Machines Can See

... but not as we do

Bruce Batchelor

Chapter 7

Sample Applications

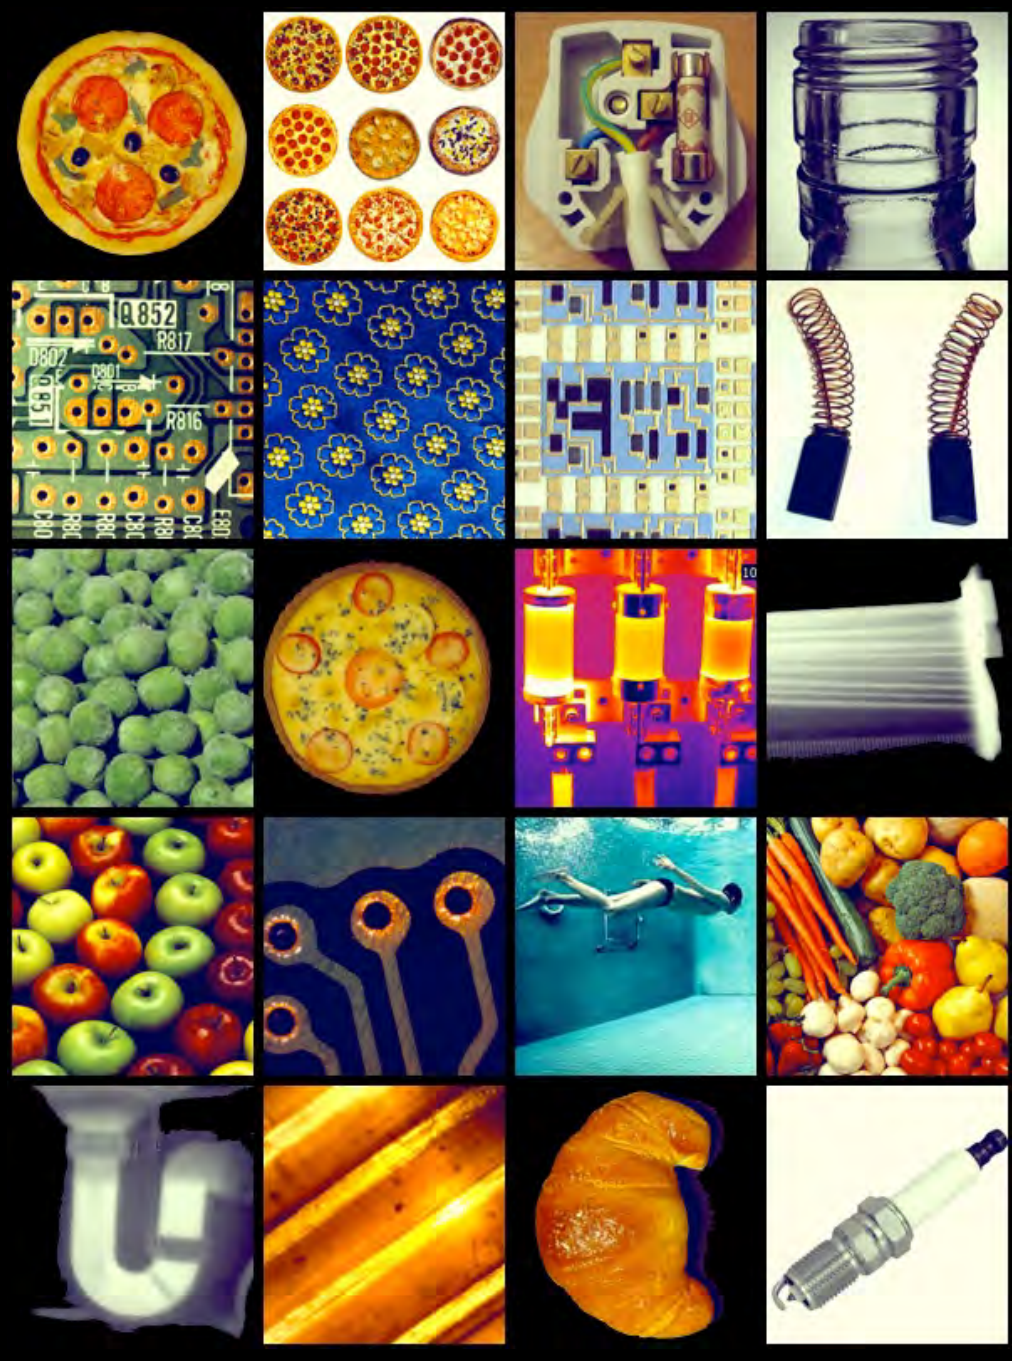

MOTIVATION: HIGH PRODUCTION SPEED

This picture illustrates a situation where automated inspection is essential: 100% visual inspection is impossible because the production rate is too high.

There are numerous places on a bottling line where a vision system might be employed, such as:

- Examining a wide scene of the transport system, to ensure its smooth operation
- Inspecting unfilled bottles for faults
- Making sure that the caps are fitted properly
- Checking the filling level.
- Checking labels

MOTIVATION: NO PEOPLE AROUND

This picture reinforces the message outlined on the previous page. Sometimes there are no personnel in place to take remedial action if anything goes wrong. Vision systems can provide vigilant monitoring of the manufacturing and transportation systems with little or no human involvement.

The vision system might be required to examine every product, to take a broader view for statistical-analysis purposes, or to monitor/control the manufacturing process.

MOTIVATION: AVOIDING UNPLEASANT/DANGEROUS CONDITIONS

Vision systems can take over from human inspectors when the working conditions are unpleasant or dangerous. Inspecting red-hot steel strip is a prime example. A camera can be placed close to the hot material, where a person cannot work safety. Moreover, a machine might be able to detect defects that a human being cannot see by using non-visible radiation (infra-red)

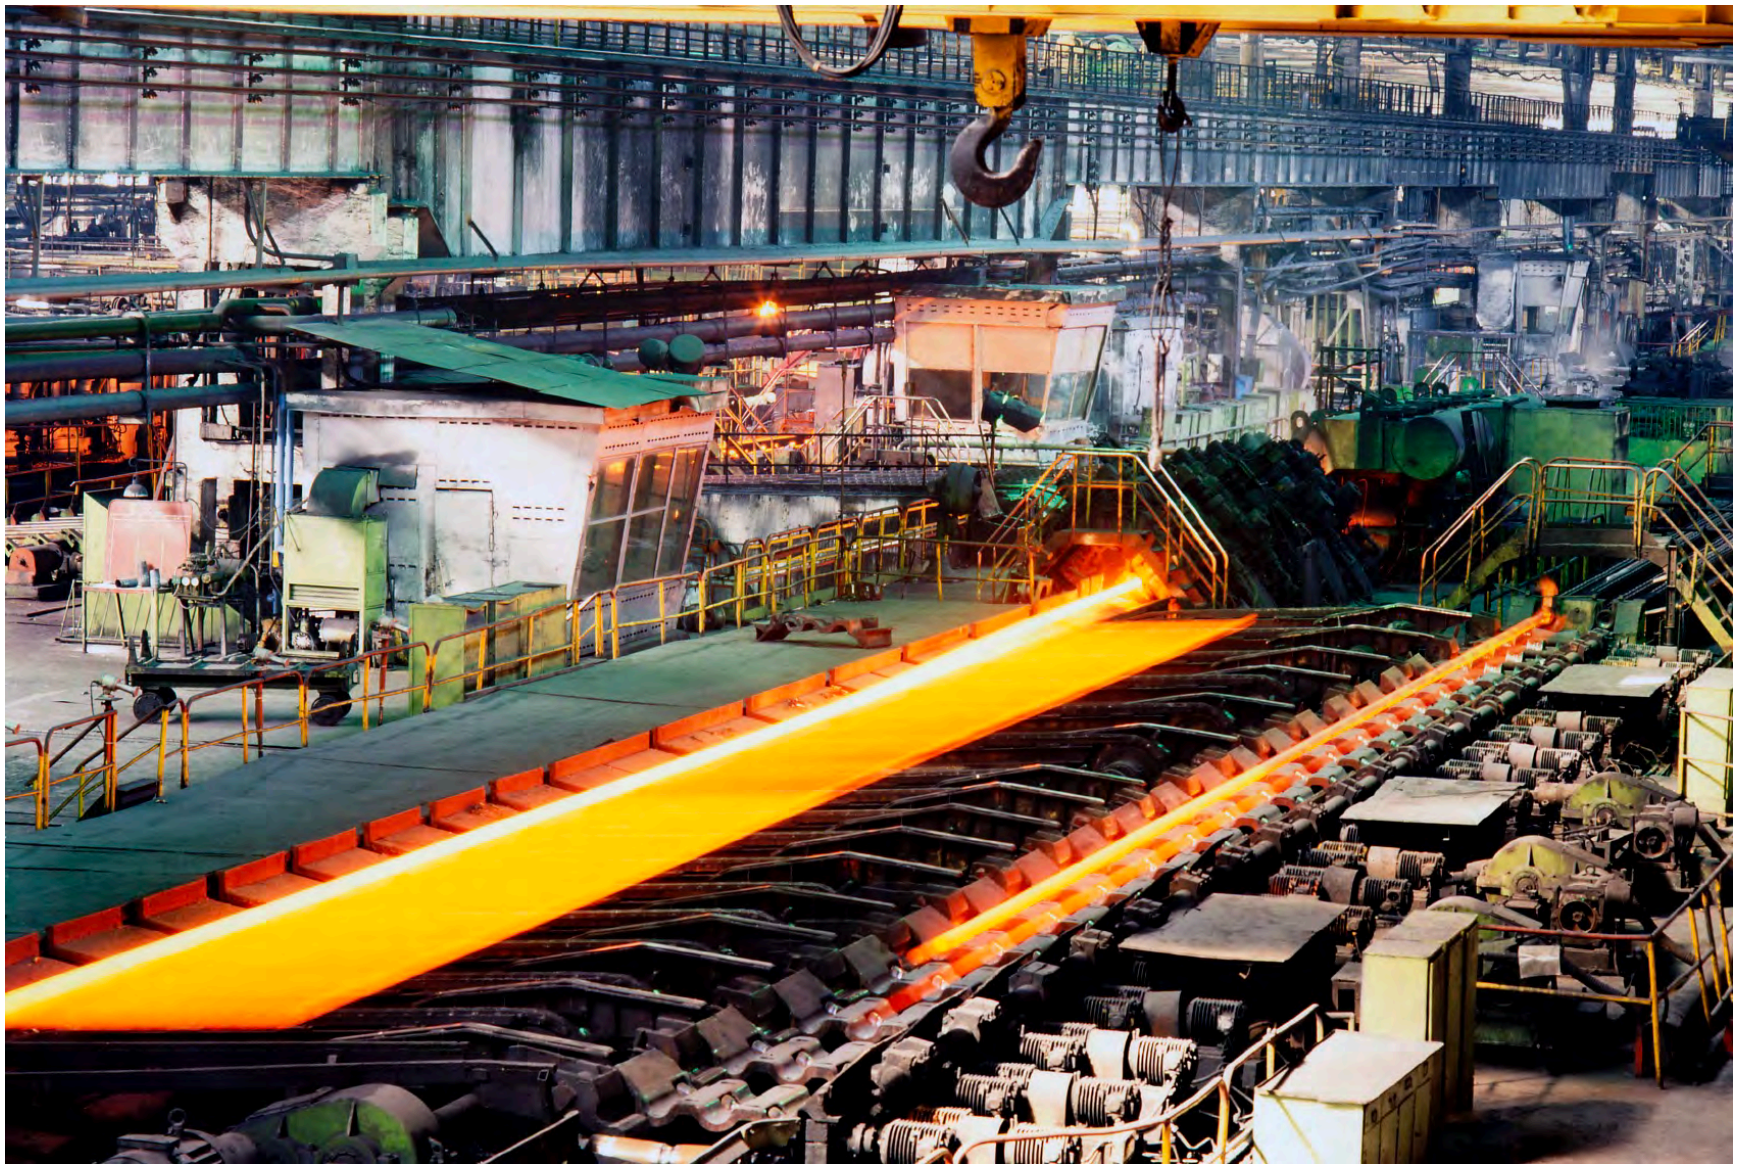

MOTIVATION: ROBOTS NEED EYES TOO!

Automated manufacturing, using robots, requires vision for a number of reasons. It cannot be assumed that a robot will pick up objects perfectly every time. They may be faulty, oily or slightly out of position. The robot's gripper may be contaminated with oil, worn, or physically damaged. For the sake of safety and to ensure continuity of production, robots need watching, by an independent camera-computer system. Vigilant monitoring of robot operation is required at all times,

A vision system can also guide robots in its operation cycle, take measurements. It can accommodate flexible materials (e.g. fabric), non-rigid packaging and highly-variable objects, such as fruit, vegetables and food products (e.g. pizzas).

MOTIVATION: INSPECTING WITHOUT TOUCHING

Machine Vision is ideal for inspecting food products. By providing non-contact sensing, it is completely hygienic and does not disturb even the softest of materials.

WRONG QUESTION: WHAT IS THIS?

In a factory, the product is always known. For this reason, an industrial vision system need not address answer general open-ended questions such as "*What is this?*". Instead, it should answer more precise questions, for example "*Is this a well made …?”*

The effect of this single point is profound: industrial Machine Vision systems can be made faster, cheaper, more reliable and more accurate.

Natural products such as fruit, vegetables and food products (e.g. pies, cakes, pizza, etc) are more variable than engineering artefacts.. Even here, specific questions can be posed. For example, tomatoes are green, yellow, orange or red: they are never blue. Neither are they ever bigger than 100mm in diameter and tomatoes smaller than 5mm are of no value.

Designing industrial vision systems should *always* take account of application knowledge;

Emulating a human being, a vision system should never be the primary goal.

* Relaxing in warm volcanic sand

VISION SYSTEMS ARE NOT PRETTY BUT THEY MUST BE TOUGH

Automated Visual Inspection systems are not photogenic but they must be tough: able to withstand very harsh environments, flying debris, dust, steam, air-born aerosol contaminants, malicious damage (by disgruntled workers), high temperatures high /variable ambient light-levels. These are all hazards that must be tolerated for a vision system designer to consider. Moreover, the ma chine must allow easy cleaning and unjamming. It must also protect factory personnel, particularly from flashing lights, laser radiation and x-rays.

The photograph shows an x-ray inspection system used in the food industry. It has a lead-lined radiation-protection enclosure and can be stripped down for cleaning in as little as 30 seconds. Cleaning involve may involve spraying the equipment with a high-pressure water jet, or steam. There is one more requirement: it must be designed to avoid trapping of food waste.

Neglect of any of these issues may render the inspection system useless, however clever the inspection software is.

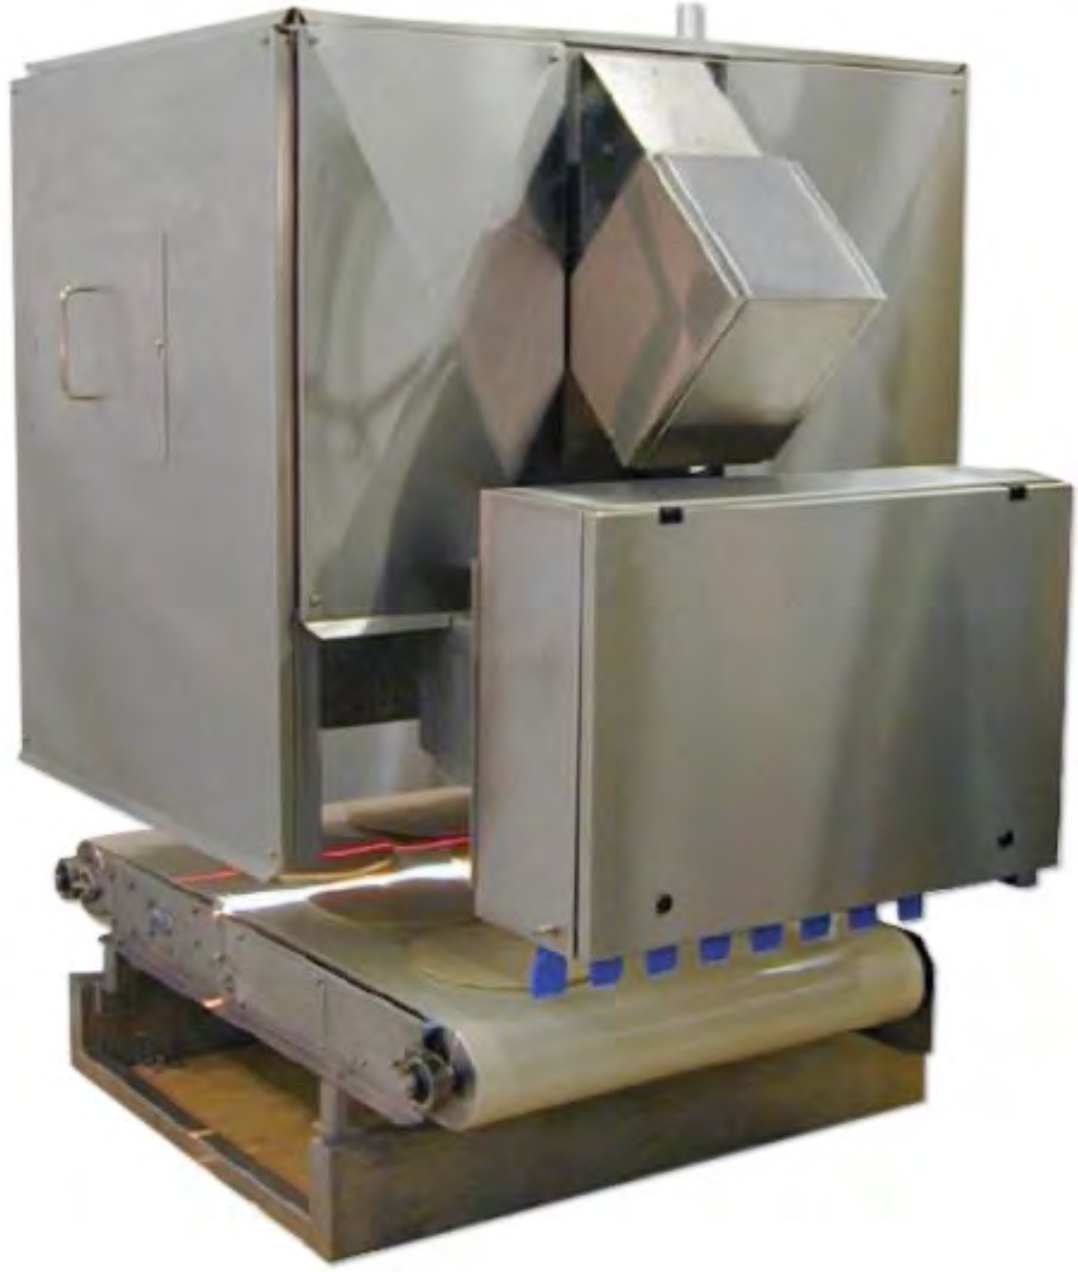

A MACHINE HAS A DIFFERENT POINT OF VIEW

To the human eye, the two oblique views of the object (an hydraulics manifold) in the top row give a good impression of its shape. However, for a vision system, orthogonal views (central and bottom rows) are more useful and allows more precise measurements to be taken. Since sustained, precise positioning is needed to obtain high-accuracy measurements the two images in the top row are of little practical use. (Obliques views may be useful to peer into holes to examine their side walls.) The general point to note is that human and machine vision have different strengths and weaknesses and different viewing angles.

