TILE SHADE CLASSIFICATION STRATEGIES FOR MAXIMUM PROFIT

Martin Coulthard,

Development Director
Surface Inspection Ltd, Great Britain

SUMMARY

This paper presents novel strategies that can be implemented in an automatic tile inspection system to carry out tile shade selection in such a way that achieves customer satisfaction and yet minimises the number of shades produced.

It reviews and analyses why tile shade classification is needed and what its objectives should be. It discusses the different production line scenarios and describes their effects on the sequence of shade variations in the inspected tiles. The requirements for effective shade measurements are listed and explained. The relevance of conventional colour measurement to tile shade measurement is discussed and additional measurements needed for tile shade are suggested. Strategies for establishing and managing shade classification schemes are then proposed and their potential results are compared with a real shade scheme created by eye. The importance of using automatic tile inspection systems is then highlighted.

OBJECTIVE

This paper aims to show that a thorough analysis of the requirements for tile shade classification leads to improved shade classification strategies which, when implemented in the software of an automatic tile inspection system, give improved customer satisfaction and reduced costs.

INTRODUCTION

The shade of a tile is the combination of a number of visual characteristics of the tile surface, including its colour and the distribution pattern of colour or decoration over the tile surface.

Sorting and classification of tiles by shade is a challenging and complicated subject. It is also a critical one for tile manufacturers.

The tile manufacturing process has many variables which can affect tile shade. These include:

- material colours, quantity and mixing;
- a wide range of printing and other decoration effects; and
- firing conditions.

The manufacturer is only able to achieve the ideal situation of a single shade for each product if he can control all of these process variables extremely well (or if his customer's expectations for shade are very low - i.e. it is a low quality, low cost product). Due to the nature of the process, with the use of natural materials, intentionally variable decoration effects and high temperature firing, it is usually not possible to guarantee the production is of a single shade. Even if it is, the investment in the necessary processing and quality control machinery may be too great. For this analysis, therefore, we will assume that the control of the process cannot be improved economically and that the amount of shade variation that occurs as a consequence is fixed.

Whenever more than one tile shade is produced shade classification is necessary and a series of problems then occur:

- quality/inspection managers find it hard to establish and control shade groups;
- manual inspectors do not give consistent, accurate shade decisions against the established shade groups;
- tiles within the same shade group have too large a visible difference (due either to incorrect inspection or to the defined shades being too large), leading to customer returns and damaging the manufacturer's reputation for quality and therefore his sales prospects;
- there are many different shades, giving high costs for storage, distribution and sales;
- during inspection many tiles do not appear to be within the current shades, and so either new shades have to be established or tiles have to be classified as non-first quality.

All of these problems incur additional costs for the manufacturer and reduce his profit.

What is needed is a rational method of establishing a shade classification system for a tile product and then applying it in a consistent way in order to avoid or minimise these problems and so increase profit. Most of the analysis that follows apply (at least in theory) to both manual and automatic shade classification.

TILE MANUFACTURER REQUIREMENTS FOR SHADE CLASSIFICATION

There are three principal requirements for shade classification:

CUSTOMER SATISFACTION

When a customer receives an order of tiles which are identified as having the same shade he wants to be able to install them on a floor or wall as they come out of the box.

The range and distribution of the shades of the tile should be such that no tile looks 'wrong', that is to say looks significantly different from the other tiles.

If there are tiles that look wrong then the customer may return the whole order and claim compensation.

A tile generally looks wrong if its appearance is distinctly different to that of the other tiles around it. If the tile type is of a constant, repeating design, such as silk-screen print, then quite small differences in the appearance look wrong. If, on the other hand, the tile has a random design, such as double-charge porcelain, then quite large shade variations are normal and only tiles with larger differences (or different types of differences) look wrong and are considered to be bad shade.

Preferably the manufacturer should not just avoid returns due to shade, but should exceed his customer's expectations and so enhance his reputation for quality, which could lead to higher sales volume or the ability to increase the selling price.

Another way to express the objective of achieving customer satisfaction is to say that the perceived difference between any 2 tiles of the shade should not be too large.

MINIMISE COSTS

Any steps the manufacturer takes to improve customer satisfaction need to be balanced against the costs of doing so.

One obvious step which should improve customer satisfaction is to sort into more, smaller shades. While this may make the customer happy it introduces considerable costs.

Different shades each need to be stacked, packed and palletised separately, and then the pallets of different shades need to be identified and stored. Any order from a customer has to be met by a single shade or, if the customer agrees, by a limited number of different, identified shades. During a particular production run the number of tiles of some shades may be too small to be saleable, and therefore need to be stored . In general, the more shades there are the greater the costs.

