Pre-processing Colour Images with a Self-Organising Map: Baking Curve Identification and Bake Image Segmentation

Leonard G. C. Hamey[1,2] Len.Hamey@mq.edu.au
Jeffrey C.-H. Yeh[1,2] Jeffrey.Yeh@mq.edu.au
Tas Westcott[3*]
Samuel K. Y. Sung[2]

[1]Cooperative Research Centre for International Food Manufacture and Packaging Science, PO Box 218, Hawthorn VIC 3122, Australia.
[2]Department of Computing, Macquarie University, Sydney NSW 2109, Australia.
[3]Westcott Consultants Pty Ltd, PO Box 334, Greenacre NSW 2190, Australia.

Abstract

Kohonen's self-organising map is used to identify the colour development of baked goods from samples taken during baking. The resulting bake curves represent the colours characteristic of a particular baked product. Images of baked goods can be segmented and foreign bodies identified using these baking curves.

1. Introduction

Kohonen's self-organising map (SOM) is a powerful method for pre-processing data to reduce its dimensionality. When presented with data samples in a high-dimensional space, the SOM nodes organise themselves according to the structure of the data, capturing topological and density features of the data in the node locations [5, 6, 7]. The present paper shows that the SOM is an effective method for pre-processing colour images of baked goods.

For many baked goods such as biscuits (crackers and cookies) the surface colour is non-uniform as a result of uneven heating or blister formation. Colour changes throughout the baking process. Extremes of baking include raw dough and over baked product [4, 15]. The colour development of baked goods follows a curve, unique to each product, which we call the "baking curve" for that product. Samples of product have surface colours which are distributed along this curve in a manner reflecting the degree of bake of the product sample. The SOM can be used to effectively recover the structure of this colour distribution. The information extracted by the SOM provides for a simple and effective segmentation algorithm, and forms the basis for bake quality assessment. These techniques could also be applied directly to other image processing tasks dealing with objects with surface colour variation or colour development over time.

Visual appeal of food is important. Food with unacceptable colour is likely to be rejected even if it has good flavour and texture [9]. Colour change in food occurs during storage [2, 8, 12], ripening of fruits [3] and cooking [10, 13]. The techniques developed below may be applied to the study of such colour change processes.

Kohonen's SOM [6] has previously been employed in a number of pattern recognition tasks [7] including character recognition [1] and estimation of heart wall location in tomographic images [11]. The present paper demonstrates that the SOM neural network can be effectively employed to characterise the colour development of baked goods. Our results show that the self-organising map recovers the structure of the baking curve. The recovered structure can then be used for image segmentation suitable as a basis for product shape analysis and foreign body detection. We have previously reported the use of the self-organising map as part of a system for bake colour quality assessment [4, 14, 15].

2. Imaging Environment

Biscuit samples were imaged using a 3-chip CCD colour video camera (JVC KY-F55BE) connected to a DT-2871 colour frame capture board in a PC-compatible computer. The images were calibrated with reference to laboratory standards of 2% and 99% reflectance (see figure 1) included in each image. The camera was operated with gamma correction disabled. A small residual non-linearity in the camera response was corrected as part of the image calibration.

*Formerly, Research Supervisor, Arnott's Research Centre, Arnott's Biscuits Ltd, Homebush NSW 2140, Australia.

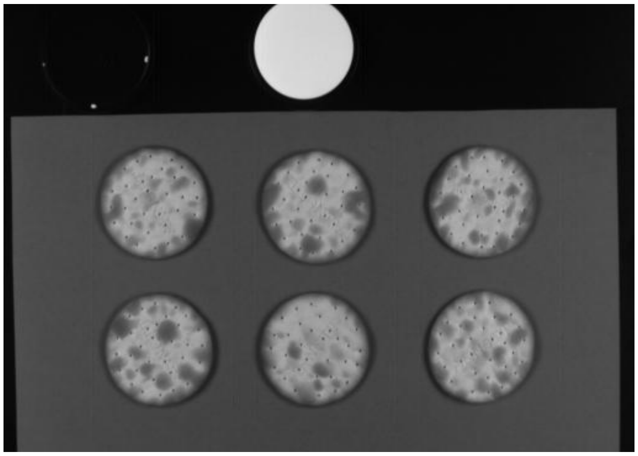

Figure 1. Image of samples of Water Cracker final product with calibration targets. The 2% target is in the top-left corner, identified by peripheral white dots.

