the RGB and HSV colour profiles.

5   Combination Colour Space Colour Profiling

The HSV and the RGB colour profiles obviously do not contain the same information, if one can miss flaws found by the other. Considerable extra computation is required to test each image using both profiles. To calculate the extra computation consider the transformation from RGB to HSV colour space:

$$V = \max(R, G, B)$$

$$S = 255 \; \frac{V - \min(R, G, B)}{V}$$

$$H = \begin{cases} 1 + \dfrac{G-B}{V - \min(R,G,B)} & for \quad V = R \\[2em] 3 + \dfrac{B-R}{V - \min(R,G,B)} & for \quad V = G \\[2em] 5 + \dfrac{R-G}{V - \min(R,G,B)} & for \quad V = B \end{cases}$$

This works out to be 13 adds and 3 multiplications per pixel. To calculate the pixel bin number takes 5 multiplications and 2 adds. It would be ideal if the information contained in the HSV profile, but missing the RGB profile, could be added to the RGB profile. In that way images could be tested in their raw RGB form with out having to be transformed to the HSV colour space. The goal is to find all the flaws that either the HSV or RGB profiles alone are able to find with this new profile. From now on this improved profile will be referred to as the multi-space RGB profile, MS-RGB profile, and the normal RGB or HSV profiles will be referred to as the single-space profiles, SS-RGB and SS-HSV.

The construction of the MS-RGB profile requires the information from both the SS-RGB profile and the SS-HSV profile to be combined. The MS-RGB profile will be RGB based and therefore initialised with the SS-RGB profile, except that all the *101*-bins are converted to negative 1.

First all the non-zero and non-*101*-bins in the SS-HSV profile are used to increase the bin values in the MS-RGB profile and then the zero bins from the SS-HSV profile are used to reduce the bin values in the MS-RGB profile.

```
for (every bin of the SS-HSV Profile with a value between 1 and 100)
{     for (every triplet that falls into the HSV bin)
    {   convert the HSV triplet to an RGB triplet
        j = the RGB triplet's bin number
        if (SS-RGB[j] == 101)
            MS-RGB[j] = HSV bin value
        else if (SS-RGB[j]  != 0 )
            MS-RGB[j] = max ( HSV bin value and the MS-RGB[j])
        else if (SS-RGB[j] == 0 and the HSV bin value is 100 )
            MS-RGB[j]++;              // this is being used as a counter }}
```

```
for (every bin of the HSV Profile with the value zero)
{    for (every triplet that falls into the HS bin )
    {     convert the HSV triplet to an RGB triplet
         j = the RGB triplet's bin number
         if (SS-RGB[j] == 101)
             MS-RGB[j]  = 0;              // The HSV bin value
         else if (MS-RGB[j] != 0  and the SS-RGB[j] == 0)
             MS-RGB[j]--;                 // this is being used as a counter }}
```

The MS-RGB profile is then compared to the SS-RGB profile. Any bin that has a zero value in the SS-RGB profile but a non zero value in the MS-RGB profile indicated this bin was being used as a counter in the MS-RGB profile. It contains a count of how many more times the bin received a vote to be increased than decreased. As it still has a non-zero value it means that it was more often voted to be increased than decreased. Bins in this situation are assigned a weight of one less than the pixel threshold value. This means that they cannot alone get a pixel labelled as a flaw pixel but since the image is smoothed can work in combination with other flaw pixels. In the last step all the MS-RGB profile bin values of -1 are replaced with 101.

In a training set of 40 flawed and 40 non-flawed components a SS-RGB and a SS-HSV profile were constructed. Then the MS-RGB profile was constructed using the above method. The SS-RGB profile had 3,829 *101*-bins, 51 zero bins and 216 other bins. The MS-RGB profile had 3,666 *101*-bins, 39 zero bins and 391 other bins. Two *101*-bins where set to zero, 161 *101*-bins were set to values between 1 and 100, 14 zero bins were increased and 22 other bins where increased. Images can now be tested for the presence of flaws without being converted to HSV by using the MS-RGB profile.
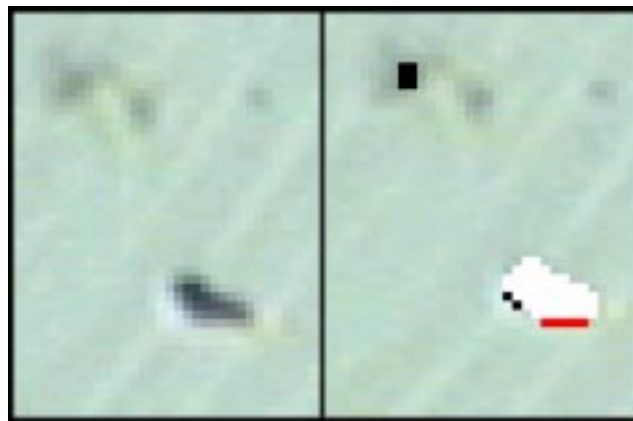
White pixels were found by both the SS-RGB and the MS-RGB profiles. The black pixels where found only by the MS-RGB profile. The red (or grey) pixels were

found only by the SS-HSV profile but missed by the MS-RGB profile. As the above image makes clear not every pixel that would have been given a non zero weight by the SS-HSV profile will also be given a non zero weight by the MS-RGB profile. This is due to the fact that the SS-HSV profile does not have unlimited power to change zero bins of the SS-RGB profile. This means that some information is lost however this is balanced by the fact that less false positives are introduced.