CLEAR, UNDERSTANDABLE AND PRACTICAL

A shade classification system that can, in theory, satisfy the customer and minimise the number of shades would be of no use unless it could be understood and used by the manufacturer's personnel in the factory environment. For this to be possible the shade system needs to be both clear and practical.

SUMMARY OF REQUIREMENTS

An optimal shade classification system should achieve a good balance between:

1. maintaining customer satisfaction;
2. minimising the number of different shades; and
3. being clear, understandable and practical

SCENARIO FOR SHADE SELECTION

There is no single best way to organise and carry out shade classification, as the optimum method depends upon a number of factors. The most crucial of these is whether or not the manufacturer needs to match shades from batch to batch.

FIXED SHADES OR NEW ONES EACH PRODUCTION RUN?

If he does need to match between batches then he must establish and manage a system of fixed shades for each product.

If he does not need to do batch to batch matching he can create new shades each production run for this product, and can afford to do so during the production run only if and when needed.

Each of these approaches has its own advantages:

Fixed shades:

- mean that the manufacturer can add tiles of a given shade from one production run to tiles of the same shade from a previous run, and so build up saleable quantities - less production is wasted due to unsaleable small quantities of tiles of a shade;
- may be preferred by customers as they can also use or re-sell tiles of the same shade together;
- allow more effective communication with customers on shade.

New shades each production run:

- means that fewer different shades are produced each run (if a good shade strategy is followed);
- no storage is required for small quantities of tiles of a given shade, as these are downgraded to a lower quality category or scrapped;
- small quantities of tiles of a given shade are less likely than with fixed shades (see later);
- is simpler and easier to administer.

The choice of which of these approaches to follow has to be made by the manufacturer by weighing up the economic costs of each of these factors. In reality each manufacturer already has an established approach and may be unwilling to change it.

ORGANISATION OF TILE INSPECTION ON PRODUCTION LINE

There are at least 3 ways tile inspection can be organised.

1. Direct from kiln. Some manufacturers convey tiles directly from the kiln past compensators or a short term buffer straight to the inspection area.

Tiles (except when coming from the compensator or buffer) arrive in the same order that they were produced and the inspectors / inspection machines can see the drift in shade produced by process changes, allowing them to make more rational decisions about the shade system and simplifies shade inspection. If we consider a particular shade property, such as the lightness of the tile, then it will normally tend to drift up and/or down gradually due to process variations. The inspectors will be able to observe these changes:

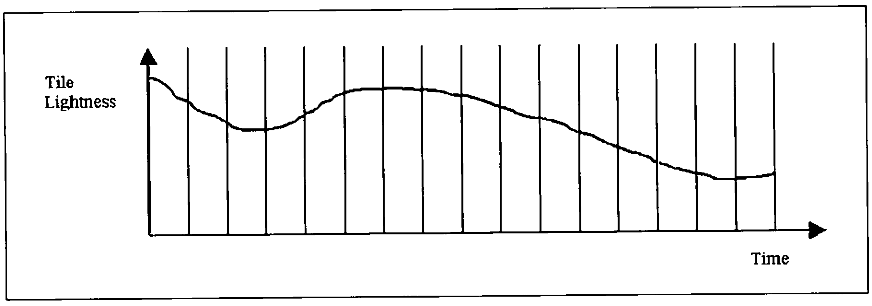

Figure 1.

If a process problem causes major shade jumps or many defects then the inspectors can notify production personnel and the problem can be fixed before a very large quantity of defective tiles are produced. This is only possible as the inspectors are seeing the tiles very shortly after they are produced.

The main disadvantages of inspecting from the kiln are that it requires 24 hour, 7 day inspection; that if there is a breakdown in the equipment in the inspection area there is nowhere to put the tiles leaving the kiln; and some problems may be experienced inspecting and packing tiles that are still hot from the kiln.

2. Storage Bed Between Kiln and Inspection. A common approach for some tile types such as uncoated porcelain tiles is to store up to 2 or 3 days worth of tiles stacked in piles on a flat storage bed after they leave the kiln. They are then unstacked and passed to the inspection area as required. This allows inspection to be carried out on just one or two shifts only on week days, as the amount of tiles can be allowed to build up when inspection is not taking place. It is not usually a practical approach for glazed tiles because stacking of the tiles can scratch the glaze, and can also be problematic with small tiles as the stacks of small tiles are less stable.