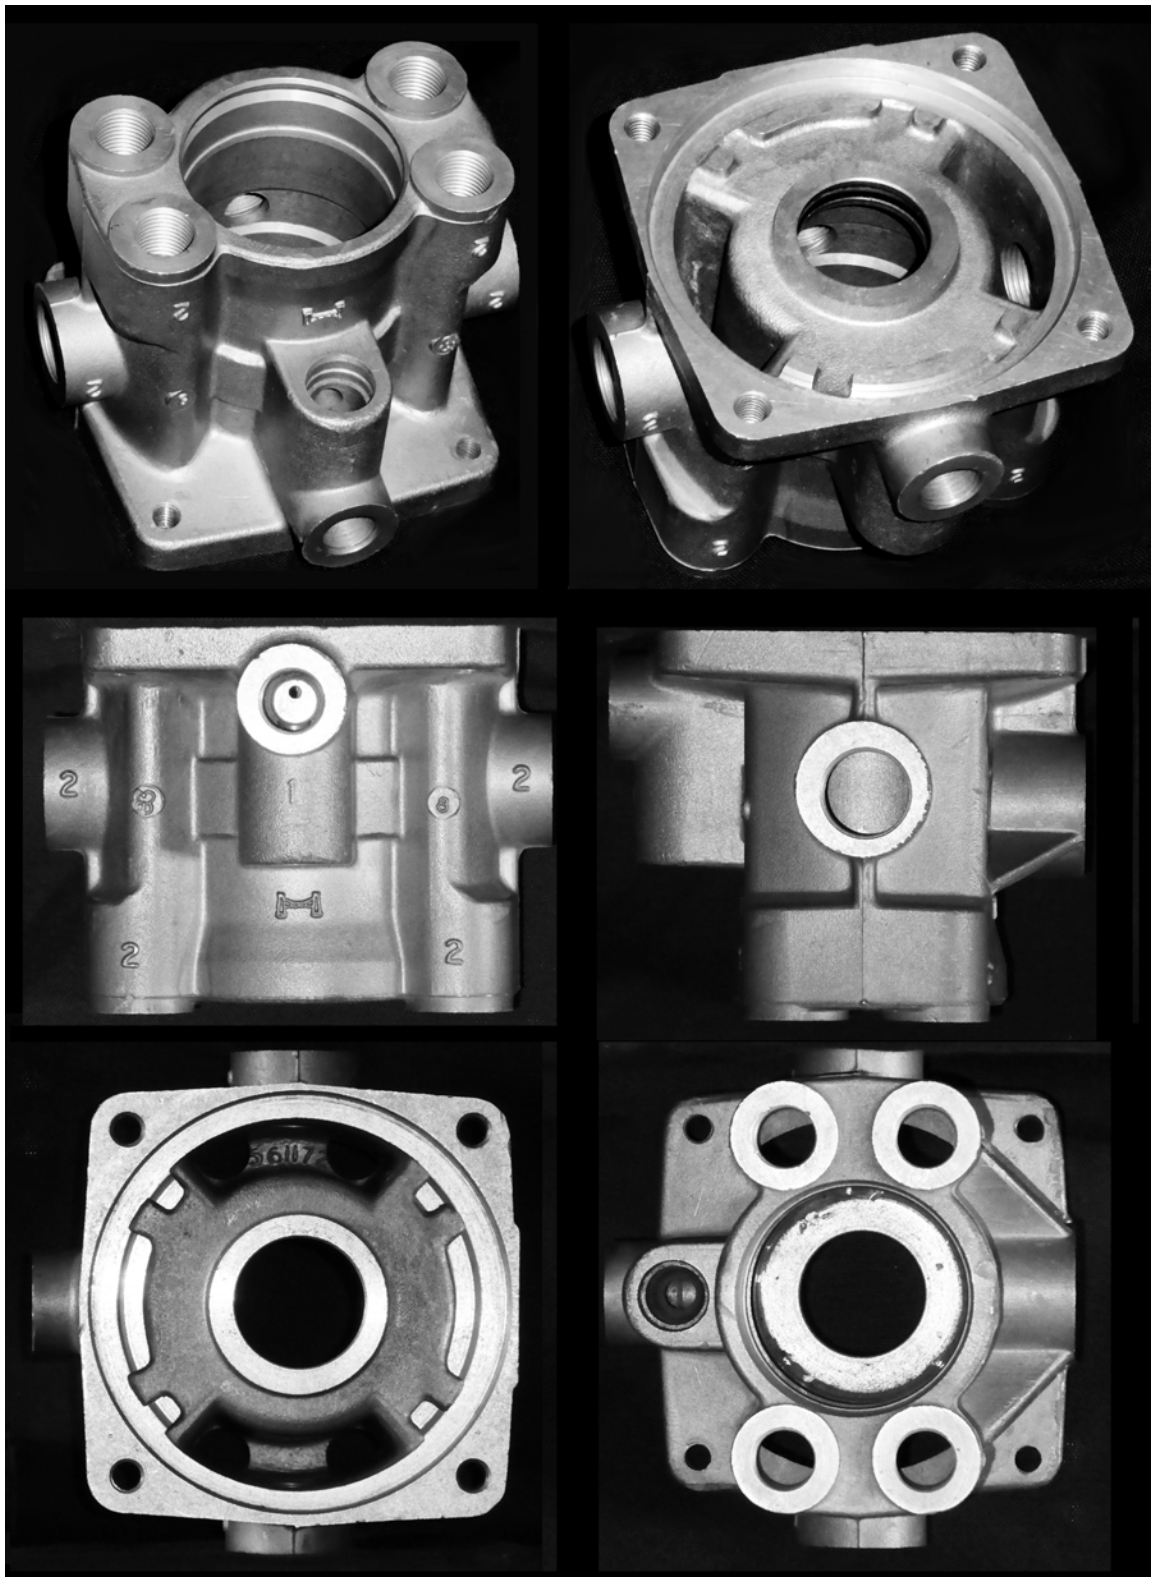

DON'T BE FOOLED BY A PRETTY PICTURE

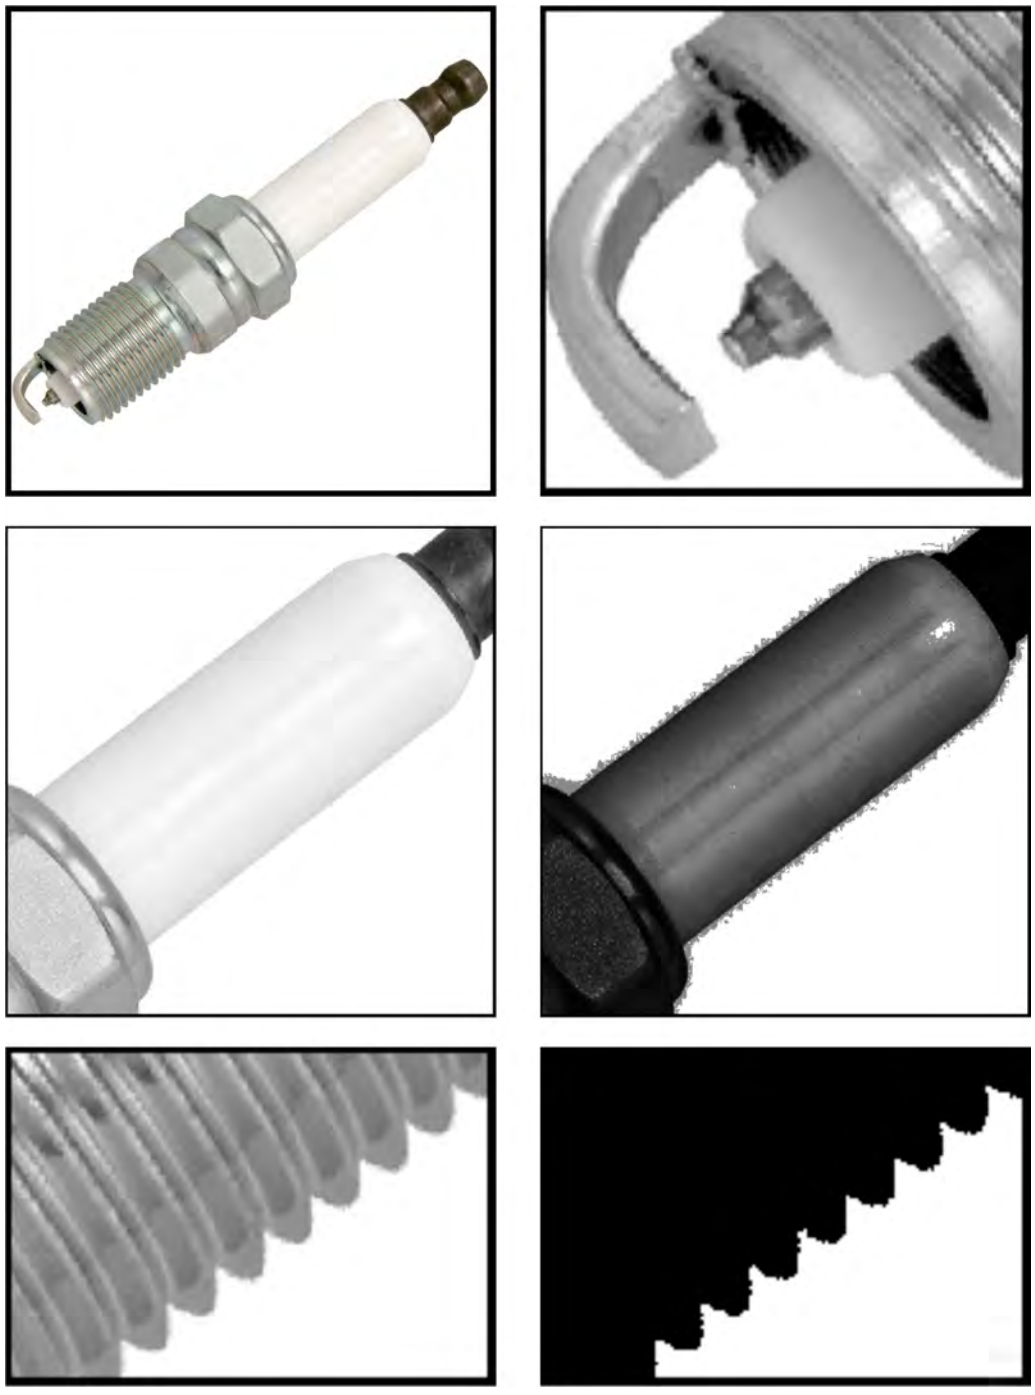

(Top-left) To the human eye, this is "a good picture" of a petrol-engine spark-plug but it is almost useless for a Machine Vision system. (This is a JPEG image. JPEG coding presents of problems for automated analysis and should be avoided where possible.)

(Top-right) An inaccurate measurement of the width of the spark gap is obtained from this expanded view of [TL].

(Centre-left) The ceramic insulator appears to have a well-defined edge in this expanded view of [TL].

(Centre-right) Contrast-enhancement emphasises a serious deficiency of JPEG coding. The edge appears ragged

(Bottom-left) The screw-tread appears to have a smooth edge in this expanded view of [TL].

(Bottom-right) Processing {BL] shows a ragged edge

DON'T WASTE PIXELS

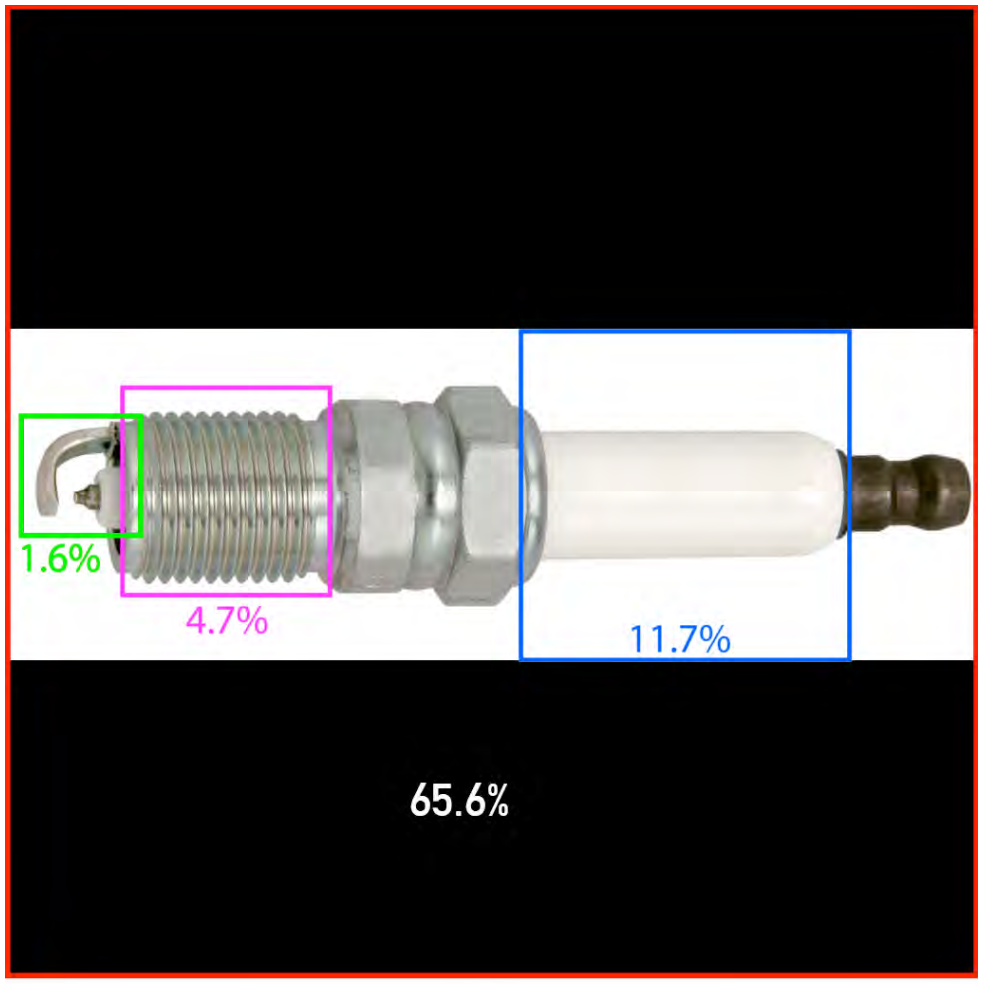

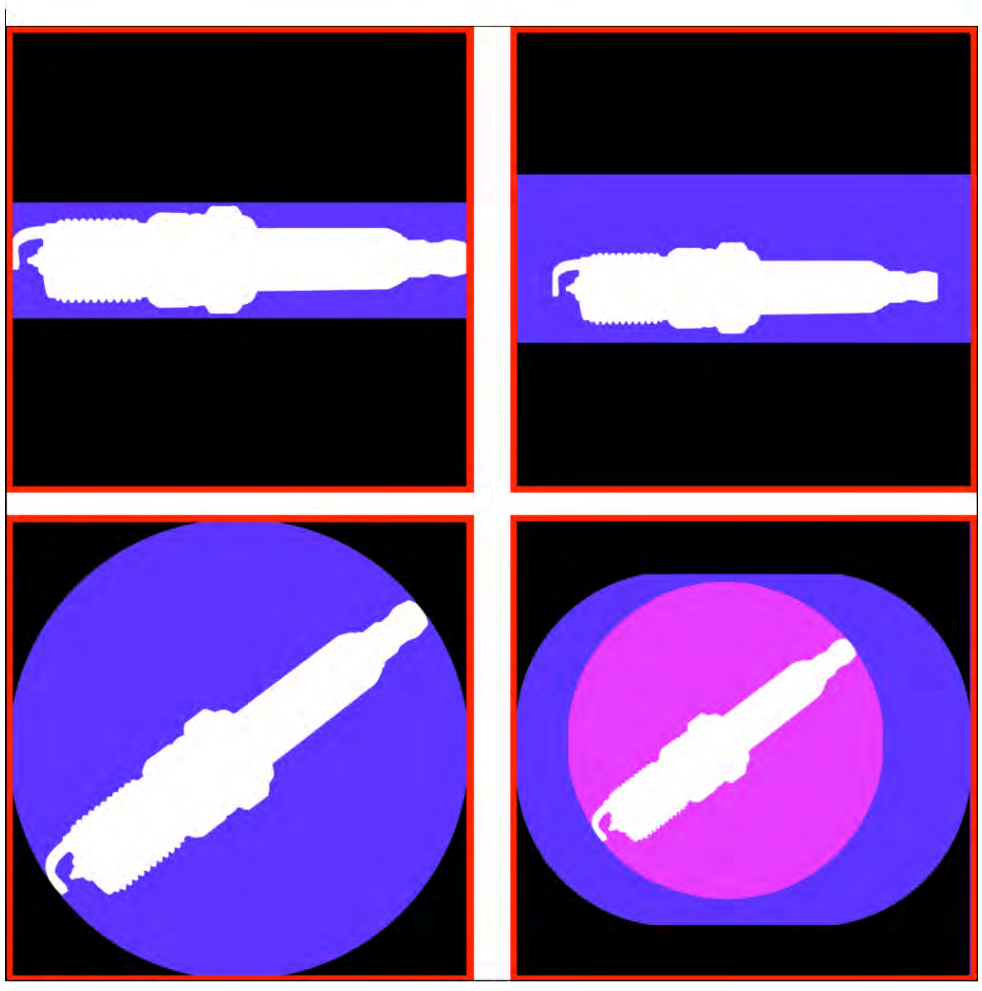

(Top) A spark-plug has three critical areas: white insulator, screw thread and spark-gap. Since it is long and thin, over two thirds of the pixels (black, 65.6%) are never required for inspection. This assumes that the spark-plug is always examined in the same position and orientation. Inspecting the insulator requires only 11.7% of the pixels (blue square); the screw thread, 4.7% (mauve square); the spark gap 1.6% (green square).

(Centre-left) If the spark-plug is always examined in the same position and orientation, black pixels are not of interest. Image processing speed can therefore be improved, simply by ignoring them

(Centre-right) If the orientation is fixed but the position is allowed to vary, more pixels must be processed.

(Bottom-left) If the position is fixed but the orientation is allowed to vary, the black pixels can be ignored.

(Bottom-right) If the position is fixed but the orientation are both allowed to vary, the number of pixels that can can be ignored (black) is decreased considerably.

IGNORE THE FACT THAT IT IS PLASTIC - IT IS GREEN & RED

A Machine Vision is concerned solely with appearance. As far an Automated Visual Inspection system is concerned, this a green and red object. The fact that is made from moulded plastic is totally irrelevant!

GENERIC CLASSES OF PROBLEMS REQUIRE ARTIFICIAL INTELLIGENCE

A system that is claimed to be capable of inspecting a broad and ever expanding class of objects (e.g. *Pizzas*) must be intelligent and involve machine learning. Its input would consist of a series of "general-purpose"measurements such as the following:

- Diameter
- A measure of deviation from circularity
- Average width of the bare crust around the edge.
- Average area and number of red/green/black blobs bigger than XXX pixels
- A (numeric) measure of the colour of the "background"
- A (numeric) measure of the "randomness"/uniformity of the topping
- etc.

These measurements can be based upon the ideas discussed elsewhere in this book.

EXPECT THE UNEXPECTED (If anything can go wrong, it will go wrong!)

An Automated Visual Inspection system must be designed to accommodate the unexpected. Defective objects presented to it for examination are likely to be very variable in form. It must be able to handle them, without becoming jammed, or damaged.

[Picture: Rail accident, Montparnasse, France 1895]

One approach to this to design a Machine Vision system that is able to learn what is “usual”. On a mass-production line, any product that deviates too far from the norm is defective and should be rejected.