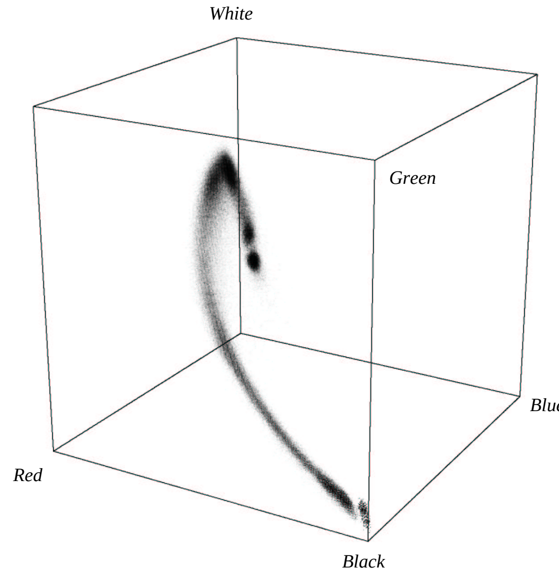

Figure 2. Baking curve for Water Cracker in RGB colour space.

The samples were illuminated with daylight-balanced fluorescent lamps, placed so as to avoid specular reflection off the biscuit surface. The image calibration procedure also corrected for illumination variations over the imaging area. Custom software was developed for image capture which enabled 16 images to be averaged together to reduce image noise levels.

3. Baking Curves

We first investigated the colour development of commercial biscuit product samples taken during baking. The biscuits studied are produced in gas-fired travelling tunnel ovens which may have between five and seven distinct oven zones. Each zone has an individual temperature control and the overall temperature profile of the oven determines the baking regime for the product.

Samples of product taken from each oven zone were digitally imaged using the set-up previously described. Raw dough samples, final product and deliberately over baked samples were also included to completely characterise the colour development process. The samples were imaged against a "sky blue" cardboard background to facilitate image segmentation. A typical image is presented in figure 1. The image contains the calibration reference standards at the top and six final product samples of a plain cracker.

After linear calibration and segmentation of the product images, we examine the colour distribution in the camera's red, green and blue (RGB) colour space. The non-background pixels in the collection of images obtained for a single product are projected into the 3-D colour space. Custom software renders each pixel as a partially opaque black point, creating an effect like a smoke cloud rendering of the colour distribution—darker plot regions represent a greater density of data. Figure 2 displays the resulting colour distribution for samples of Water Cracker. It is readily apparent that the distribution falls along a distinct curve in colour space which we call the baking curve. This curve displays the colour changes that occur during baking as a result of changes in moisture and chemical reactions, especially the non-enzymic browning (Maillard) reaction. Other products exhibit similar baking curves.

The baking curve shown in figure 2 represents the entire baking process. At any point in the process, the colours observed in biscuit samples occupy a distribution along the baking curve. In order to visualise the distributions that arise as baking takes place, we plotted outlines of the distributions for each stage of baking. The outlines include 80% of the biscuit pixels for each baking stage. The resulting plot, shown in figure 3, demonstrates the colour progression from raw dough through to over baking.

4. Recovering the Baking Curve

A 1-dimensional SOM is trained to identify the baking curve. The training data consists of image pixels. Only images of product samples are required. Since the biscuits are imaged against a blue background, a simple colour segmentation is used to remove the background to produce the