The stacks of tiles that are sent to the inspection area are those which were stacked first, but the tiles within each stack are unstacked and inspected in the reverse order to the production. If the shade of the tiles is changing this gives a confusing pattern of shade changes, with the changes during the time for each stack (perhaps 30 minutes) reversed:

While this pattern of shade changes can be confusing, staff can go to the store, pick off samples, see how the shade is varying and so plan how they are going to sort the shades.

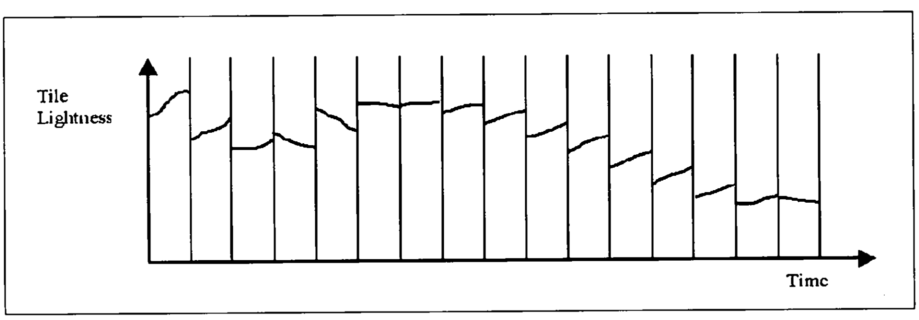

Figure 2.

A disadvantage is that there can be up to 3 days between the production and inspection of a tile, so it is too late for inspectors to identify process problems.

3. **Storage Wagons Between Kiln and Inspection.** The most common way of producing glazed wall tiles is for the tiles leaving the kilns to be loaded onto storage wagons. The wagons are moved to a storage area, and then later moved to the inspection area to be unloaded. Often, however, the wagons are not loaded and unloaded in the same order, so the shade changes seen by the inspectors can be very confusing, as the parts of the shade process curve are shuffled almost at random.

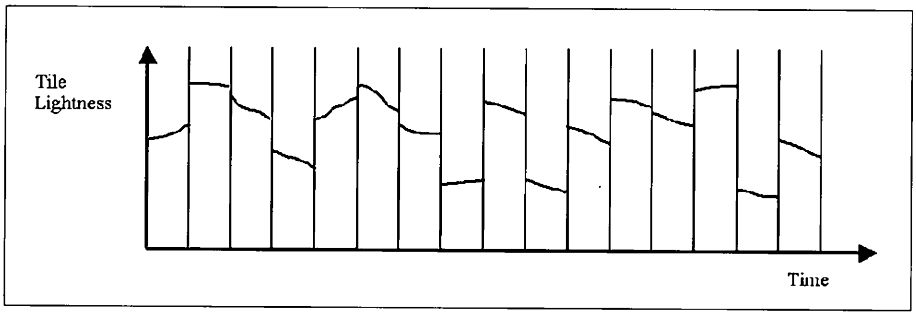

Figure 3.

Shade sorting under these conditions can be extremely difficult, and again feedback from inspection to production is too late to be useful in most cases.

SHADE MEASUREMENTS

Until now the discussions have been equally applicable to manual shade sorting and automatic shade classification. To take the analysis further, however, we need to start to consider how we can measure shade.

Appearance measurement instruments and systems take objective measurements which correspond with these descriptions. They illuminate the tile surface with one or more light sources of controlled intensity and colour, collect the light reflected from the tile surface, convert it to digital signals and analyse them using a digital processor.

To be useful for tile shade classification an **individual** shade measurement should:

1. measure an appearance variation that is clearly visible to the human eye, in a way that correlates with human perception;
2. give reliable, i.e. repeatable, measurement values;
3. measure shade variations which occur or could occur during the production of the tile - there is no need to measure something that cannot change;
4. be linear with human perception - a 1 unit difference at one part of the measurement scale should appear to a person to be just as large as a 1 unit difference at another part of the measurement scale (at least within the range of measurement values that occur);

Further to this the **set** of shade measurements selected for tile shade classification of a particular product should:

1. be able to detect and measure all of the shade variations that occur or could occur - there must be enough measurements;
2. be orthogonal to each other - a change in the appearance characteristic measured by one measurement should not affect any of the other measurements;
3. be isotropic - a 1 unit difference of one of the measurements should be perceived as being just as large as 1 unit difference of any of the other measurements.