MEASUREMENT BY PROXY

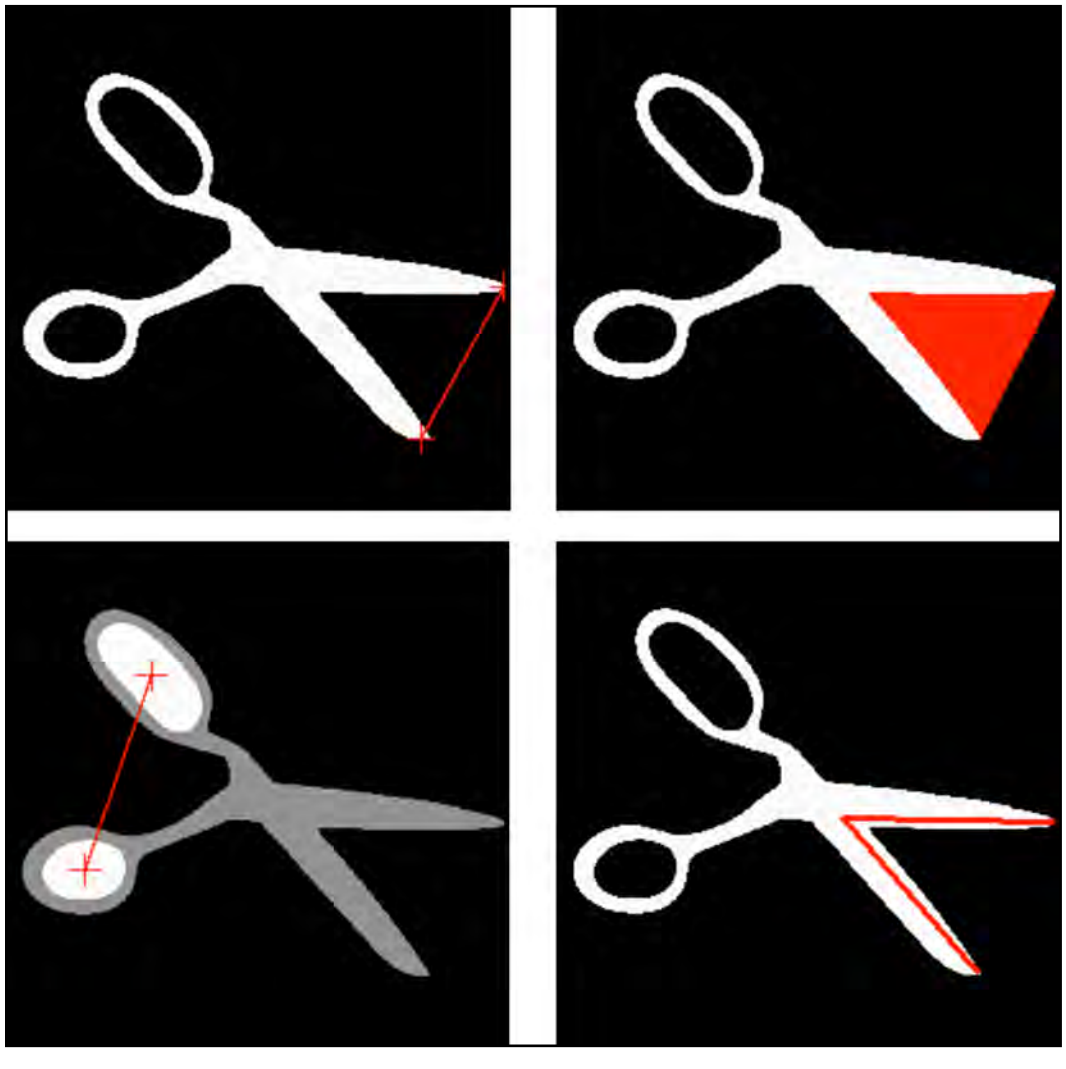

Measurement by proxy offers an important “escape route” when it is not possible to estimate an important parameter directly For example, the volume of an apple is a good estimator of its weight and an approximate estimate of volume can be calculated from two orthogonal side views, For an egg, just one would suffice, because it is rotationally symmetrical.

In the illustration opposite, four measurements are directly (non-linearly) related to each other. Measuring any one of them uniquely determines the values of the other three.

(Top-left) The distance between the tips of the blades

(Top-right) The area of one of the major bays (red).

(Bottom-left) The distance between the centroids of the lakes (white).

(Bottom-right) The angle defined by the red lines.

GLASS & CLEAR PLASTIC

BOTTLE SHAPE (checking that the sides are parallel and vertical)

(Top-left) Original image, back illumination. The corners are darker that the rest of the background, due to poor lighting.

(Top-centre) Edge detector. Eliminates back-ground shading.

(Top-right) Binary image, de-rived from [TC].

(Centre-left) Straight lines fit-ted to the sides. These allow us to check that the bottle has parallel sides and that its axis of symmetry is vertical. Geometric distortion is due to the lens.

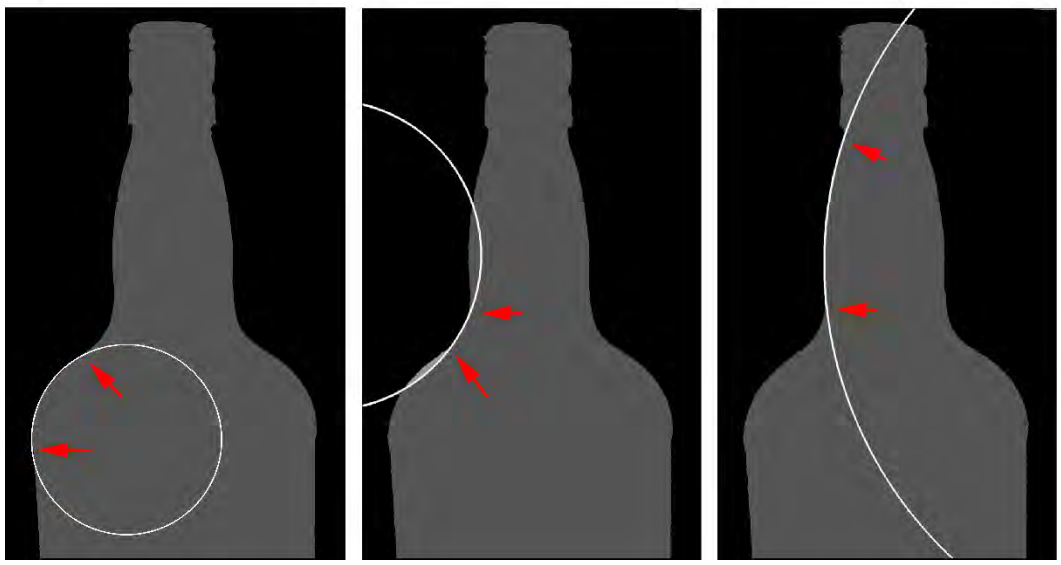

(Centre-centre) Checking that the bottle is vertical & symmetrical. Thite arc is equidistant from the left- and right-most edges. Ideally, it should be a vertical straight line.

(Bottom-left & Bottom-centre) Checking that the shoulder is well formed. Each circle was fitted to three

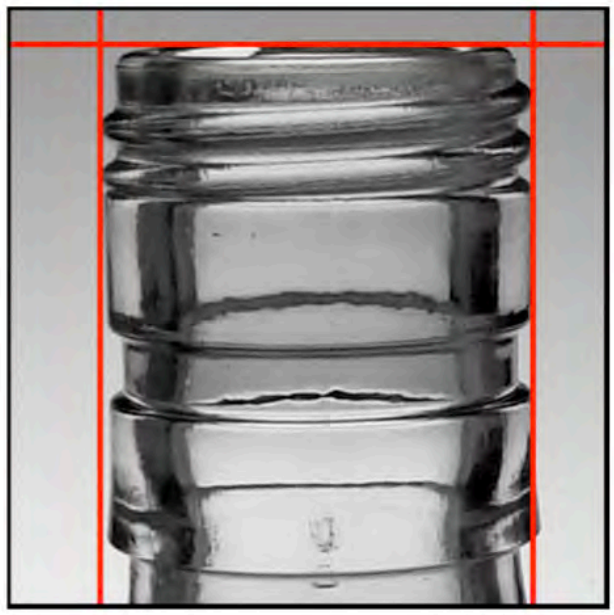

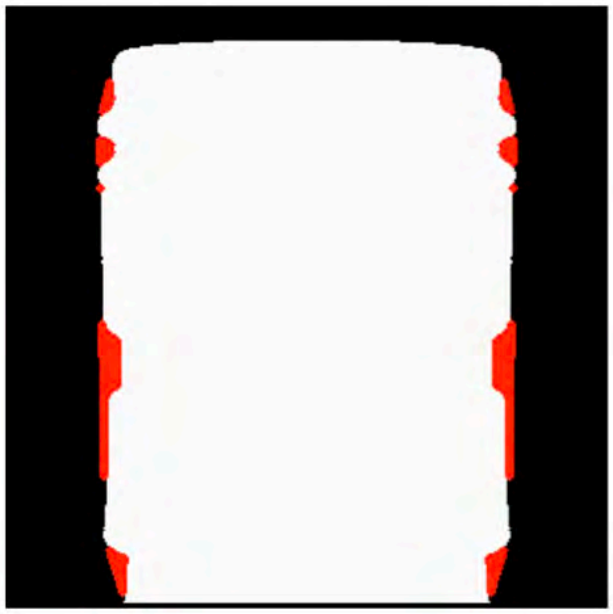

THREAD ON A

(Top-left) Original image, obtained using back illumination.

(Top-right) Top and sides of the mouth region identified.

(Bottom-left) Differences between the silhouette and its convex hull.

(Bottom-right) Dark horizontal streaks identified. One mild blurring filter and another strong blurring filter were applied and the resulting images subtracted. The resulting image was then thresholded. Unwanted clutter at the sides has been eliminated, by masking.

CLEAR PLASTIC BOTTLE

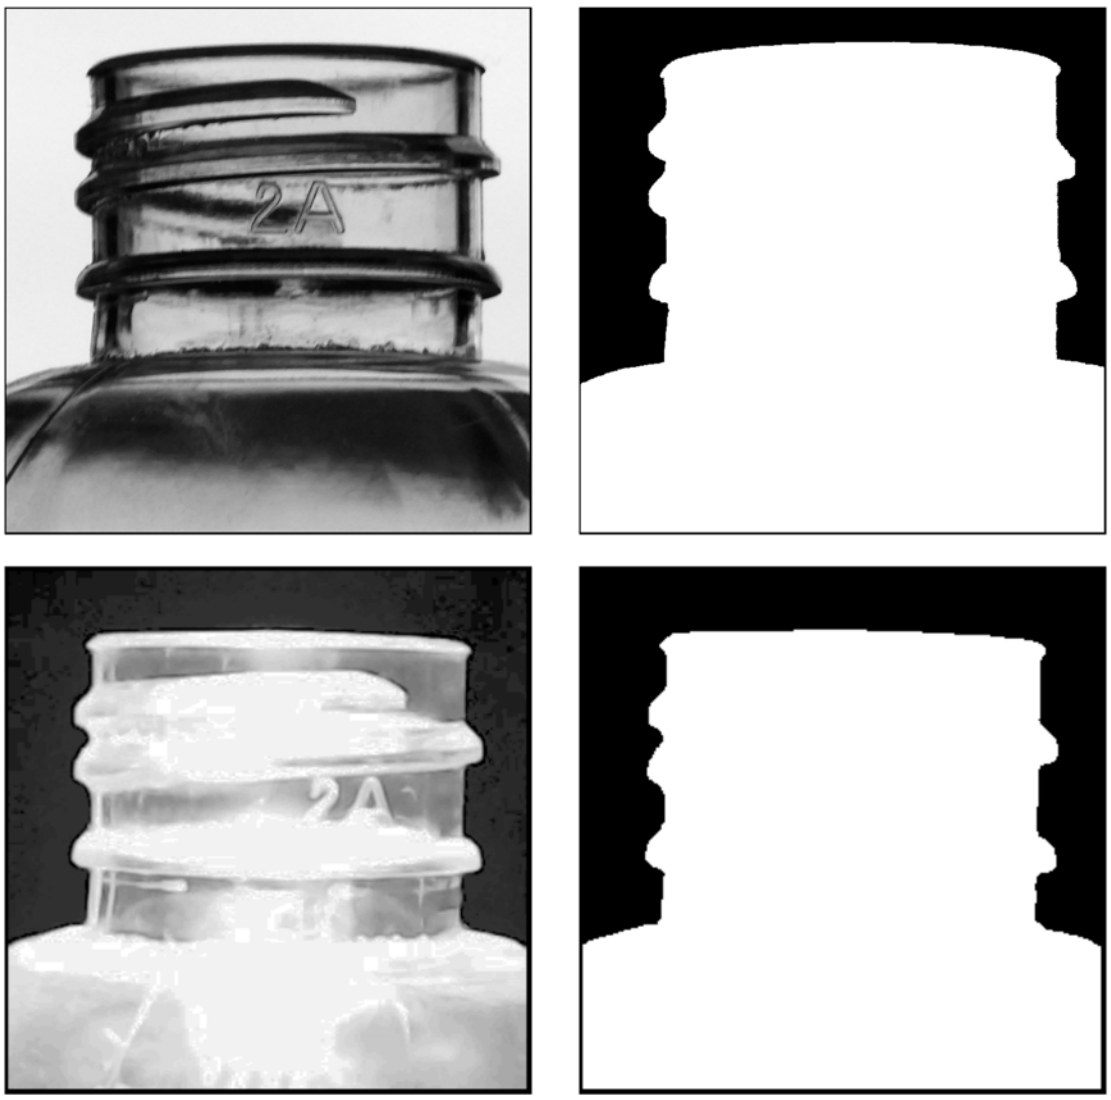

Many of the techniques used to inspect glass can also be used on clear plastic bottles. Although their refractive indices are different, the main difference is the thickness of the walls of the bottle. These images also demonstrate the difference between the images obtained using back-and dark-field illumination

(Top-left) Original image, back illumination.

(Top-right) Binary image derived from [TL].

(Bottom-left) Original image, dark field illumination.

(Bottom-right) Binary image derived from [BL].

BOTTLE-WALL DEFECTS (bird-swing or birdcage)

BIrd-swing (or Bird-cage) is a filament of glass spanning the interior of a bottle.It is created by a manufacturing fault and can only be detected visually. A 2-stage moulding process is used to make bottles.

(Top-left) BIrd-swing, side view, back lighting.

(Top-right) BIrd-swing, end-on view, back- lighting. [A colleague gave me this bottle, which had contained black-currant juice.]

(Bottom-left)) BIrd-swing, dark-field illumination.

(Bottom-right) Spike, dark-field illumination. This is a “stalagmite” of glass. Again, this is a manufacturing fault and occurs in jars, made by pressing the semi-molten glass into a mould. A spike can be formed as the plunger is withdrawn. [My daughter was given milk to drink from this bottle. Fortunately, my wife realised in time!]

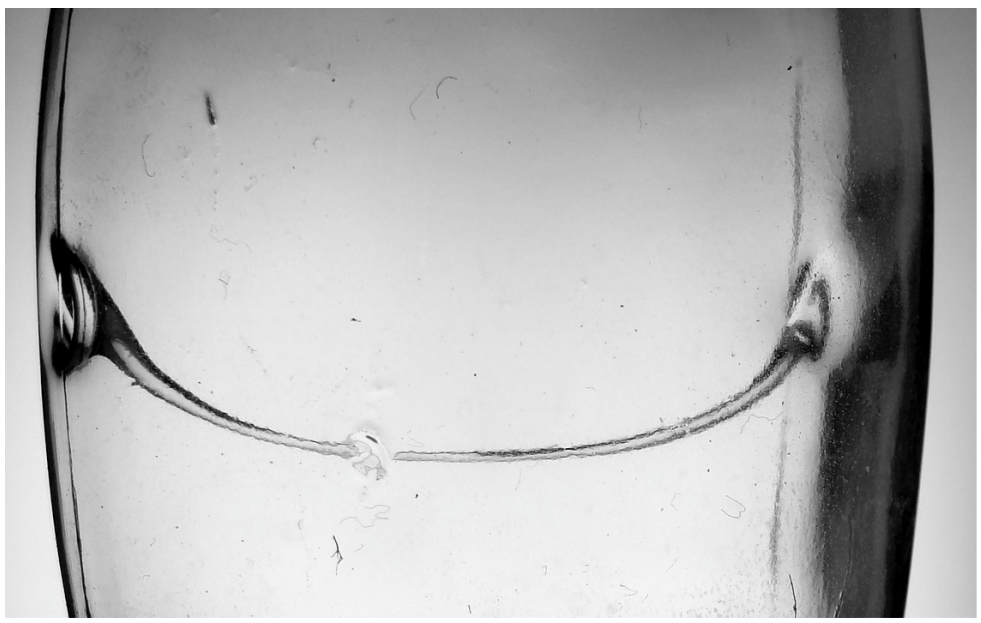

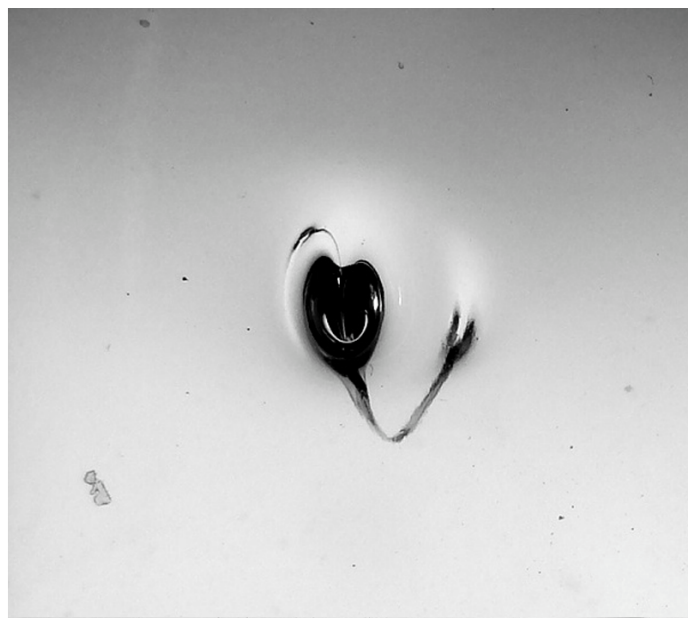

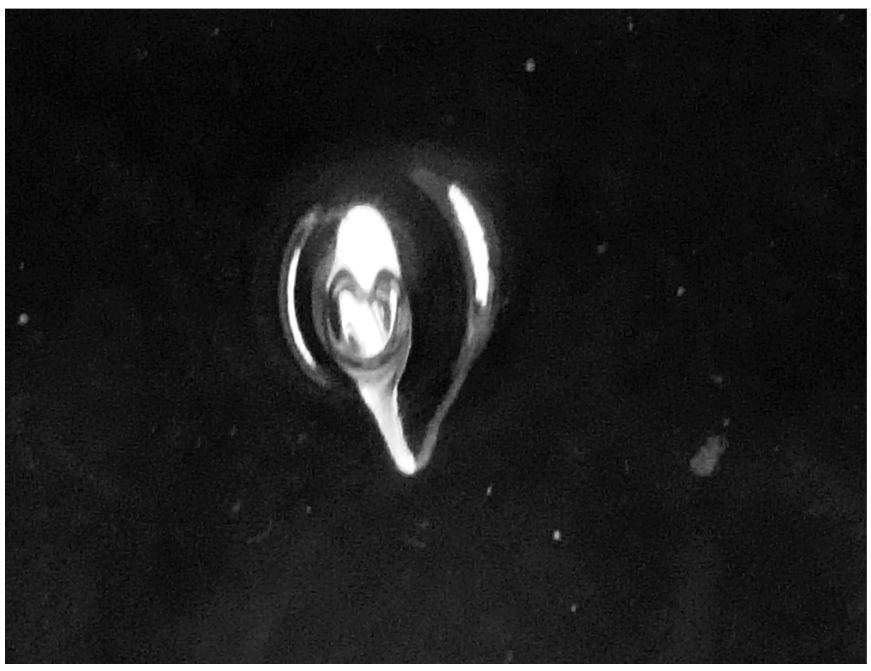

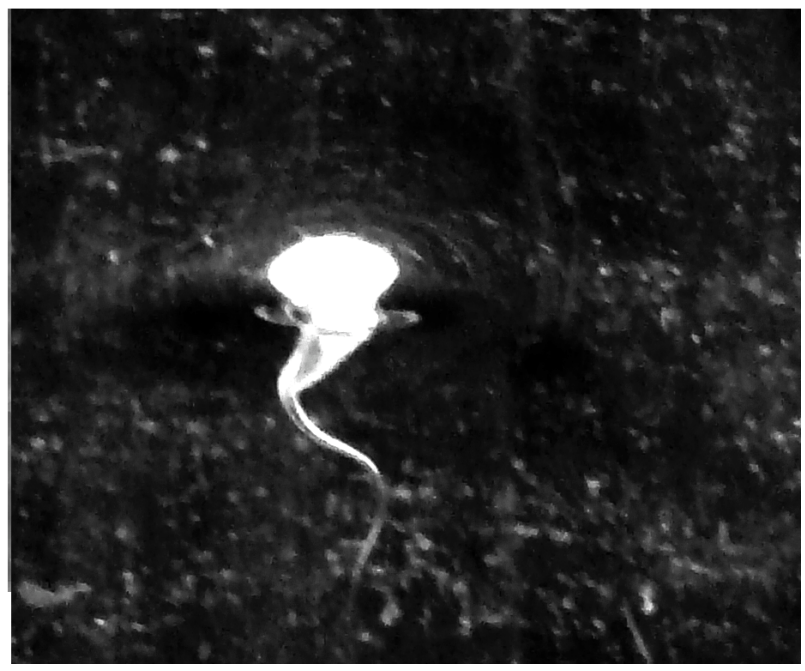

SHATTERED SAFETY GLASS / CRACKLE GLASS

The fracture pattern of shattered glass is an important feature of automobile windscreens, determining its safety in an accident. The illustration here uses Crackle Glass, which is used for its decorative effect.