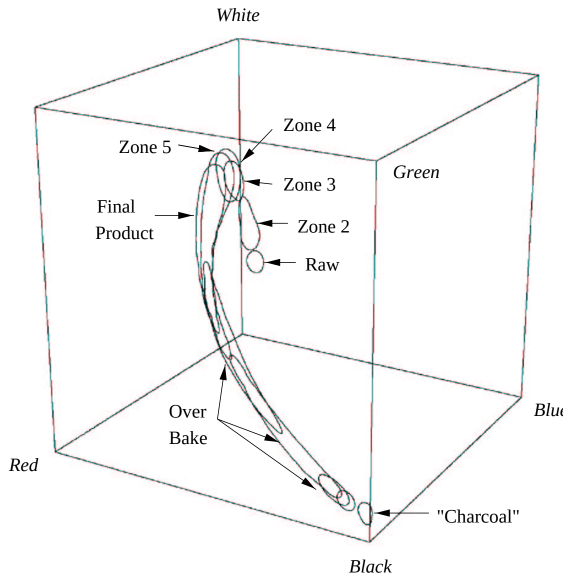

Figure 3. Colour development of Water Cracker during baking and over baking.

training data. This initial segmentation is not to be confused with the segmentation later performed on the basis of the trained SOM: that segmentation does not require any particular background colour as is shown in the next section of this paper.

The distribution of data in a typical baking curve has a dense central core of biscuit colour pixels with noise producing a surrounding scatter. A baking line can be constructed through the centre of the baking curve to represent the true path of the colour development during bake. This can be achieved by utilising the properties of a one-dimensional SOM.

A SOM consists of a layer of input nodes and a second layer of Kohonen nodes, with full connection between the layers. In our case, the input layer has three nodes—one for each of the red, green and blue colour bands. Each pixel of a biscuit is presented to the SOM as a distinct training pattern, so only a few samples are required for each degree of baking. The collected biscuit pixels are presented to the SOM in random sequence to prevent the SOM being unduly influenced by a particular biscuit.

An adaptive weight is assigned to each connection between an input node and a Kohonen node. The set of weights from the input nodes to a particular Kohonen node form the weight vector for that Kohonen node. The Euclidean distance between an input vector and a weight vector determines the similarity between the input and the node. For each input vector, the winning Kohonen node is the node with the smallest Euclidean distance to the input vector. The wining node and its neighbouring nodes are updated according to Kohonen's learning rule:

$$W_i(t+1) = W_i(t) + N_{iw}(t)\,[I(t) - W_I(t)]$$

where $W_i(t)$ is the weight vector of Kohonen node i at time t and $I(t)$ is the input vector at time t. Thus, the Kohonen node's weight vector is moved toward the input vector by an amount controlled by the neighbourhood function $N_{iw}(t)$ where w represents the winning node. For our experiments, N was the bubble neighbourhood kernel which has value $\alpha(t)$ for $|i-w| \leq T(t)$ and value zero elsewhere. Here, $T(t)$ is a user-specified neighbourhood distance threshold and $\alpha(t)$ is a user-specified learning rate. Both of these parameters may be varied over time.

A SOM is characterised by its topology and density preserving features. The topology preserving feature ensures that when the trained SOM nodes are projected onto the RGB colour cube, the topology, or baking line, created by connecting neighbouring nodes closely resembles the shape of the baking curve. The density preserving feature ensures that the distribution of trained SOM nodes matches the distribution of the pixels in the RGB colour space. Therefore, a SOM with a limited number of nodes will position nodes in the dense centre of the baking curve. This ensures the one-dimensional baking line created lies along the centre of the baking curve and hence correctly represents the baking curve.

Figure 4(a) displays the fit of a SOM with 30 nodes to the water cracker baking curve of figure 2. Because of the U-shaped curve at the top of the baking curve, the SOM has found a sub-optimal fit in which the end of the baking curve has not been correctly identified. A similar problem was observed by Manhaeghe et al [11] who found that an open-ended SOM model of the heart often failed to correctly match the open end in the tomography data. For this reason, they preferred a closed model which more reliably fitted the heart data. A similar technique could be applied here.

In our case, the U-shaped curve is a result of including raw product and early oven zones in the baking curve. Since our application interest is in final baked product colour, the fit of an open SOM curve is reliably facilitated by the exclusion of radically under baked product (figure 4(b)).

Two further problems confront the use of the SOM in this application: the selection of the number of nodes and the training duration. We found that a SOM with insufficient nodes will not cover the ends of the baking curve because of the low density of pixels representing extremes of bake. On the other hand, a SOM with excessive nodes will start to position the nodes amongst the noise scattered around the baking curve. We experimented with different values for these parameters, and selected the values that gave the longest smooth SOMs.