COLOUR MEASUREMENTS

The requirements listed above may seem complex and unrealistic. Fortunately, however, we are able to build our tile shade classification system on an area of appearance measurement that has seen extensive research and development over more than 70 years - the science of **colour measurement.**

The work in this field has been directed at finding ways of measuring, specifying and communicating colour and colour differences. It has resulted in several well documented and understood colour measurement systems, defined in international standards.

The basis of colour measurement is an understanding of how people see and describe colours. We normally do so using three aspects of colour, which research has found can be used to define a colour:

Hue: this is whether the colour is red, yellow, green, blue or some other colour;

Saturation: how strong or vivid the colour is;

Lightness: this is how light or bright the colour is.

These 3 values are normally referred to as colour co-ordinates, and can be used to determine a 3 dimensional colour space in which any colour can be represented by a point. In this space lightness is treated as the vertical axis, the saturation is the distance of the colour point from the lightness axis, and the hue is the angle between the line in the horizontal plane (the colour plane) from the lightness axis to the colour point and a fixed reference axis in that plane.

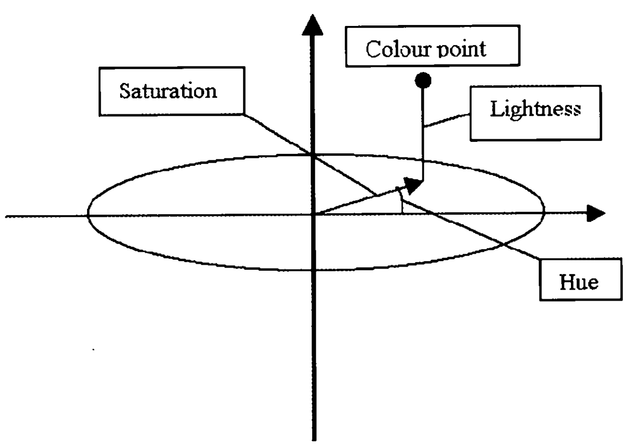

Figure 4.

While this system allows the colours to be defined it does not allow easy measurement of colour differences. Further developments in colour science aimed to overcome this difficulty and arrived at improved systems such as the L*a*b* colour measurement system which is probably the most widely used today.

In this system lightness (L*) is still the vertical axis, but two measurements a* and b* are defined as the linear axes, with a* defining the colour on a blue to red scale, and b* indicating a green to yellow scale.

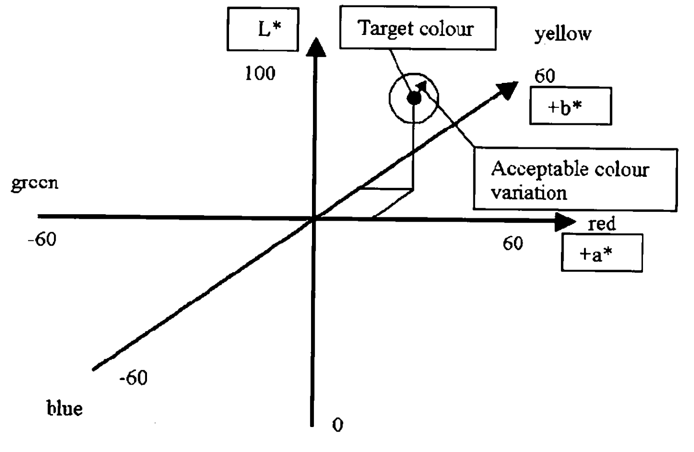

Figure 5.

Measurements a* and b* are in the horizontal colour plane and are perpendicular to each other, so the system has 3 perpendicular axes (i.e. the axes are orthogonal) and forms a Cartesian space. This L*a*b* space is also referred to as the CIELAB colour space.

Further to this the 3 measurements L*, a* and b* are defined in such a way that 1 unit on any of the scales has about the same visibility, making them both linear and isotropic. They therefore meet the requirements for tile shade measurements listed above. L* has a scale from 0 (black) to 100 (white) while a* and b* have scales from -60 to +60.

The reasons for the linear, orthogonal and isotropic requirements can now be seen. Let as consider two colours C1 and C2 with Lab colour co-ordinates (L1, a1, b1) and (L2, a2, b2) respectively. The magnitude of the colour difference ΔE is then the distance between the two points in the Lab colour space:

$$\Delta E = \sqrt{[(L1 - L2)^2 + (a1 - a2)^2 + (b1 - b2)^2]}$$

A colour difference (ΔE) of 1 unit has about the same visibility at any part of the colour space.