Clear Top Filter Cases	SS-RGB		SS-HSV		MS-RGB	
Correctly Classified Flawed	90	93%	91	94%	95	98%
Misclassified Flawed	7	7%	6	6%	2	2%
Correctly Classified Small Flawed	11	32%	11	32%	16	47%
Misclassified Small Flawed	23	68%	23	68%	18	53%
Correctly Classified Non-flawed	57	100%	57	100%	55	96%
Misclassified Non-flawed	0	0%	0	0%	2	4%

■ Training Set of 80 (40/40), Pixel threshold = 33.

6 Results and Conclusion

With non-optimised code on a P166 running NT the system is able to test a component in 0.7 seconds using the MS-RGB profile.

Flaws whose largest dimension does not exceed 0.2 mm (4 pixels) are called small flaws. Flaws below that size are very unlikely to be detected by humans, and are therefore not considered real flaws. On the other hand casings with small flaws also cannot be viewed as non-flawed. To avoid this ambiguity the small flawed casings have been tested in a separate category. The small flaws figures are an artefact of the resolution of the system therefore not a measure of the technique.

In a total of 90 components with flaw pixels found by the single space profile 10 had no pixels found by the MS-RGB profile or the SS-HSV profile that weren't already found by the SS-RGB profile. Of the remaining 80, 14 had more pixels found by the SS-HSV profile that by the MS-RGB profile and 66 more by the MS-RGB profile than the SS-HSV profile. In several images that contained two or more flaws one of the flaws was found only by the MS-RGB profile. The MS-RGB profile has shown an improvement over the SS-RGB profile in finding flaws, especially the small flaws. It does however have a higher false positive rate.

A limitation of colour profiling is that it does not find every flaw pixel. It is only testing for the presence of a flaw. If the area of the flaw is needed this method could be combined with a region growing technique that would segment out the flaw with a much greater accuracy than colour profiling alone. In developing this system much consideration

was given to the final application. Ideally the system would not only accurately classify components as flawed or non-flawed in real time but also be easy to use. This system does not have to be reprogrammed for different products to be inspected; it just has to be trained. It is also possible to view the problem the other way round and not test negatively for the presence of a flaw but positively for the presence of a potentially missing part of an object. Training is not only easy, but if the non-flawed training set is presented before the flawed training set, the system is able to predict how successfully it can classify future unknown images. To do this it estimates how many of the flawed images in the training set it can find flaws in. This is done in the following way. After all the non-flawed images have been added, the non-flawed training histogram is not altered again. The binary colour histogram of each flawed image can be compared to the non-flawed training histogram to see if it has filled any bins that are empty in the non-flawed training histogram. If this is the case the colour profile will have a bin value of 100 and will be filled by at least one pixel from this image. It is reasonable to expect that the flaw in this image would be found. The percent of flaw images that can be expected to be found is recorded and given to the user as an estimate of how well the colour profile will be able to label pixels in future unknown images.

This system could easily be applied in many automated quality control problems in which flaws can be located by their colours with no change to the software. It is accurate, real time and easy to use.

Acknowledgements

References

[1] Ballard, D.H. and Brown, C.M. Computer Vision. *Prentice Hall*: New York, 1982.

[2] Duffy, Nicola and Lacey, Gerard. Colour Profiling. *Irish Machine Vision and Image Processing Conference (IMVIP-97)*, September 1997.

[3] Funt, B.V. and Finlayson, G.D. Color Constant Color Indexing. *IEEE Transactions on Pattern Analysis and Machine Intelligence.* 17:5, May 1995.

[4] Hafner, J., Sawhney H.S., Equitz, W., Flickner, M. and Niblack, W. Efficient Color Histogram Indexing for Quadratic Form distance Functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence.* 17:7. July 1995.

[5] Smith, J. R. and Chang, S.F. Automated Image Retrieval Using Colour and Texture *Columbia University Technical Report #414-95-20,* 1995.

[6]   Smith, J. R. and Chang, S.F.   Single Colour Extraction and Image Query *International Conference on Image Processing (ICIP-95),* Washington, DC, October 1995.

[7]   Stricker, M. and Orengo, M.   Similarity of Colour Images *SPIE Proceedings: Storage and Retrieval for Image and Video Databases III,* 381-392, 1995.

[8]   Stricker, M. & Swain, M.   The Capacity and Sensitivity of colour Histogram Indexing *Technical Report, Dept of Computer Science, University of Chicago, Number TR-94-05,* March 1994.

[9]   Swain, M. J. & Ballard, D. (1991).   Colour Indexing *International Journal of Computer Vision.* 7:1, 11-32.