(Top-left) Original image, back lighting.

(Top-right) Binary image, showing individual glass fragments.

(Bottom-left) Fragments shaded according to size: largest are brightest. (8 size categories) Fragments on the edge of the image have been removed, since they would distort the measured distribution of sizes. The aspect ratio of these fragments is an important safety parameter and can be estimated in a number of ways.

(Bottom-right) Rendering [BL] in pseudo-colour for easier viewing

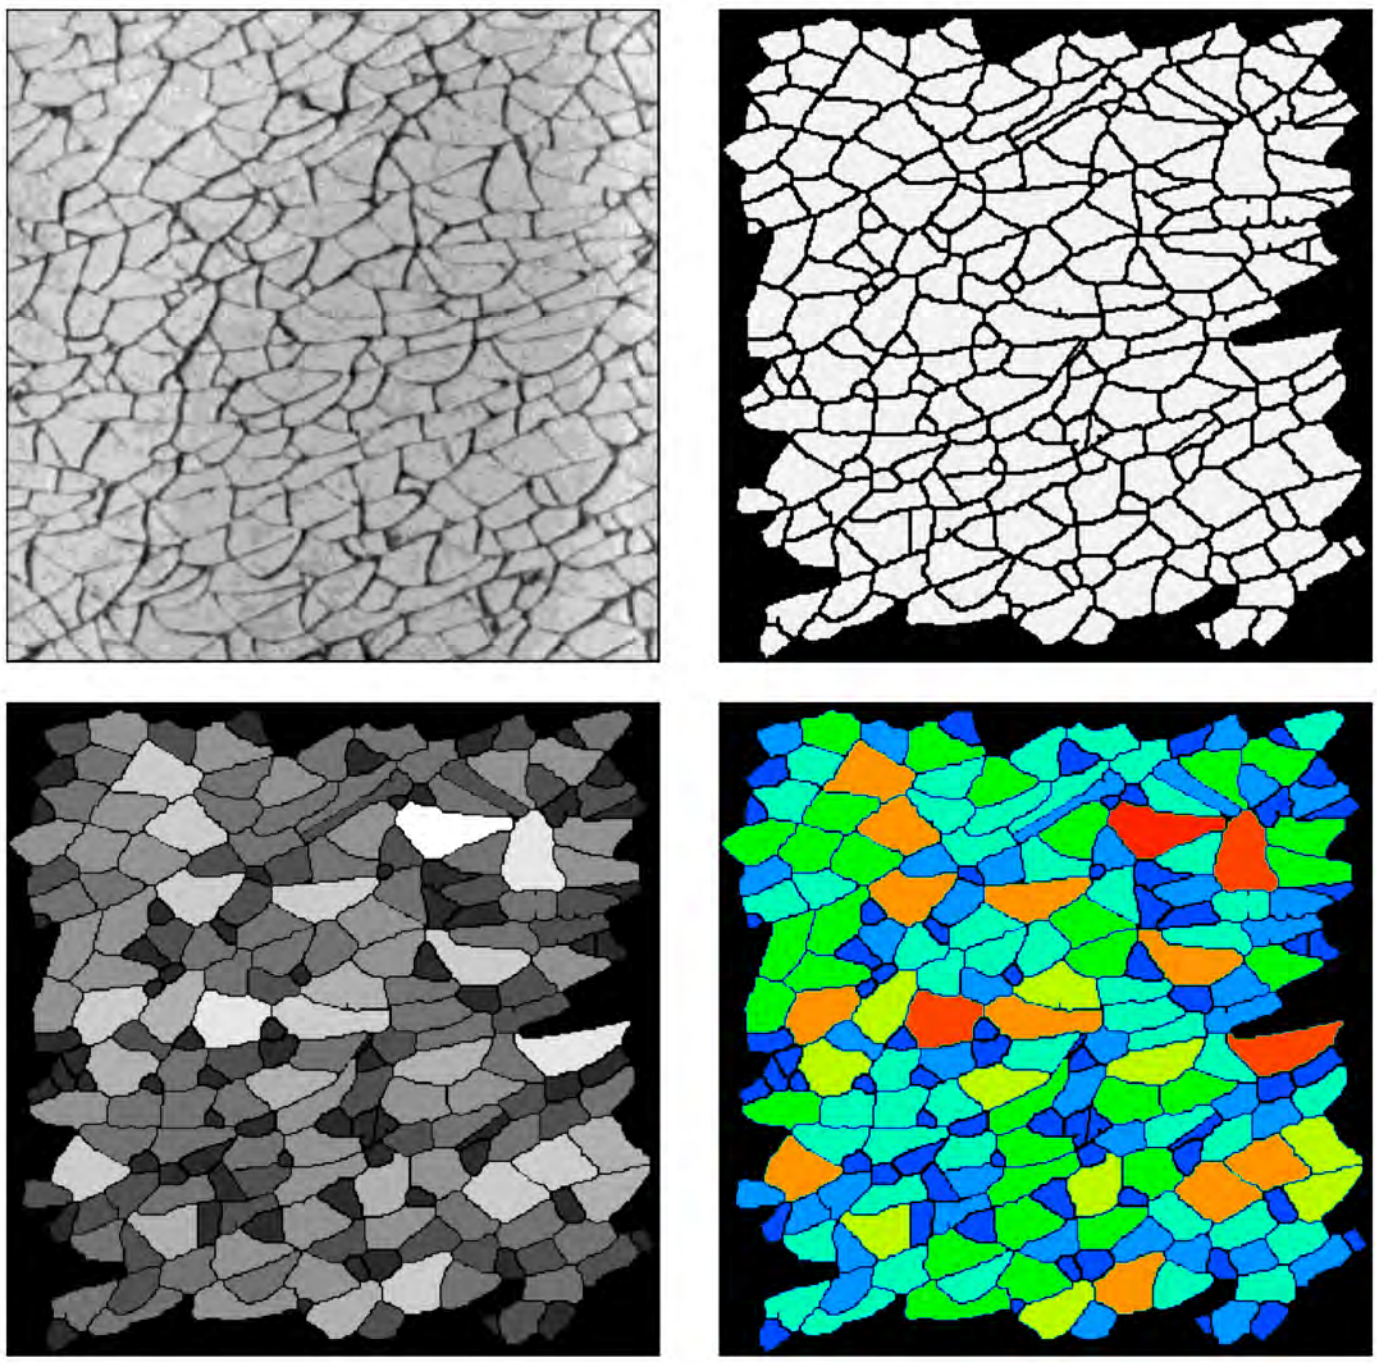

LENS (visualise the warping effect of the lens)

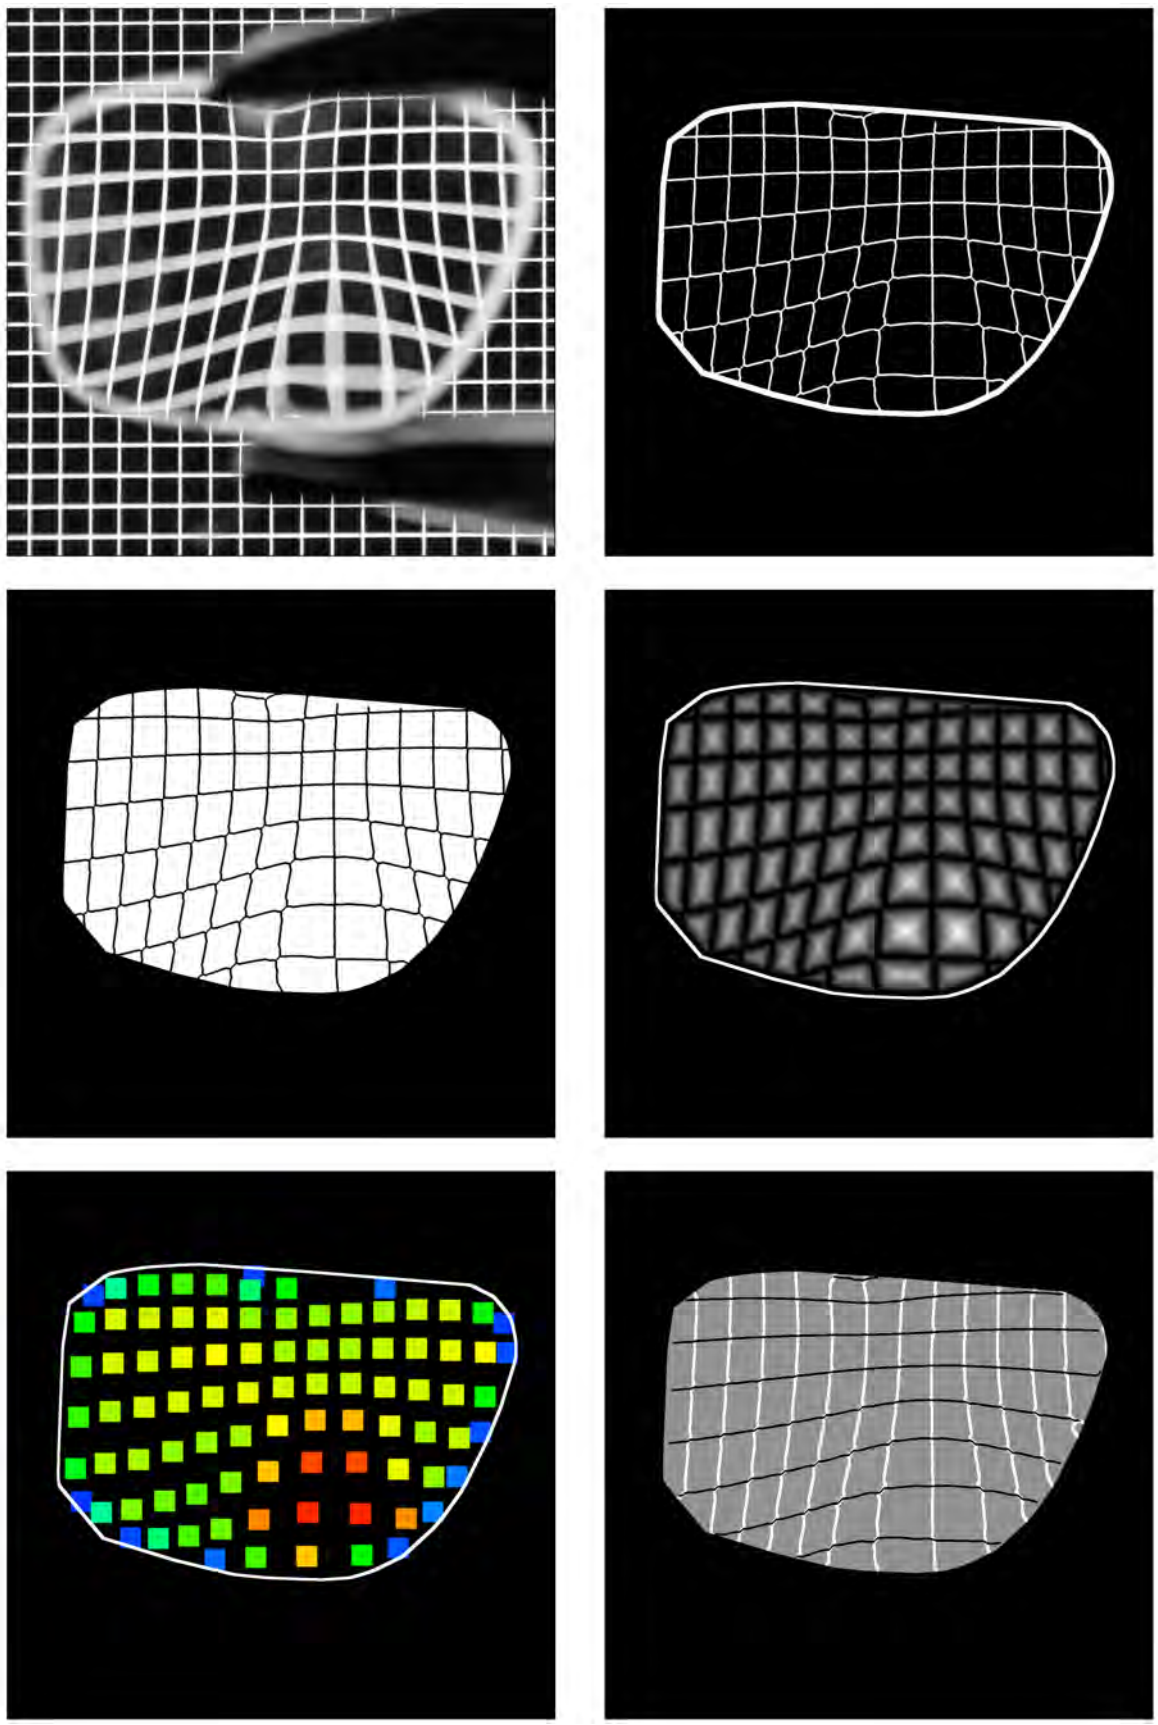

Using a patterned background to visualise the' warping effects of a variable-focus lens.

(Top-left) Original image. The lens is in front of a regular square grid.

(Top-right) Mesh has been reduced to thin arcs

(Centre-left) Invert the contrast, ready for next operation in [CR].

(Centre-right) Grass-fire transform applied to [CL]. Notice that big "cells" are brighter than small ones.

(Bottom-left) Find the local maxima in [CR] and represent the peak intensity in pseudo colour. Red spots indicate large "cells; green and blue, smaller ones.

(Bottom-right) Horizontal and vertical arcs in [TL] have been separated. The curvature of these arcs can be measured to quantify the warping effect of the lens.

BOTTLE SIDE-WALL THICKNESS

(Top- left) Original image, glass tumbler with a thick base (Below A) and partially filled with water(A to B). The background stripes are distorted by irregularities in the thickness of the glass and are magnified by the cylinder of water, which acts as a crude lens.

(Top-right) Binary image, with edge contours superimposed.

(Bottom-left) The vertical centre lines of the bright stripes have been detected. The straightness and thickness of the vertical white lines provide a quick, easy and approximate way to check the thickness of the side walls of the the tumbler.

(Bottom-right) A jar filled with water viewed against background with a single dark stripe (between the red arrows). The lens formed by the cylinder of water magnifies the dark stripe, providing a simple check that the jar is full.

CHECKING FILLING LEVEL OF A GLASS VIAL

(Left) Original image, obtained using back-lighting.

(Right) Top and bottom of the meniscus. detected using row integration and simple thresholding.

BOTTLE-SHAPE ANALYSiS (dark brown, glass vial)

Inspecting a small dark-brown glass vial.

(Top-left) Original image, dark-field illumination.

(Top-right) Edge detector.

(Bottom-left) Binary image, derived from [TR] by thresholding and filling lakes. Red pixels indicate the centres of the horizontal white chords. Ideally, this should be a vertical straight line.

(Bottom-right) The binary image in [BL] was flipped about a vertical line passing through the centroid. The resulting image was then subtracted from [BL]. Ideally, this should contain only grey pixels.