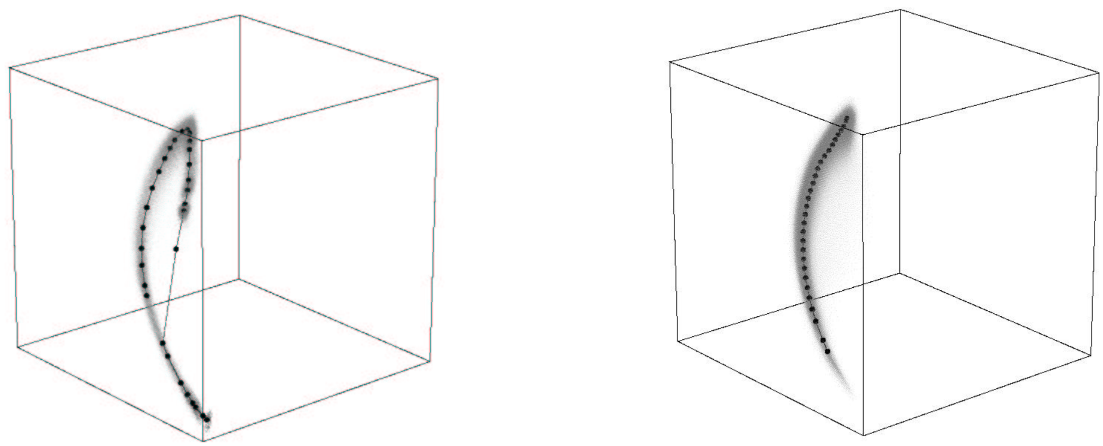

(a) Including early oven zones

(b) Final product only

Figure 4. SOM fit to Water Cracker baking curve.

5. Image Segmentation

The baking curve produced by the SOM consists of a sequence of points in RGB colour space. These points represent the core of the colour development curve with the biscuit pixels scattered around the curve. The basis of our segmentation technique is to determine the distance of each image pixel from the (linearly interpolated) baking curve. Only pixels that are close to the baking curve are likely to be biscuit pixels. Rather than making a binary segmentation decision, we assign to each pixel a weight which is a Gaussian function of the pixel's shortest distance from the baking curve.

To determine the standard deviation of the Gaussian weight function, an estimate of the scatter variance is required. The scatter is modelled as a circular Gaussian distribution with unknown variance σ^2, in two dimensions u and v locally perpendicular to the baking curve. Given a large sample from this scatter distribution, the mean square of u over the sample, $s_u^2 = \sum_i u_i^2/N$, estimates σ^2. Similarly, s_v^2, the mean square of v, also estimates σ^2. Neither u nor v is directly accessible, but the Euclidean distance d of each pixel from the baking curve is readily computed. Since the mean squared deviation of the scatter s_d^2 is given by $s_d^2 = s_u^2 + s_v^2$, it follows that s_d^2 estimates $2\sigma^2$. Thus, σ^2 may be estimated by half the mean square deviation of biscuit pixels from the baking line. This deviation is computed over the set of pixels in the SOM's training set and includes only biscuit pixels.

Figure 5 shows a colour input image of biscuits and distracters. Figure 6 shows the segmentation results achieved by our approach. The segmentation obtained is highly suitable for the task of bake quality assessment. It will be noted that the background and non-biscuit objects are completely segmented since their pixels are far from the baking curve, while the shadow features and distracters on the product sample (such as docker holes and embossing) have reduced segmentation value (figure 6(a)). This is because the shadowed pixels fall near the baking curve but not on the baking curve since they have reduced overall brightness. The application of simple morphological operators dilate and erode produces a final segmentation which is suitable for product shape analysis and detection of foreign bodies (figure 6(b)).