This system is widely used across many industries to specify colours and the acceptable variations of them. For instance a manufacturer might specify that the plastic components he is buying should have an Lab colour of (24.2, 33.2, -14.8) and that the tolerance or acceptable variation of this is a ΔE of 1.0 from these co-ordinates. This requirement can be considered to define a sphere in the Lab colour space, with its centre at co-ordinates (24.2, 33.2, -14.8) and a radius of 1.0. This sphere is drawn as a circle in Figure 5. Any point that is inside this sphere represents a colour that is acceptable.

As the measurement system is standardised both the manufacturer and his supplier now know what colour is required and can check, using standard instruments, whether or not the specification is being met.

This system can be applied to tiles and for uniform, plain coloured tiles it can give an effective system of shade control (though even on these tiles other shade characteristics such as gloss may need to be considered and measured).

However, in general, tiles are not uniform and plain. They have a wide range of decoration types, including silk-screen printing, rotary printing, spray and mixed materials to give speckles or blobs. While the average L*, a* and b* values for the tile surface are often useful shade measurements they are rarely sufficient to fully define the shade of a tile and a range of other measurements are needed.

OTHER SHADE MEASUREMENTS

In the tile industry the terms 'shade' and 'tone' are used to encompass a wide range of appearance characteristics. The measurements that are required to define a tile's shade depend on the tile type and in particular on its decoration. Some of the measurements are Lab **colour** values, but different sets of measurements may be needed for the base material/colour and for other decoration materials or print layers.

The **geometric distribution** of colours is critical for decorated tiles. A wide range of measurements can be used to define the position, contrast, size and shape of this decoration.

The **reflectivity** of the tile's surface - its gloss - is another key appearance characteristic. Variations in the shininess of the surface are likely to be noticed by customers. If one or more tiles in a batch have a much lower gloss than the others and all are laid on a floor then these tiles will look clearly different and so may cause returns. In some cases variations in the texture of the surface are also perceivable and should therefore be measured.

On **plain tiles** typical shade measurements are: L, a, b - average gloss - surface texture.

On **printed tiles** possible measurements include: L, a, b of base colour - weight/contrast of print - sharpness of edges of print - x and y registration of print - (sometimes it is also necessary to measure the colours of the print).

On tiles with a **random speckle pattern**: L, a, b of base material - average size of speckle - number of speckles per unit area.

An effective automatic tile inspection system needs to be able to take all of these measurements. Each measurement should be designed in such a way that it meets the list of requirements given in this paper.

SHADE SPACES

A tile's shade is defined by the values of its shade measurements. Unlike simple colour measurement, typically more than three shade measurements, and sometimes up to about 10, are needed to fully define a tile's shade. We can however extend the concepts of colour measurement to cover these additional measurements.

All of these shade values can still be considered to be shade co-ordinates that define the position of the tile's shade in n-dimensional shade space (where n is the number of shade measurements). If all of the measurements are linear, isotropic and orthogonal the shade difference between two tile shades is still equal to magnitude of the distance between the two shade points in this shade space, and this is still equal to the square root of the sum of the squares of the differences of the individual measurements.

As with colour measurement a single tile shade group can be defined by the co-ordinates of the central or target shade and the maximum difference from it. The shade is therefore a hyper-sphere in n-dimensional space. We can guarantee that no 2 shade point within this shade group are more than a given distance apart, equal to the diameter of the sphere.

If we believe that our customers will be satisfied with a shade if any two tiles within the shade group always have a shade difference of less than d then the largest volume of space that can satisfies this requirement is a hyper-sphere of diameter d.

We will now examine how these concepts can be used to establish a strategy for shade classification that is optimised to ensure this maximum shade difference within a shade group and also to minimise the number of shade groups required.

A STRATEGY FOR SHADE CLASSIFICATION WHEN WE CREATE NEW SHADES EACH PRODUCTION RUN

We can describe an optimum strategy for shade classification in the case when we wish to create new shades each production run, and we inspect the tiles in the same order that they were produced.

When this is so we can expect that, during a short period of time, the tiles produced will have an average shade and a limited amount of shade variation about this average shade. If we view a plot of the two shade measurements these points can be represented as a cluster of points, as shown.

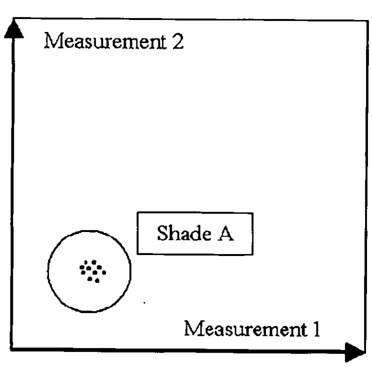

Figure 6.