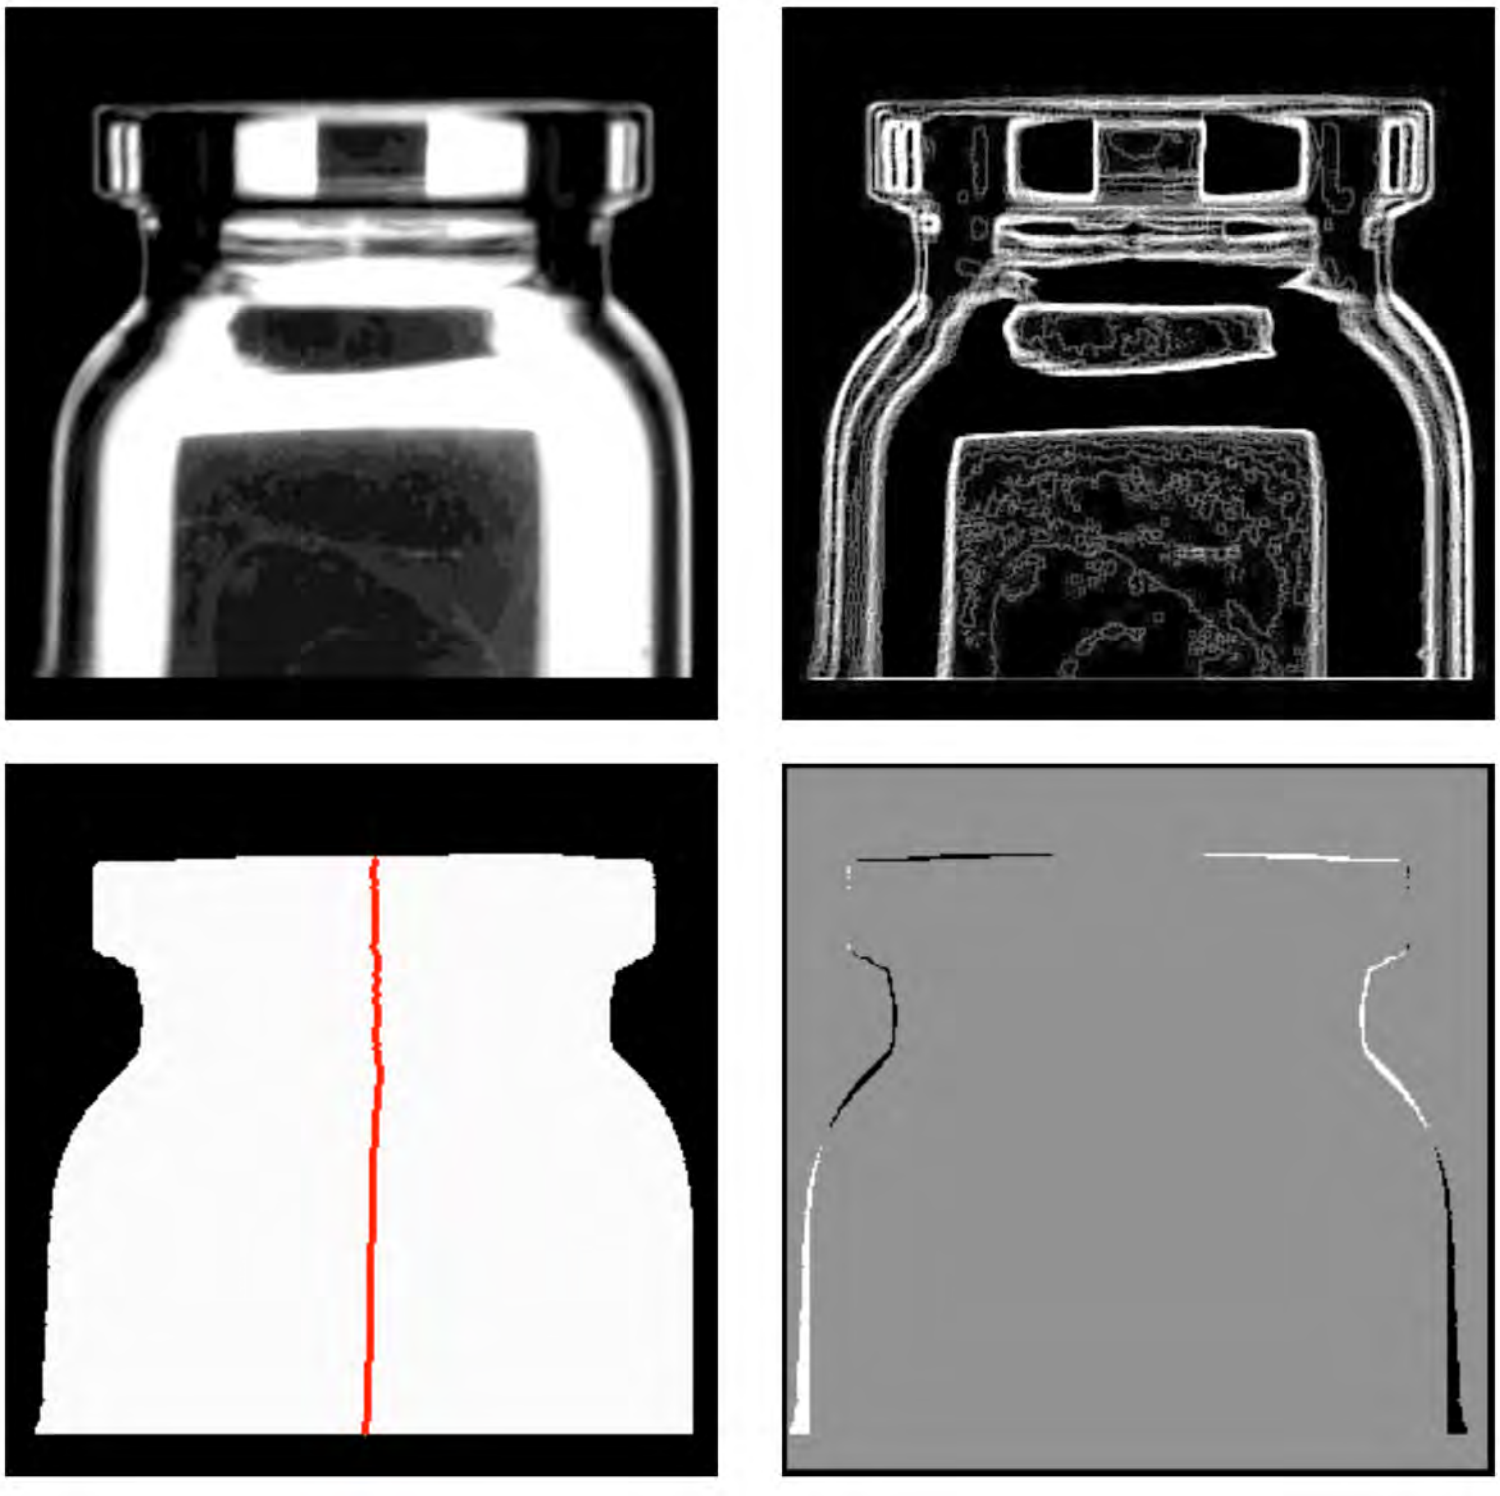

BOTTLE-MOUTH INSPECTION

Faults around the mouths of bottles present a serious health hazard as sharp glass fragments can be swallowed.When I was a boy, glass from the broken top of a milk bottle got stuck in my throat. Fortunately, I was able to cough it up with only minor bleeding. There are several types of critical defects:

- Cracks and splits
- Chipped
- Fragments of glass attached, loosely or firmly, to the inner edge of the mouth.Either of these can be dislodged when a filling tube is inserted into the mouth

BOTTLE MOULDING FAULTS

Moulding fault. This spectacular “fin’ was created when the two halves of the mould did not close properly; molten glass flowed into the gap between them. A fault like this can easily be broken, creating a sharp edge capable of causing serious hand injury.

Also see

"A Vision for the Future", J Victor, URL *http://old.emhartglass.com/files/A0035.pdf*, Accessed 1st May 2020

"New Standard for the Inspection of Glass Bottles", Miho, URL https://miho.de/en/applications/glass-bottles/, Accessed 1st May 2020.

Section 2

NATURAL PRODUCTS

ESTIMATING ANIMAL WEIGHT (pigs)

A system for estimating the size and shape of live pigs was described by J. A. Marchant and C. P. Schofield (*Machine Vision for the Inspection of Natural Products*, Chapter 13)

In practice, several cameras would be needed to obtain a reliable weight estimate. It is helpful if the animal is constrained during image capture.

(Top-left) Original image.

(Top-right) The animal's silhouette was obtained using a colour recognition filter.

(Bottom-left) Outline of the silhouette superimposed onto the grey-scale image.

(Bottom-right) The outline of the animal's body, was derived from [TR] using morphological erosion.

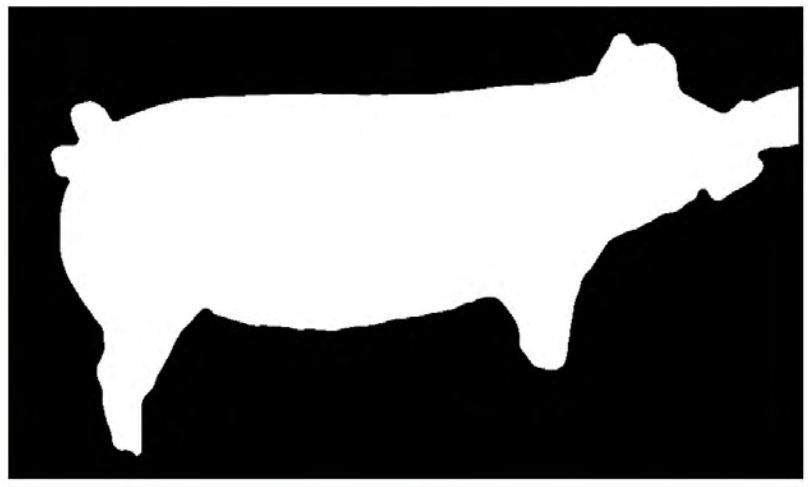

CUTTING POINTS FOR PLANT MICROPROPAGATION

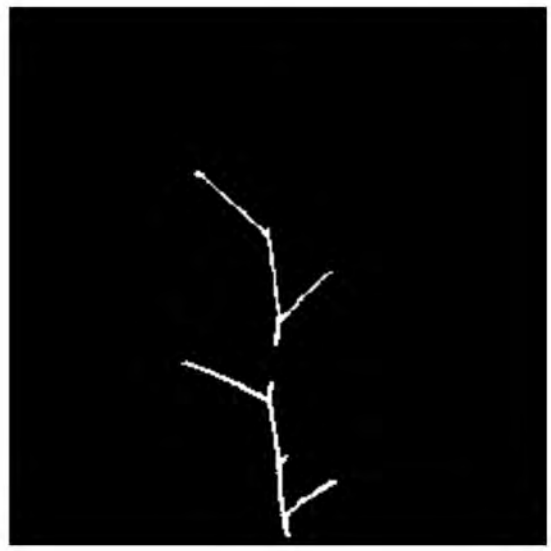

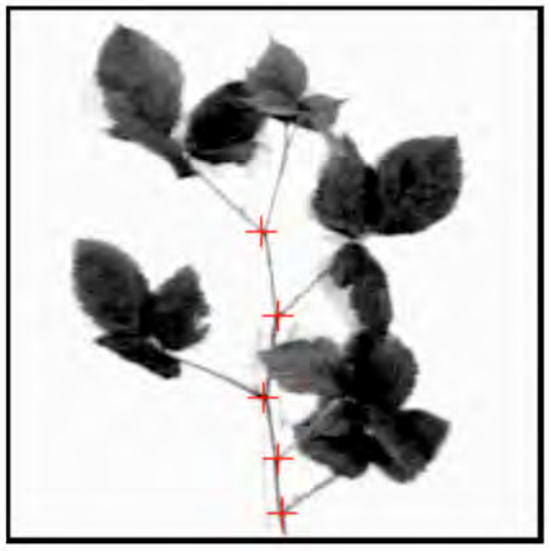

Dissecting an open-structure plant for micropropagation.

(Top-left) Original image.

(Top-right) After applying the “crack detector” algorithm.

(Centre-left) After thresholding [TR].

(Centre-right) Very small white blobs in [CL] have been eliminated.

(Bottom-left) Joints on the skeleton in [CR] indicate the cutting points at nodes on the stem. By cutting the plant just below and just above such a point, we obtain a tiny "Y" of plant material that can be replanted in a nutrient material. This process is the repeated for each plantlet every 3-4 weeks.

PLANT SHAPE (estimating the health of a decorative house plant)

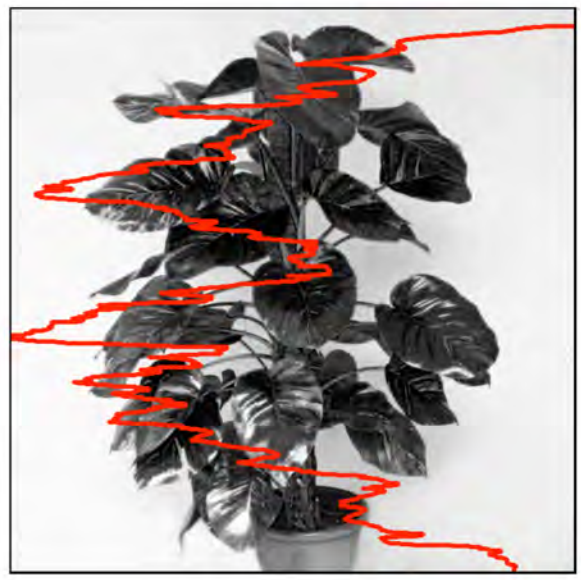

(Top-left) Original image.

(Top-right) B-image, provides best contrast.

(Centre-left) [TR] after thresholding.

(Centre-right) Principal axis and contour showing row-by-row centre point of the leaf cover. (Ideally, this should be a straight vertical line.)

(Bottom-right) Convex hull of [CL]. The ratio of the area of the leaf cover, compared to that of the convex hull is a measure of “solidity”.

(Bottom-right) Graph showing the total amount of leaf cover in each row.

SEEDLINGS

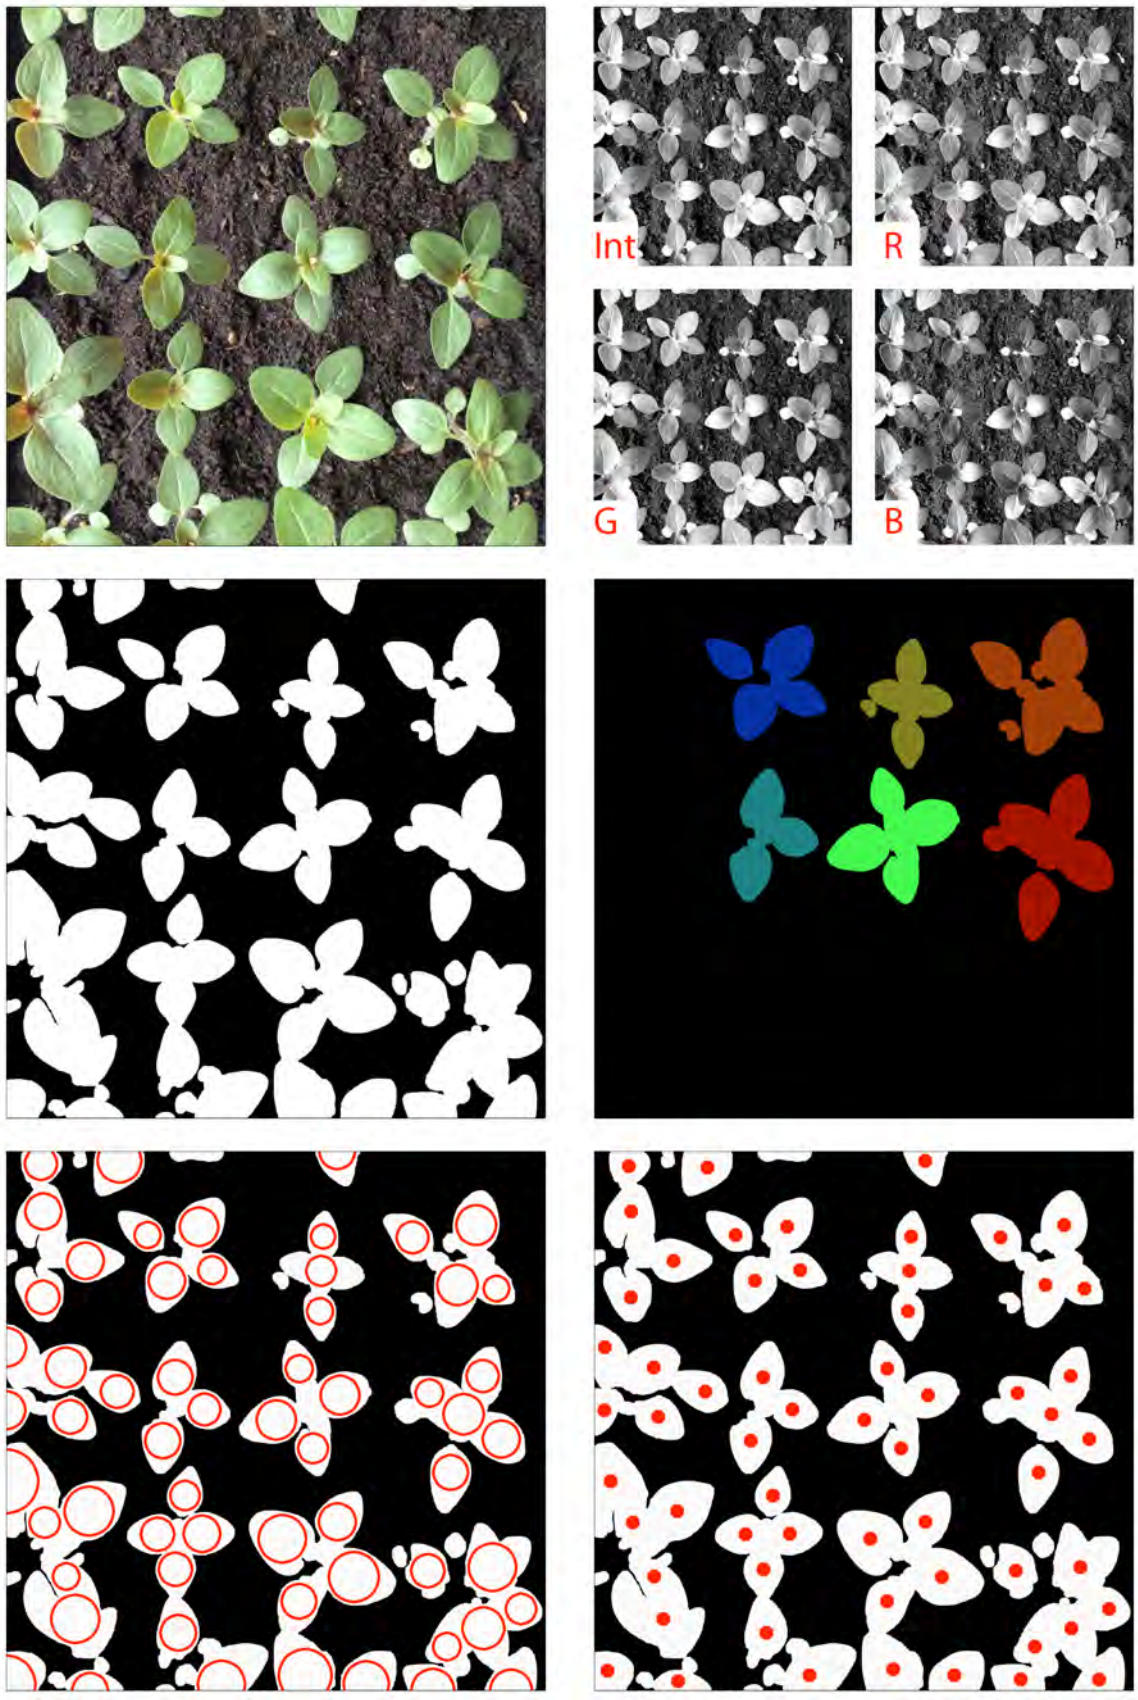

(Top-left) Original image.

(Top-right) Intensity and RGB component images.

(Centre-left) Output of a colour filter designed to detect leaf-green.

(Centre-right) When the seedlings are well separated, or overlap slightly, it is possible to locate and analyse each plant individually.

(Bottom-left) A program recursively places circles into the white areas. These circles are crude indicators of leaf size.

(Bottom-right) Centres of the circles. These are included here to emphasise

MONITORING FLOWER DEATH

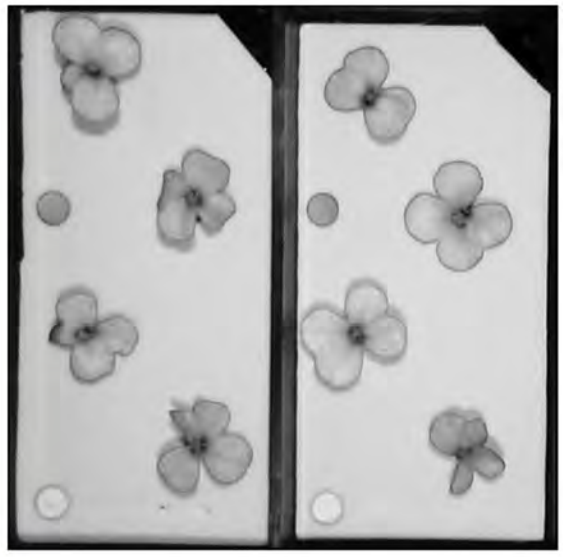

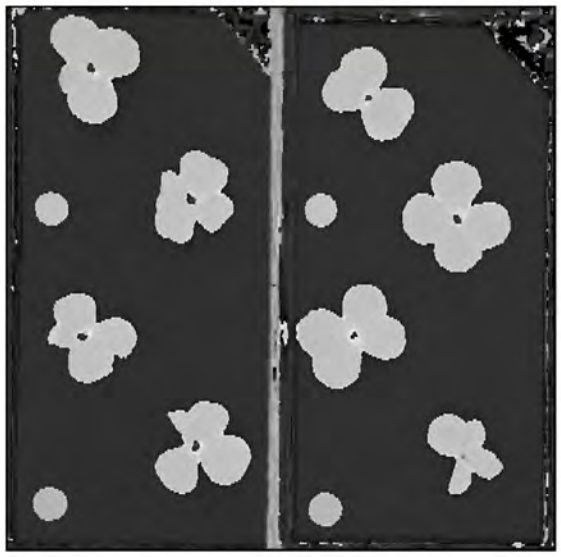

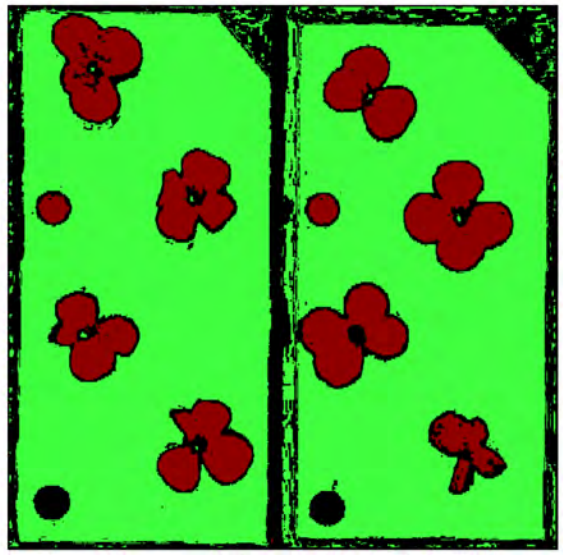

Identifying wall-flowers using a colour-recognition filter. Image [TL] was created during a research study of flower-death processes, by Hilary Rogers & Andrew Harrison, Cardiff University.