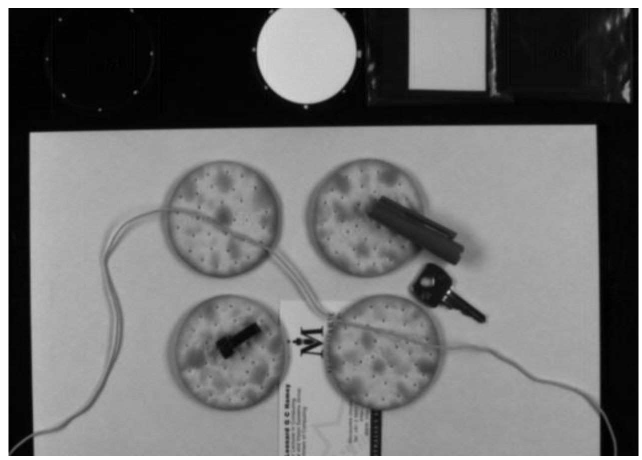

Figure 5. Original colour image with distracters.

(a) Baking curve segmentation

(b) Final segmentation

Figure 6. Baking curve based segmentation of a colour biscuit image.

6. Conclusion

The colour development of baked goods can be described by a baking curve in colour space. Kohonen's self-organising map is suitable for identifying baking curves of baked goods. An estimate of the noise distribution around the baking curve facilitates image segmentation into biscuit and non-biscuit areas. This provides for product shape analysis and identification of foreign bodies. The techniques developed have application to other food and non-food products.

7. Acknowledgements

This research was supported and partly funded by Arnott's Biscuits Ltd and the Australian Research Council. The authors wish to acknowledge the help and support of Arnott's personnel.

References

[1] Z. Chi, J. Wu, and H. Yan. Handwritten numeral recognition using self-organising maps and fuzzy rules. *Pattern Recognition*, 28:59–66, 1995.

[2] H. Corke, L. Au-Yueng, and X. Chen. An automated system for the continuous measurement of time-dependent changes in noodle color. *Cereal Chemistry*, 74:356–358, 1997.

[3] R. Gómez, J. E. Pardo, R. Varón, and F. Navarro. Evolution of color during the ripening of selected varieties of paprika pepper (Capsicum annuum L.). *Journal of Agricultural and Food Chemistry*, 44:2049–2052, 1996.

[4] L. G. C. Hamey, J. C.-H. Yeh, and C. Ng. Objective bake assessment using image analysis and artificial intelligence. In *Cereals '97: Proceedings of the 47th Australian Cereal Chemistry Conference*, pages 180–184, Perth, Australia, 1997. Royal Australian Chemical Institute.

[5] R. Hecht-Nielsen. *Neurocomputing*. Addison-Wesley Publishers, 1990.

[6] T. Kohonen. *Self-Organization and Associative Memory*. Springer-Verlag, 1984.

[7] T. Kohonen. *Self-Organising Maps*. Springer, New York, second edition, 1995.

[8] W. C. Lin, J. W. Hall, and A. Klieber. Video imaging for quantifying cucumber fruit color. *HortTechnology*, 3:436–439, 1993.

[9] P. P. Ling, V. N. Ruzhitsky, A. N. Kapanidis, and T.-C. Lee. Measuring the color of food. *Chemtech*, 26:46–53, 1996.

[10] J. E. Lozano and A. Ibarz. Colour changes in concentrated fruit pulp during heating at high temperatures. *Journal of Food Engineering*, 31:365–373, 1997.

[11] C. Manhaeghe, I. Lemahieu, D. Voglaers, and F. Colardyn. Automatic initial estimation of the left ventricular myocardial midwall in emission tomograms using kohonen maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 64:259–266, 1994.

[12] E. J. Sacks and D. V. Shaw. Color change in fresh strawberry fruit of seven genotypes stored at 0c. *HortScience*, 28:209–210, 1993.

[13] M. G. Scanlon, R. Roller, G. Mazza, and M. K. Pritchard. Computerized video image analysis to quantify color of potato chips. *American Potato Journal*, 71:717–733, 1994.

[14] C. T. Westcott and L. G. C. Hamey. Data recognition system. Patent Specification PCT/AU95/00813, Arnott's Biscuits Limited, 1995.

[15] J. C.-H. Yeh, L. G. C. Hamey, C. T. Westcott, and S. K. Y. Sung. Colour bake inspection system using hybrid artificial neural networks. In *Proceedings of the IEEE International Conference on Neural Networks*, volume 1, pages 37–42, Perth, Australia, Nov. 1995. IEEE.