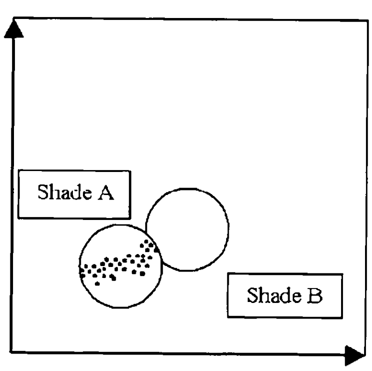

Figure 7.

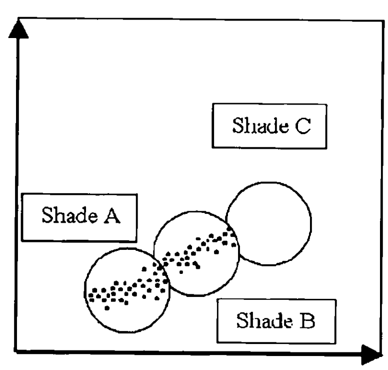

Figure 8.

A shade group in this 2 dimensional plot is the cross-section of the hyper-sphere, in other words a circle, of diameter d. We could fix the position of this first shade circle (shade A) at the start of the batch, for instance with its centre at the average shade of the tile at that time. If we did so then we would need to establish a new shade when the average shade had drifted only by a distance of d/2, the radius of the shade circle.

A new shade can actually be created later if we allow the shade circle to drift to cover the tile shade points until the average shade has moved a distance of about d. During this process outliers - tiles with shade quite different from the current production - are downgraded to second or reject quality. Only when it is no longer possible for the shade A circle to cover, say, 90 out of the last 100 tiles, as well as all of those previously classified as shade A, is it necessary to create shade B.

The degree of overlap of the shade B circle with that of shade A should be determined so that the normal variability of tile shade about its current mean is covered.

This sequence is then repeated. Shade B drifts around shade A to cover the changing shades until the current shade reaches its opposite side. As the mean shade drifts further it is no longer possible for shade B to cover the shade points of current tiles and shade C is established.

After the production run has continued for some time a zoomed out view of the shade plot might look like this.

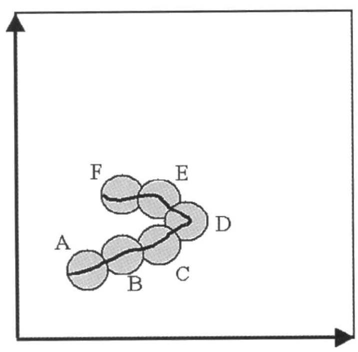

Figure 9.

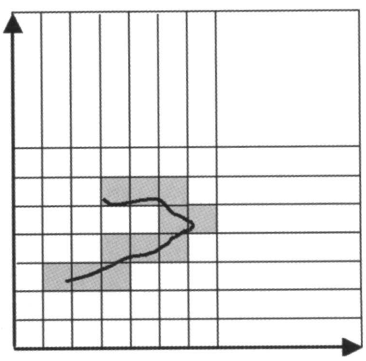

Figure 10.

There is a chain of slightly overlapping circles which follow the movement of the mean shade. Almost all the tiles have been classified into one of the six shades. No two tiles within each shade are more than a distance of d apart (i.e. none of them have a shade difference of more than d).

For comparison, we could have taken our shade group to be square, then it would have to be a square with a diagonal of d to ensure that no two tiles within the shade were more than d apart. The length of the side of the square is therefore $d / \sqrt{2}$. If a fixed grid of squares of this size is plotted over the same shade process curve we see that at least eleven shade groups would have been required - twice as many as with the optimum strategy. Using fixed square shade groups is clearly inferior.

The strategy of chaining together circular shades gives the minimum number of shades (at least until the shade curve doubles back onto itself). Moreover the strategy can be applied not only in two dimensions, as illustrated, but also in one, three, four or more by chaining together hyper-spheres in n-dimensional space.

A STRATEGY FOR SHADE CLASSIFICATION WITH FIXED SHADES

If the previous strategy is followed in order to establish a system of fixed shades to be used for many production runs then problems occur. The shade curve will inevitably pass through the same area of shade space many times and the spherical shade groups start to overlap. If the first established shades are left intact then the later shades which are filling in the space around the earlier shades become small and irregularly shaped. The arrangement of shades no longer minimises the number of shades required and is also disorganised and unclear, making it hard to manage.