(Top-left) Original image. The small round spots are coloured paper markers, used for image registration.

(Top-right) Intensity. {I-image)

(Centre-left) Hue. (H-image)

(Centre-right) Colour scattergram in the HS plane.

(Bottom-left) Colour map, obtained by processing

TREE GROWTH RINGS

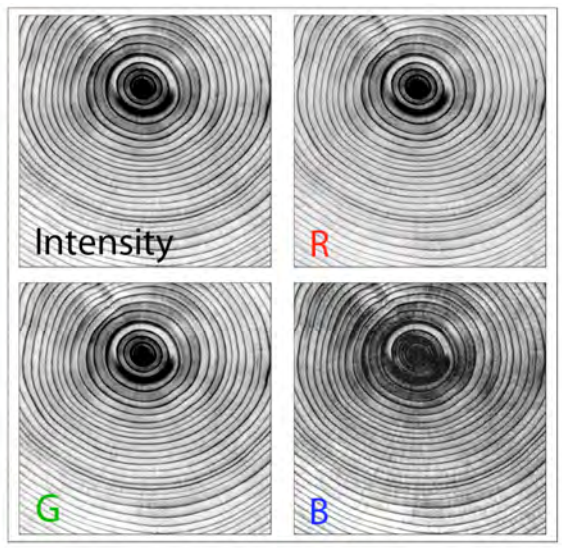

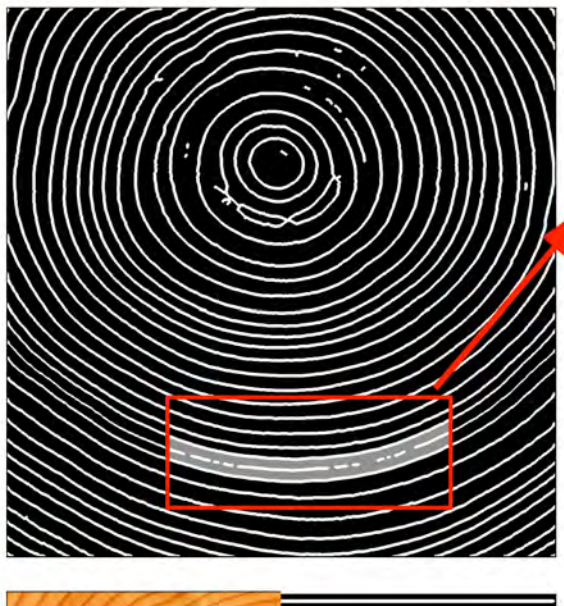

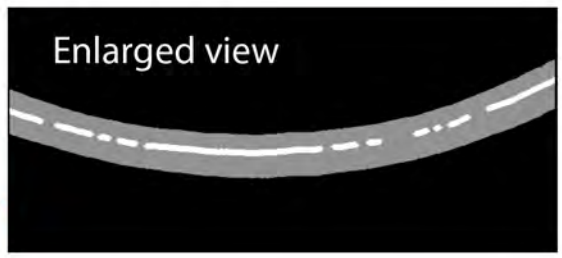

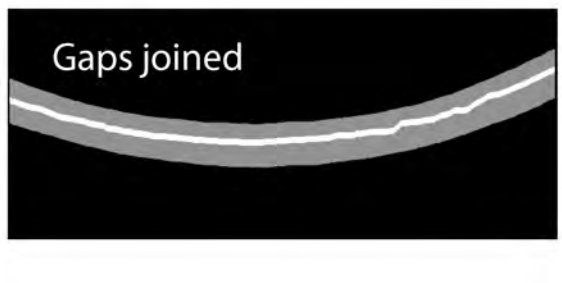

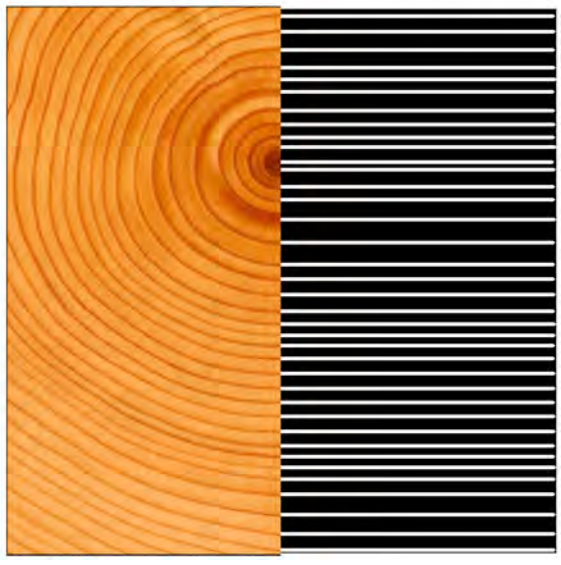

(Top-left) Original image.

(Top-right) Intensity and RGB component images.

(Centre-left) After filtering, thresholding and thinning. The broken contour in the red box is shown enlarged in [BR].

(Centre-right) Joining gaps.

(Bottom) One white bar has been found for each growth ring. This creates a "bar code" that can be compared with other samples for dendrochronology

Section 3

FOOD &
FOOD PRODUCTS

PORK PIE (checking fat content)

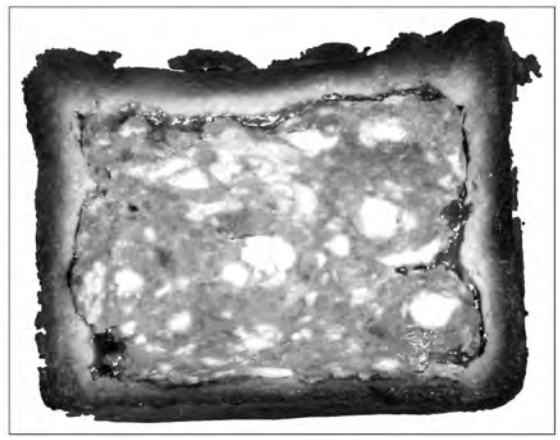

(Top-left) Original image.

(Top-right) Intensity image (I-image)

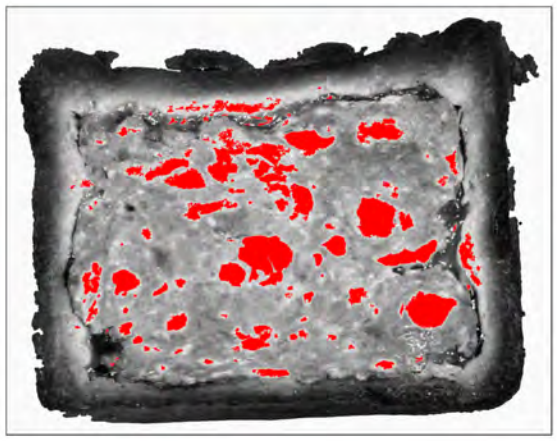

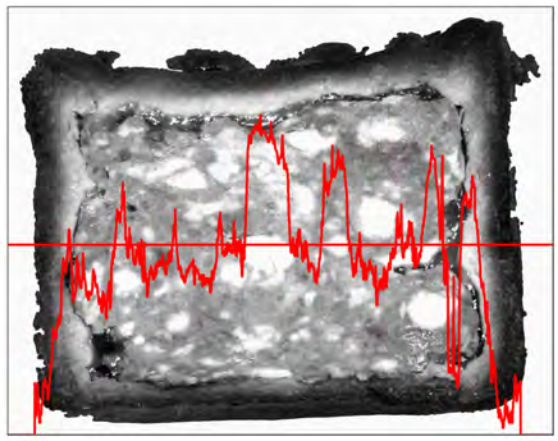

(Centre-left) Bright spots fat) in [TR] detected.

(Centre-right) Intensity profile along the horizontal red line.

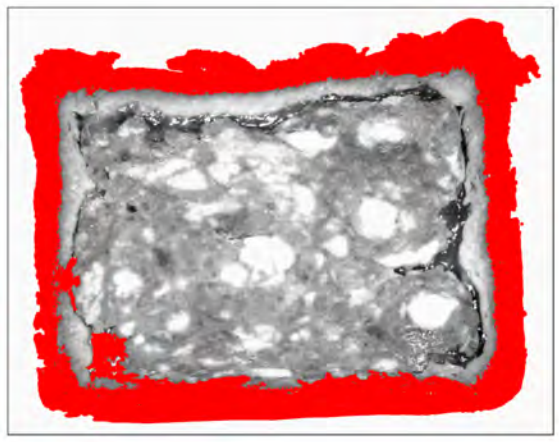

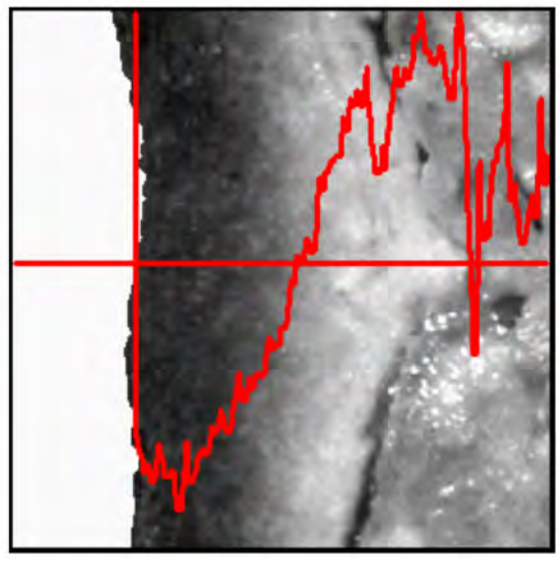

(Bottom-left) Outer crust

(Bottom-right) Expanded view of the edge of [TR]. The intensity profile shows how the fat is penetrating into the pastry crust..

PROCESSED MEAT (checking fat content)

A quick web search shows that there that there are several different approaches to measuring the fat content of meat. These can be based on: visible light, x-rays (singe- and dual-wavelength), or microwaves. Even a standard colour image produces promises results.

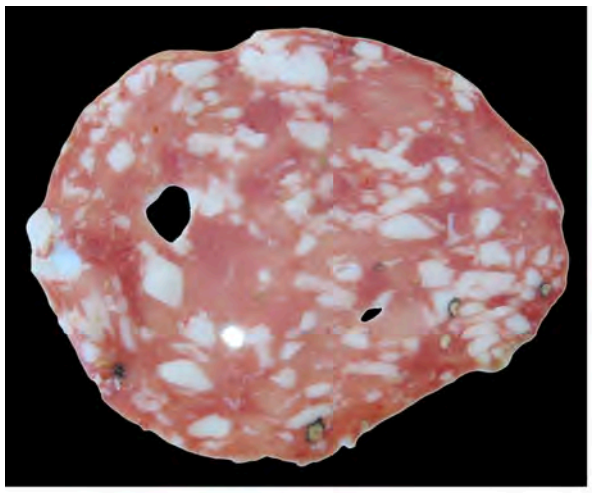

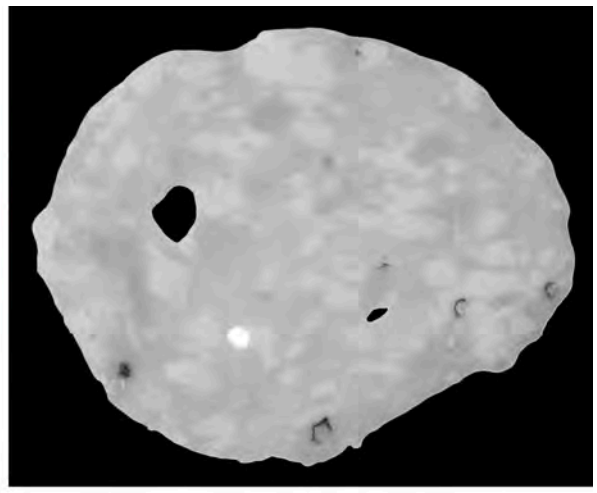

(Top-left) Original image.

(Top-right) R-image

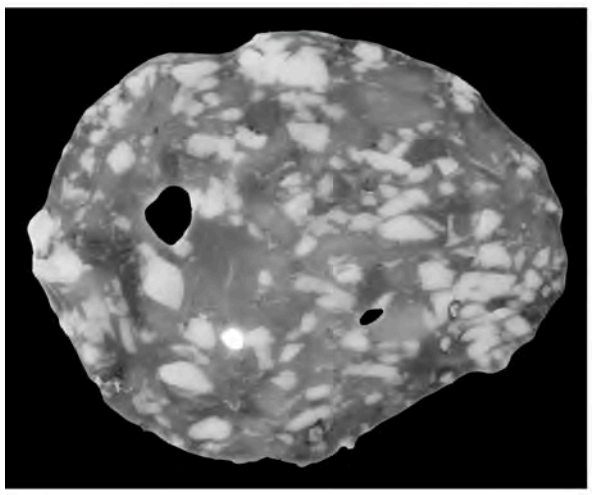

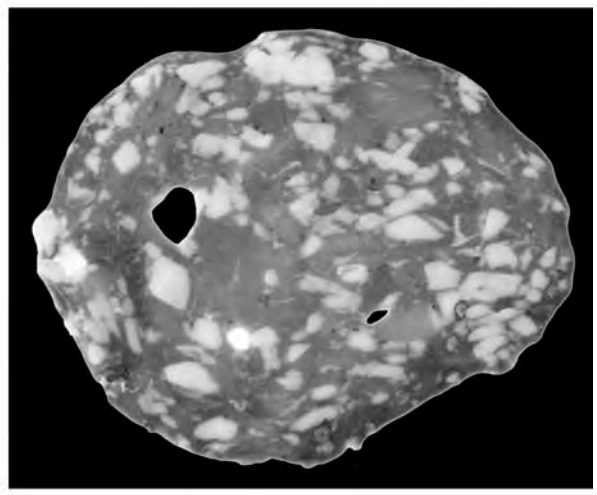

(Centre-left) G-image

(Centre-right) B-image

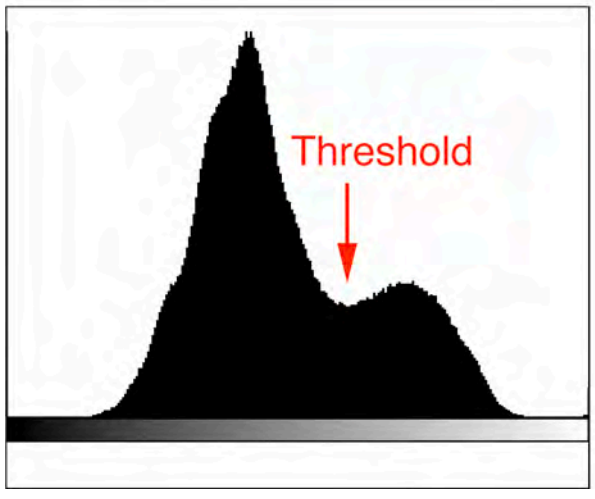

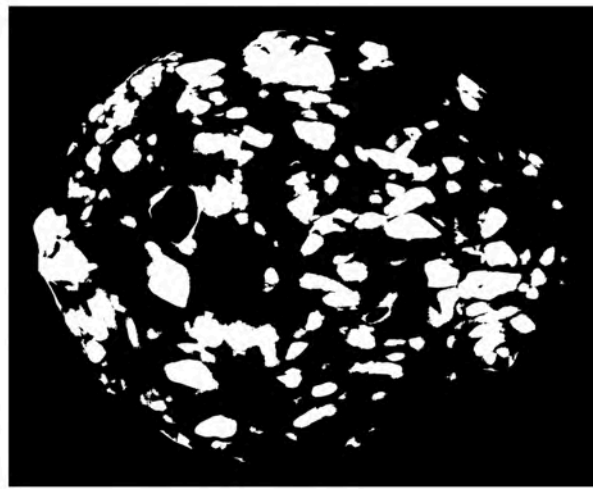

(Bottom-left) Intensity histogram of the B-image

(Bottom-right) [CR] thresholded as indicated in [BL].

LOAF SHAPE - 1

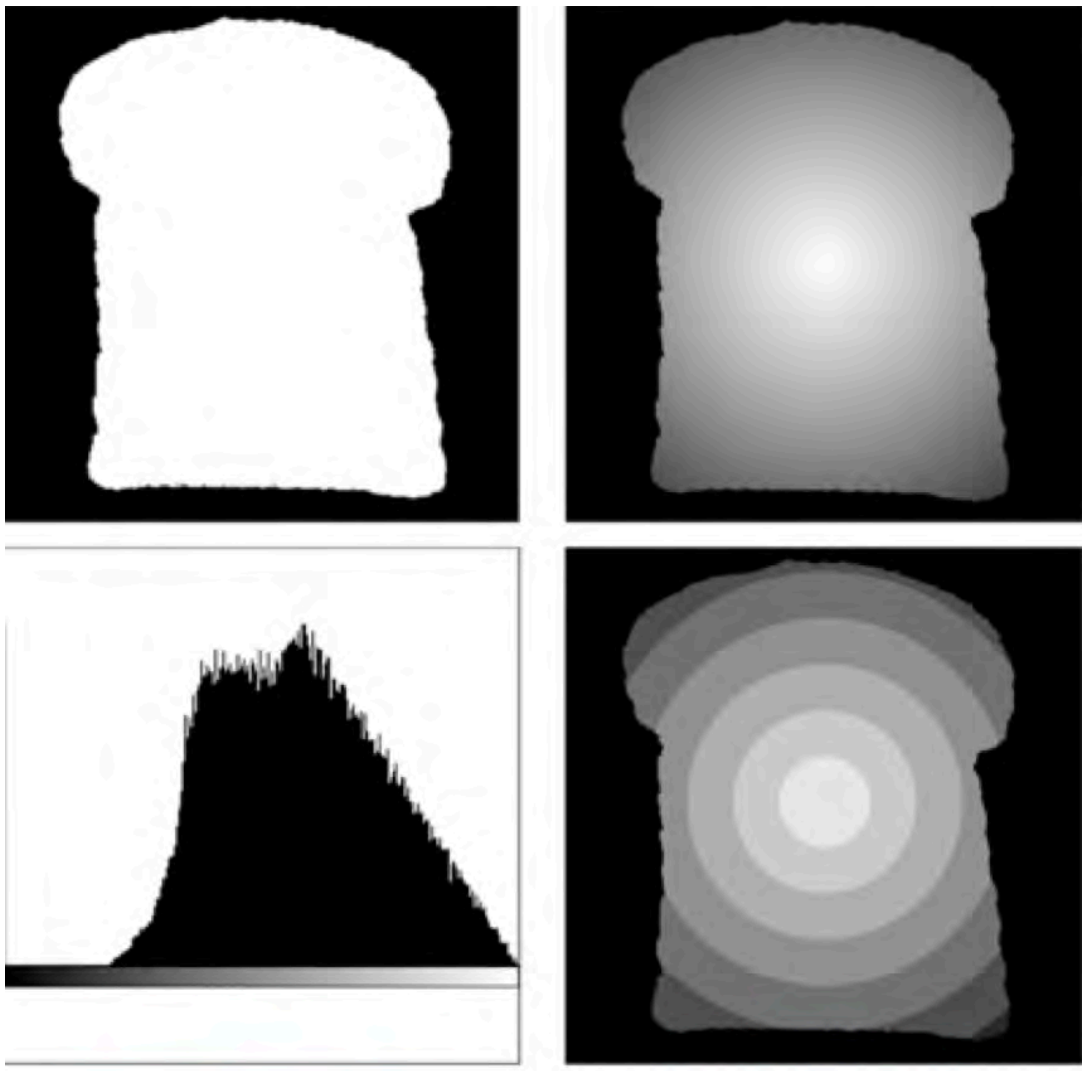

(Top-left) Original (binary) image.