The fundamental problem is the circles do not 'tile' together to cover a plane, spheres do not pack together to fill a three dimensional volume and, similarly, hyper-spheres do not pack together without leaving spaces between them. It is therefore necessary to find a shade shape that does tile / pack without gaps, and yet still covers as large amount of shade space as possible while keeping the maximum distance between two shades within the group below the threshold value.

In two dimensions the problem is fairly straightforward. As the shade space is isotropic the best shape will be symmetrical and have equal length sides. Only three of these tile together without gaps:

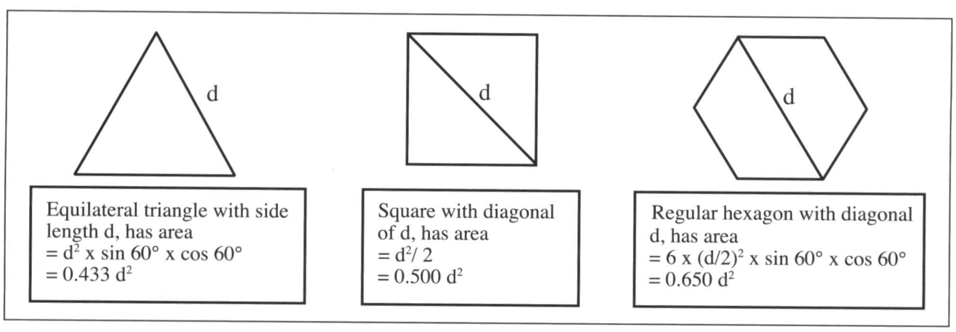

Figure 11.

It seems clear that the regular hexagon is the shape of shade group that both tiles together and gives the minimum number of shade groups (as it has the largest area for a given value of d).

In the general case where the central shade for a product is at point P, any shade group must not have tiles within it that are more than d apart, and to be acceptable as this product type the shade must not be more than r from shade P, the shade space can be tiled with hexagons as shown below. In this case 7 shade groups are required to cover almost all the shade space of the product.

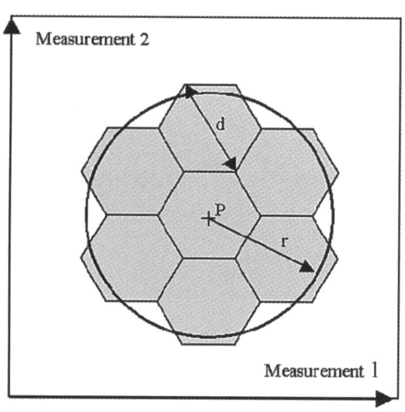

Figure 12.

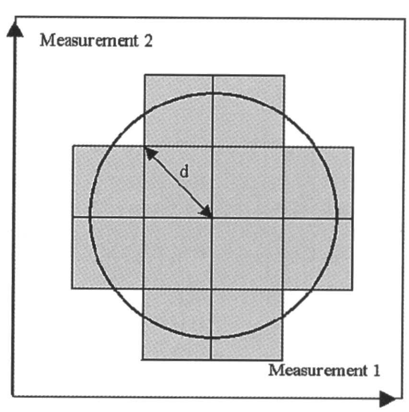

Figure 13.

If square shades, of diagonal d, were used then 12 shades would be needed to cover a similar proportion of the product shade space.

Typically 30% more shades would be required if the shade groups were square than if they are hexagonal. This can be calculated from the ratio of the areas: 0.650 / 0.500 = 1.3.

The situation is more complex in three and more dimensions, and it is unlikely that any regular shape more complex than a cube is able to pack together without gaps. This is an area for further research.

WHAT HAPPENS NOW WITH MANUAL INSPECTION?

When shade groups are established purely by eye it is impossible to create a shade scheme that even approaches the ones described above.

The graph below shows a scheme of twelve fixed shades that has been developed by one manufacturer, with great care and attention, over a period of several years. Each shade group is indicated by an ellipse. It is noticeable that:

- each of the shades is a different size and shape;
- the shades overlap; and
- there are gaps between the shade groups.

It is not possible to create an optimised shade scheme without the aid of an automatic inspection system.