(Top-right) "Intensity cone". The brightest point of the "cone" was placed at the centroid of the slice.

Centre-left) Intensity histogram of [TR]. Notice that this is independent of orientation of the slice.

(Centre-right) As [TR] but with a reduced number of grey levels. The intensity histogram of this image is still independ

ent of rotation.

(Bottom-left) Grassfire transform, with a reduced number of grey levels. The intensity histogram is independent of rotation.

(Bottom-right) Circular wedge, with a reduced number of grey levels. This method can easily be extended to make it independent of rotation.

LOAF SHAPE - 2

(Top-left) Minimum-area rectangle.

(Top-right) Circumcircle..

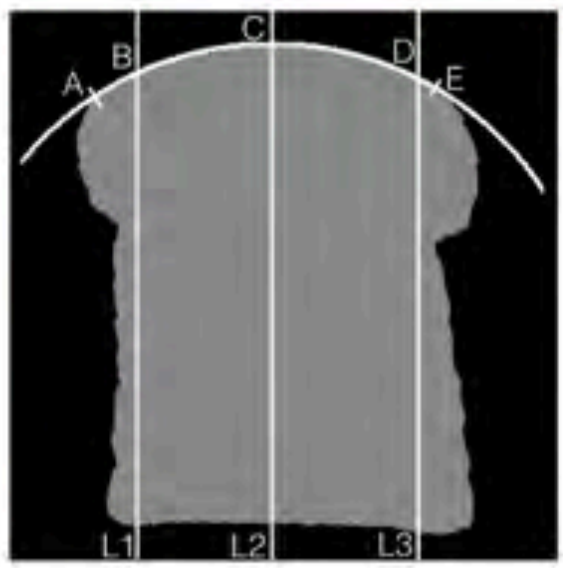

(**Centre-left)** Convex hull. The two coloured regions are bays that should have approximately equal areas.

(Centre-right) Circle fitted to the top of the slice, fits accurately between points A and E.

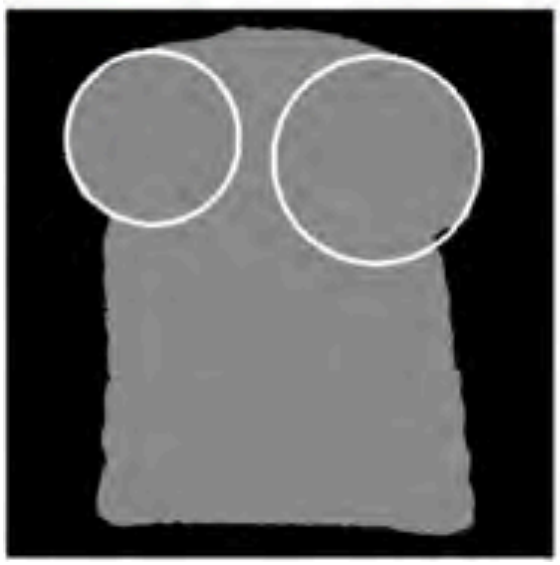

(Bottom-left) Circles fitted to the overspill areas.

TEXTURE (bread crumb)

It is possible to generate a large number of crumb-texture measurements, based on a combination of grey-scale filtering, binary-morphology operators and thresholding. The five binary images shown here were generated by varying only the grey-scale filtering. Thresholding was then applied, so that 50% of the pixels are white. (Edge effects are ignored.) We can then apply several directional binary morphology operators to each of these. Finally, a multi-element measurement vector is formed, simply by counting the blobs in these images. A learning program can then be trained to detect “unsatisfactory” textures.

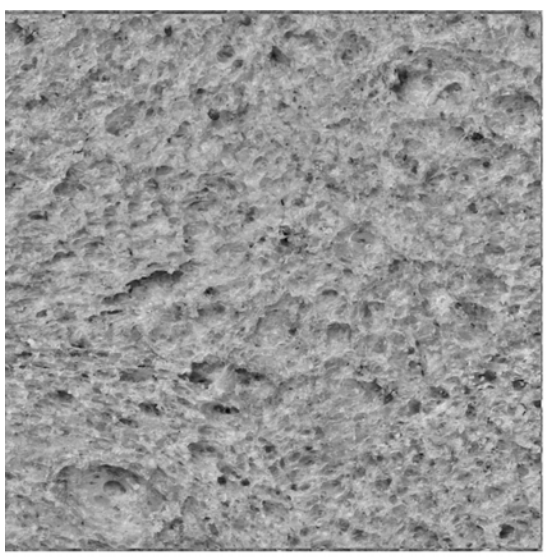

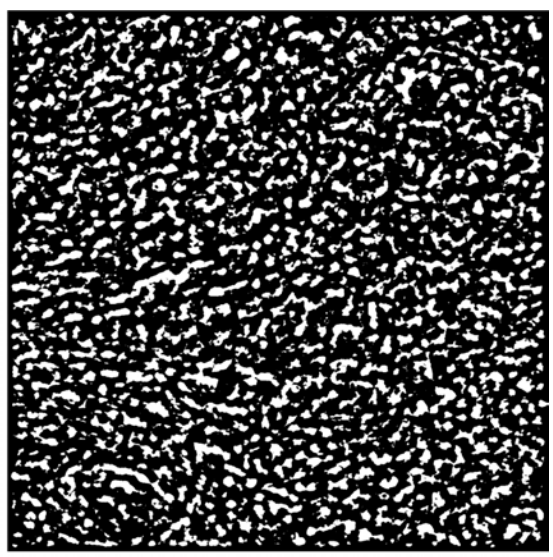

(Top-left) Original image, obtained from a machine-cut bread slice. Uniform front lighting was used. Grazing illumination can provide higher image contrast.

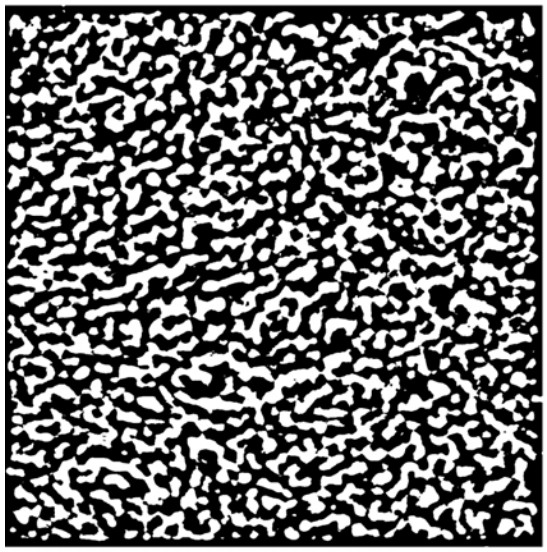

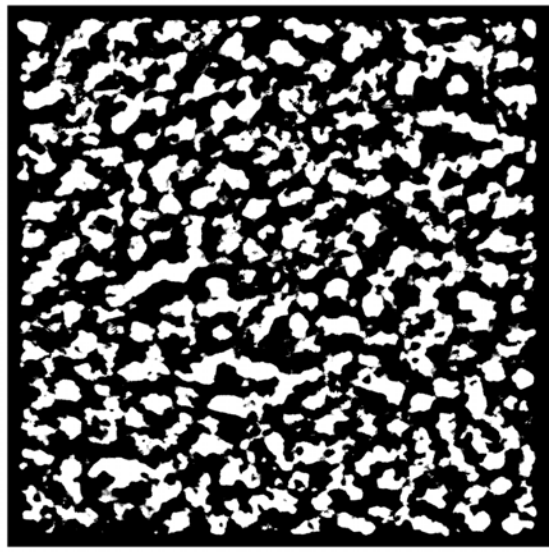

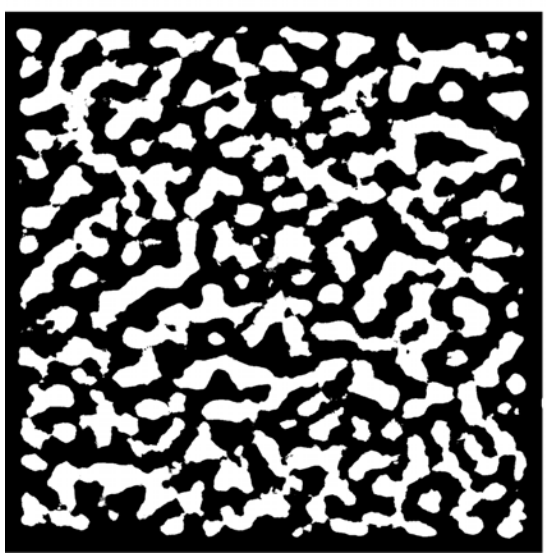

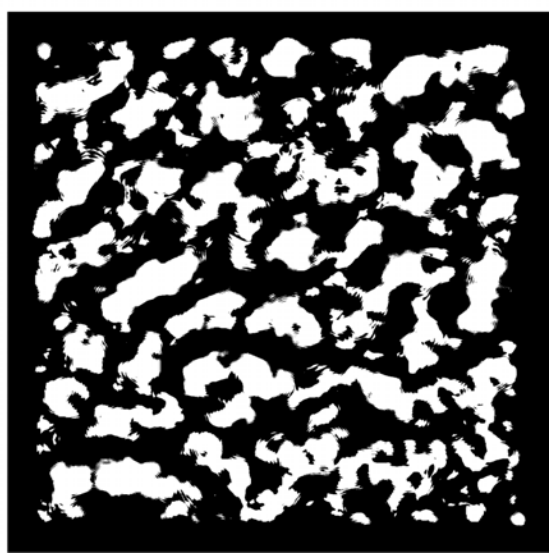

(Others) Binary images with varying levels of grey-level morphological filtering, followed by thresholding so that 50% of each image is white (ignoring edge effects).

LOAF SHAPE (projecting light stripes, "structured light")

A series of parallel light stripes is projected onto the surface of the loaf. (They can be coloured to make it easier to separate them and to join “broken” stripes.) The camera views the projected pattern from an oblique angle. The surface height follows from a simple calculation.

(Top) Optical arrangement.

(Bottom) Left-to-right: (i) Side view: (ii) Top surface; (iii) Light stripes projected onto the top surface; (iv) After processing.

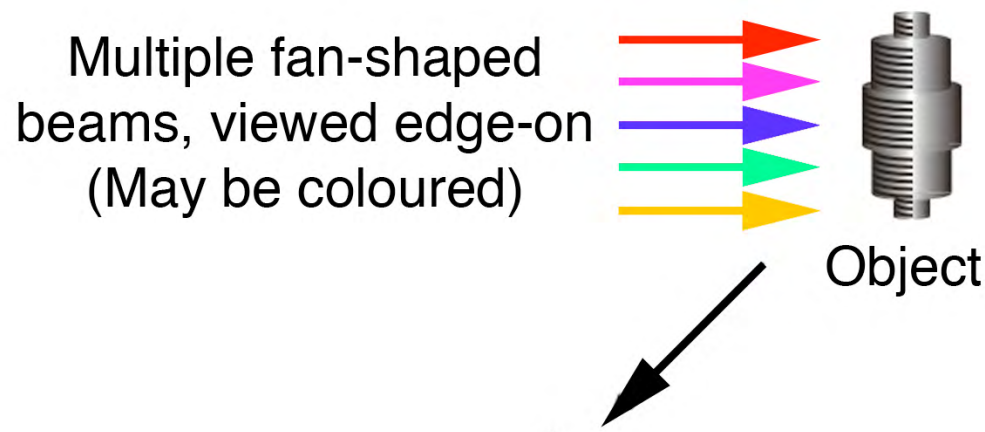

ESTIMATING SIZE (peas)

(Top-left) Original image.

(Top-right) Background intensity variations eliminated, by subtracting a blurred version of the (grey-scale) image from itself.

(Bottom-left) Thresholding [TR] so that 50% is white. (Minor noise-removal filtering as also been applied.)

(Bottom-right) Red spots indicate the centres of the blobs in [BL].

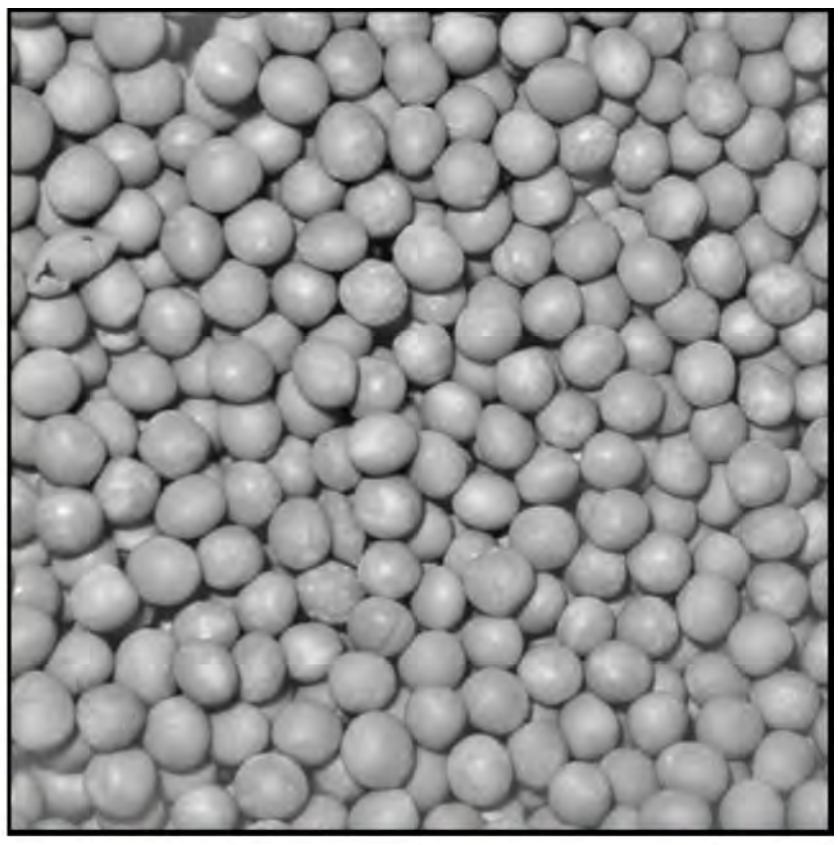

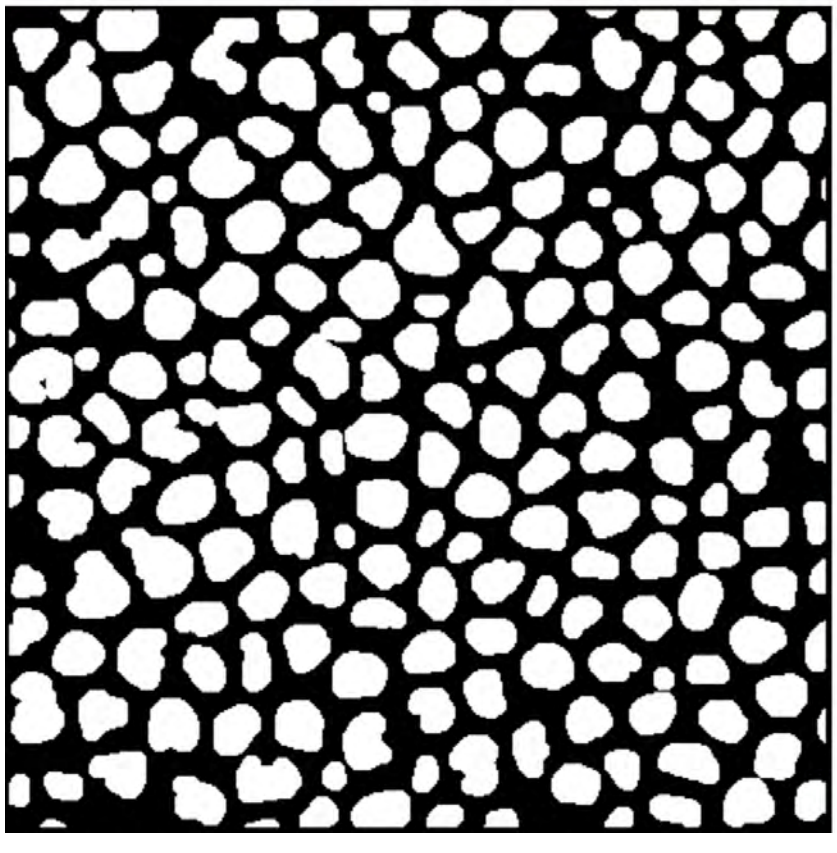

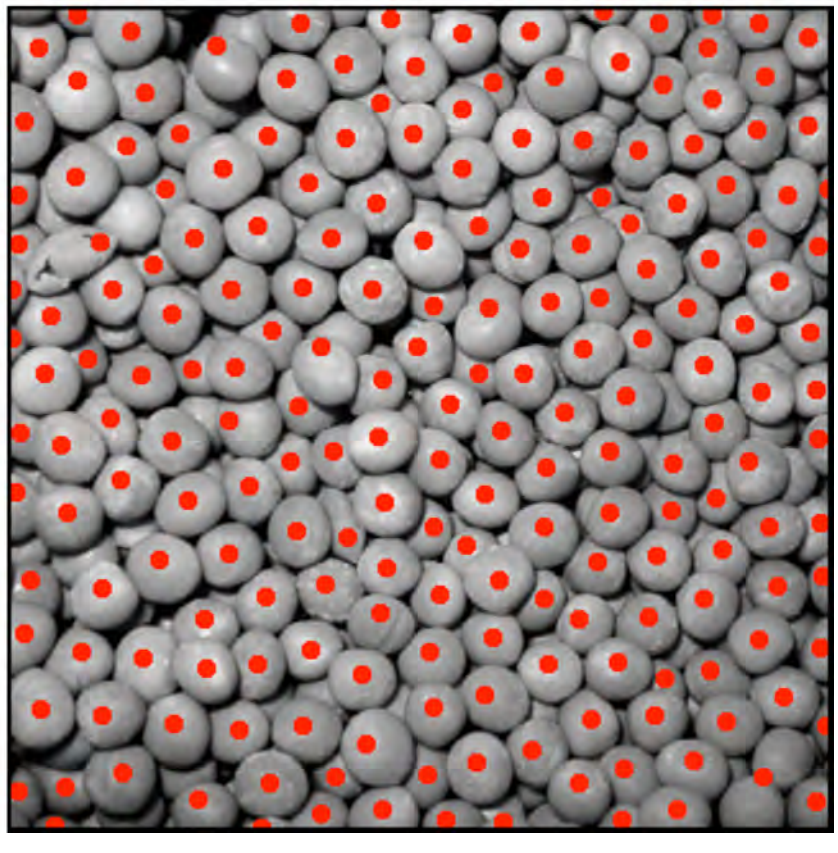

BONES IN CHICKEN MEAT (x-ray image)

(Top-left) Original image Notice the low-contrast bone, bottom-left. [Image supplied by Dr. Mark Graves.]

(Top-right) Contrast enhanced.