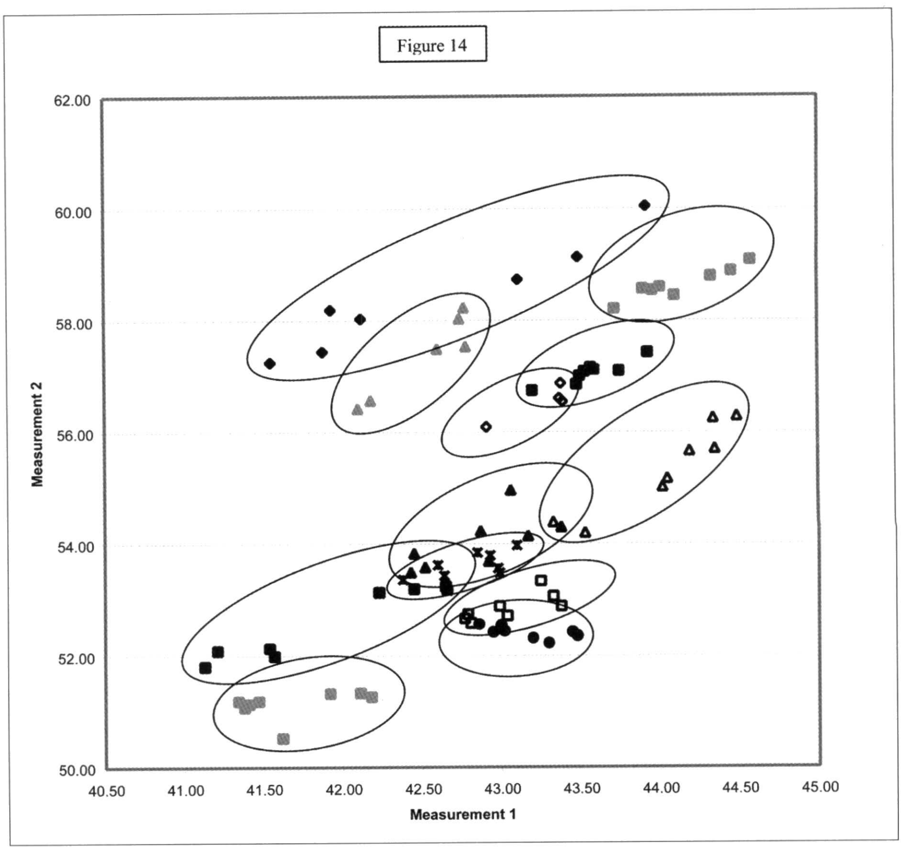

Figure 14.

USE OF AUTOMATIC INSPECTION SYSTEMS FOR SHADE CLASSIFICATION

Automatic tile inspection systems use electronic cameras and digital image processors to acquire and process digital images of tiles. They are able to calculate shade measurements and take these measurements in a way that is reliable and repeatable.

Most current systems, however, concentrate on learning a shade classification scheme that has been developed by human inspectors. They are able to learn the sizes and positions of shade groups by passing sample tiles through the system and then to apply this scheme to achieve automatic inspection. At best this only serves to mimic the manual scheme in a more repeatable, automatic way and it therefore has most of the shortcomings of a manual shade classification.

Far greater benefits can be achieved if the automatic inspection system is programmed to guide the manufacturer in setting up an optimal shade classification scheme. This is only possible, however, if the system supplier has a clear understanding of the requirements and theory of tile shade classification, and experience of what shade measurements are required for different tile types.

The system should guide the manufacturer's personnel through a sequence of steps in order to establish the shade scheme. The interface of the system needs to be clear and simple, using non-technical terms to describe shade differences and clear graphics to show progress in the sequence.

By using such a system the manufacturer can carry out shade classification with a good balance between the key objectives of customer satisfaction, minimising costs and having a clear, practical shade scheme. He is, therefore, able to reduce his costs and improve his profits.

CONCLUSIONS

A review of the need for tile shade classification has arrived at three main requirements: customer satisfaction, minimising the number of shades, and having a clear, understandable and practical system.

The optimum strategy for meeting these requirements varies depending on whether the tile manufacturer establishes new shades each production run or wishes to match shades from one production run to the next production run. The strategy described for new shades each run is similar to that normally followed intuitively during visual inspection, though its implementation on an automatic tile inspection system can be expected to give more consistent, reliable results at a lower inspection cost.

The strategy proposed for shades which must be used over many production runs, however, is less intuitive and depends on the use of regular hexagon shaped shade areas in two-dimensional shade space. This strategy typically requires about 30% less shades than if square shade areas are used, and has further reductions in the number of shades relative to manually established shade schemes. It also does not produce the many anomalies which are inevitable when shade schemes are set up by eye.

The key objectives of customer satisfaction and minimum number of shades can therefore be more fully achieved if automatic tile inspection systems do not just re-create manually set up shade schemes, but rather guide the user in setting up rational schemes that follow the proposed strategies.