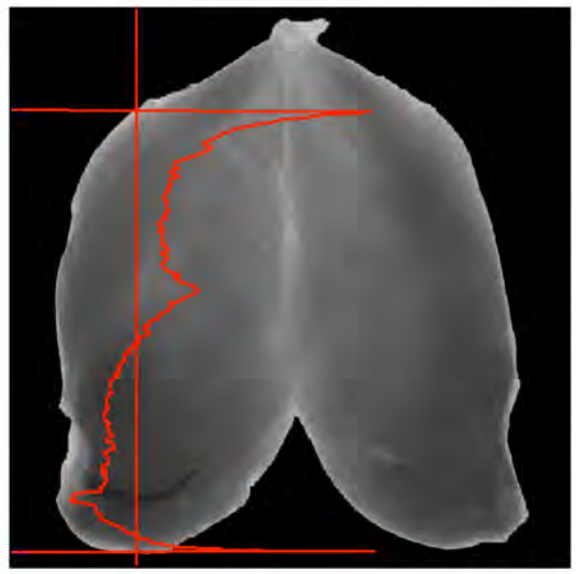

(Centre-left) Binary image.

(Centre-right) Intensity profile along the vertical red line.Notice the low-contrast bone.

(Bottom- left) “Crack detector” filter applied to [TL].

(Bottom-right) Bone detected successfully by eliminating small blobs from [BL].

AUTOMATED HARVESTING OF MUSHROOMS

(Top) Original image.

(Bottom) Mushrooms detected. Notice the false “hit” (green cross).

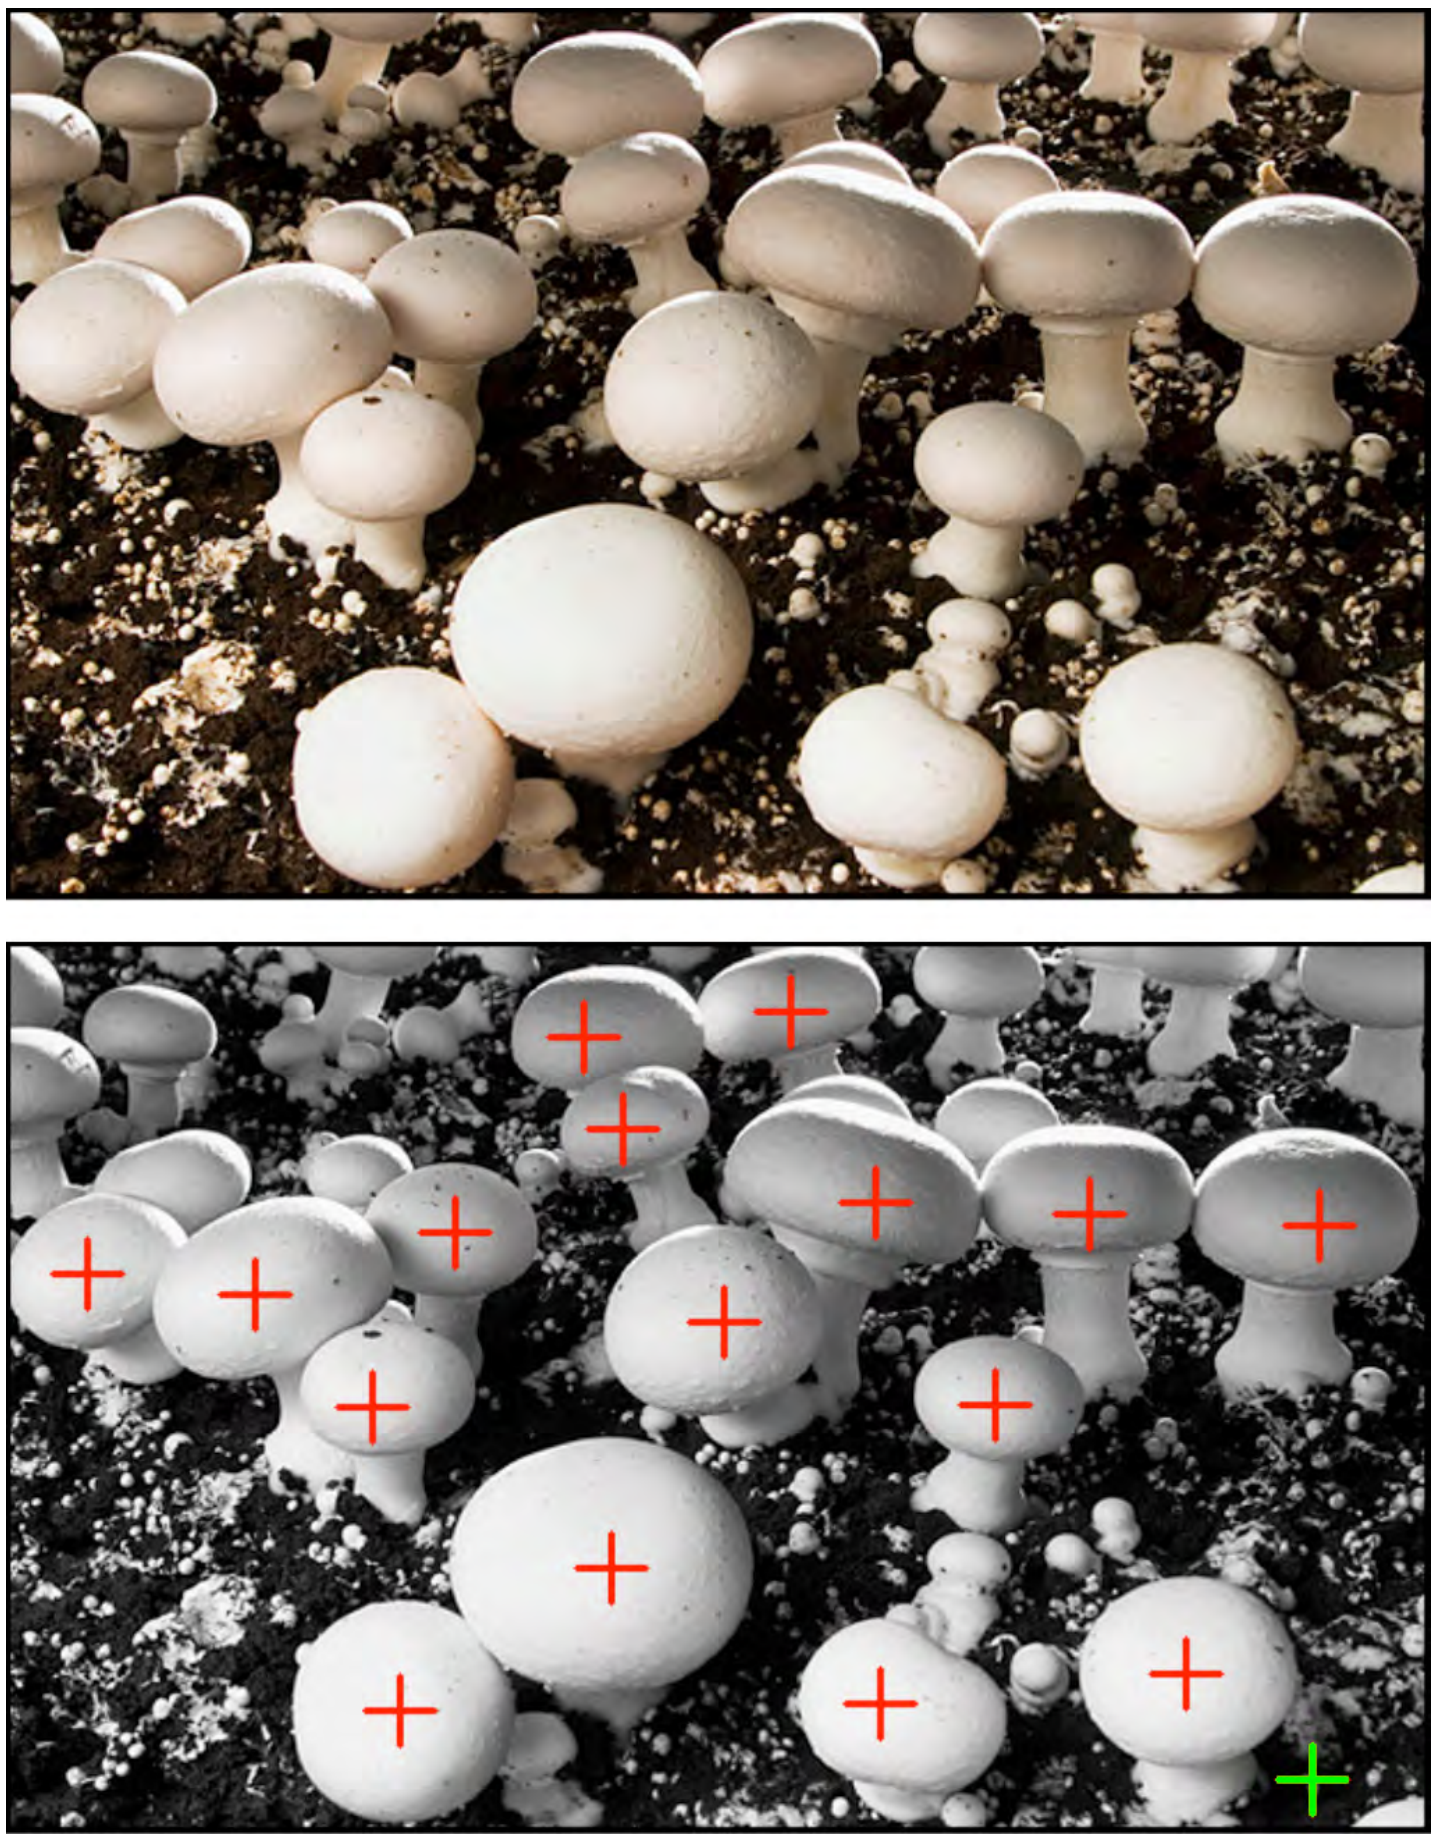

GRADING FRUIT BY SHAPE (curvature of a banana / cucumber)

Bananas are classified by quality and size so they can be traded internationally. EU Regulation 2257/94 states that bananas must be "*free from malformation or abnormal curvature. Class 1* allows"*slight defects of shape"; Class 2* allows "*defects of shape".* Visit

www.europarl.europa.eu/unitedkingdom/en/media/euromyths/bendybananas.html

(Top) Casual visual examination is not an adequate basis for international trade; objective criteria are needed to avoid and resolve disputes.

(Bottom-left) Eliminating the ends of a banana, using morphology, to estimate the extent of the edible part of the fruit. (Could this form the basis for estimating its weight?)

(Bottom-right) Circles fitted to the “inner” and “outer”edges and the central “spine”. The radii of these circles effectively determine the curvature, over the edible part of the fruit.

A slightly different method would be required for the cucumber.

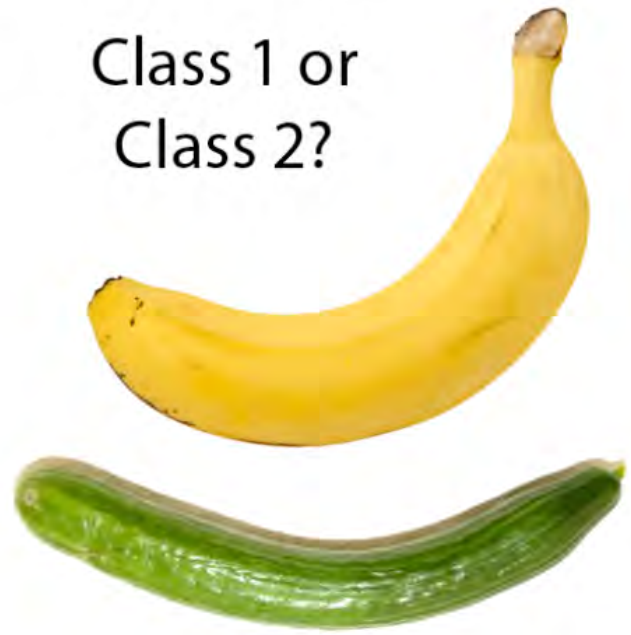

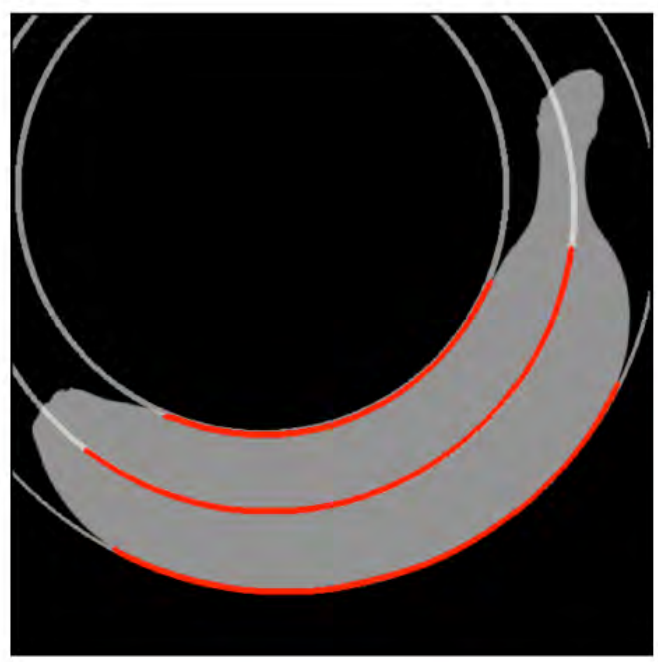

GRADING FRUIT (bruising, using IR sensing)

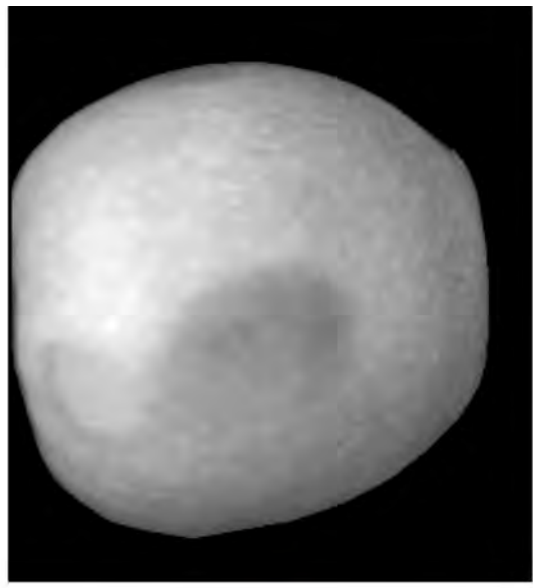

(Top-left) Apple, viewed under white light.

(Top-right) Same apple using near IR sensing. The dark patch indicates bruising

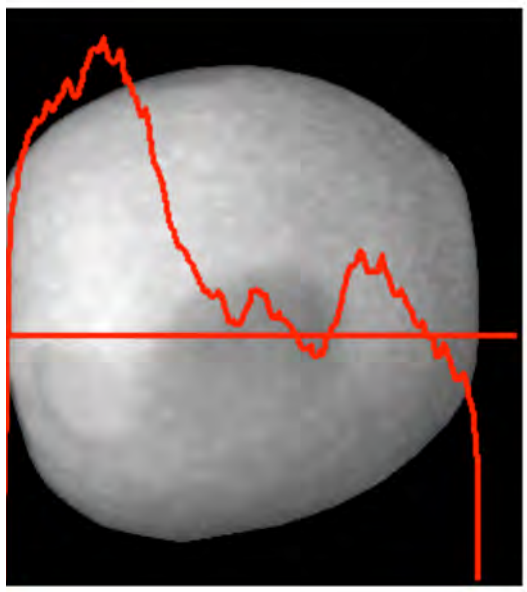

(Centre-left) Intensity profile along the horizontal red line.

(Centre-right) Intensity profile along the vertical red line.

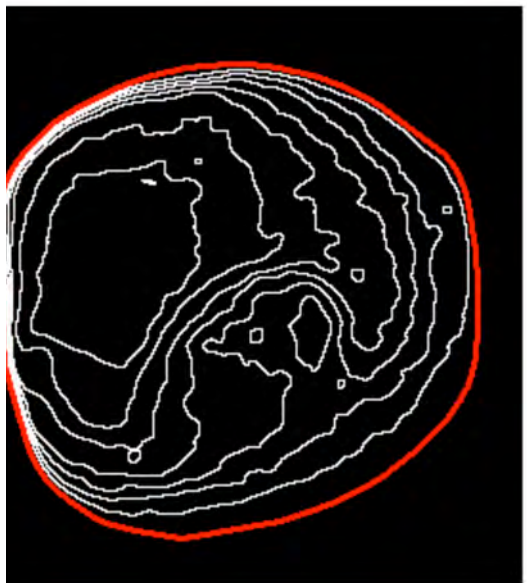

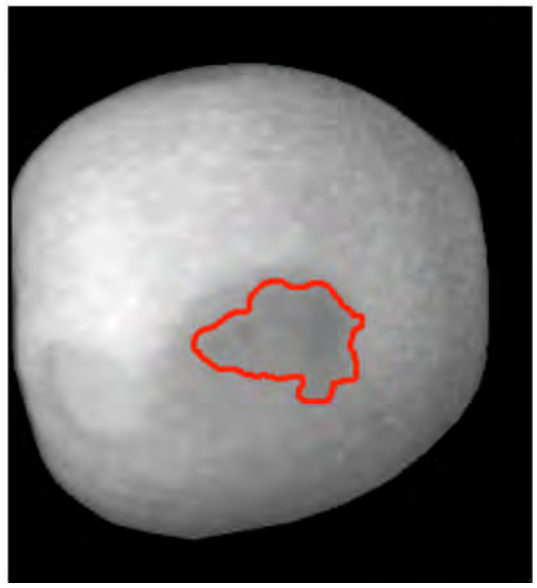

(Bottom-left) Ispohotes in [TR].

(Bottom-right) Outline of the bruised area, obtained

CAKE DECORATION PATTERNS (continuous extruded strip-line products)

Original (binary) image. All crossing points should lie within the green stripes and the centroids of the lakes should all lie within the red stripes.

Detecting diagonals. The red arrow indicates a break in the white contour. Pattern narrowing.

Detecting lakes. The centroid of the “excess” lake (red cross) lies outside the red stripe. Pattern thickening.

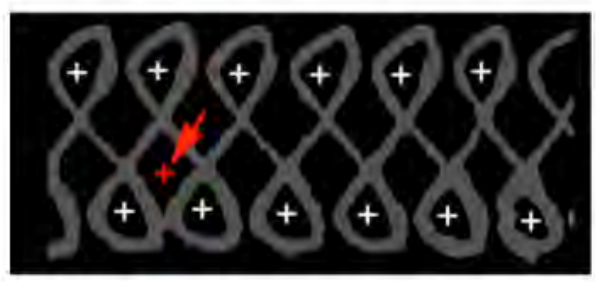

Detecting narrow gaps (white crosses) Pattern thickening.

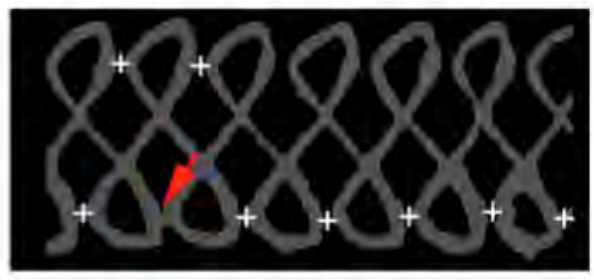

Detecting crossings. The structuring element (inset) was derived from one of the crossings after thinning. (Red arrow)

Identifying small background regions. The structuring element (red diamond) is too big to fit into one of the “coves”. (Red arrow).

FISH FILLETS (x-ray image)

Using x-rays allows the thickness and total weight to be measured, as well as allowing bones to be detected. The original image was kindly supplied by Dr. Mark Graves]

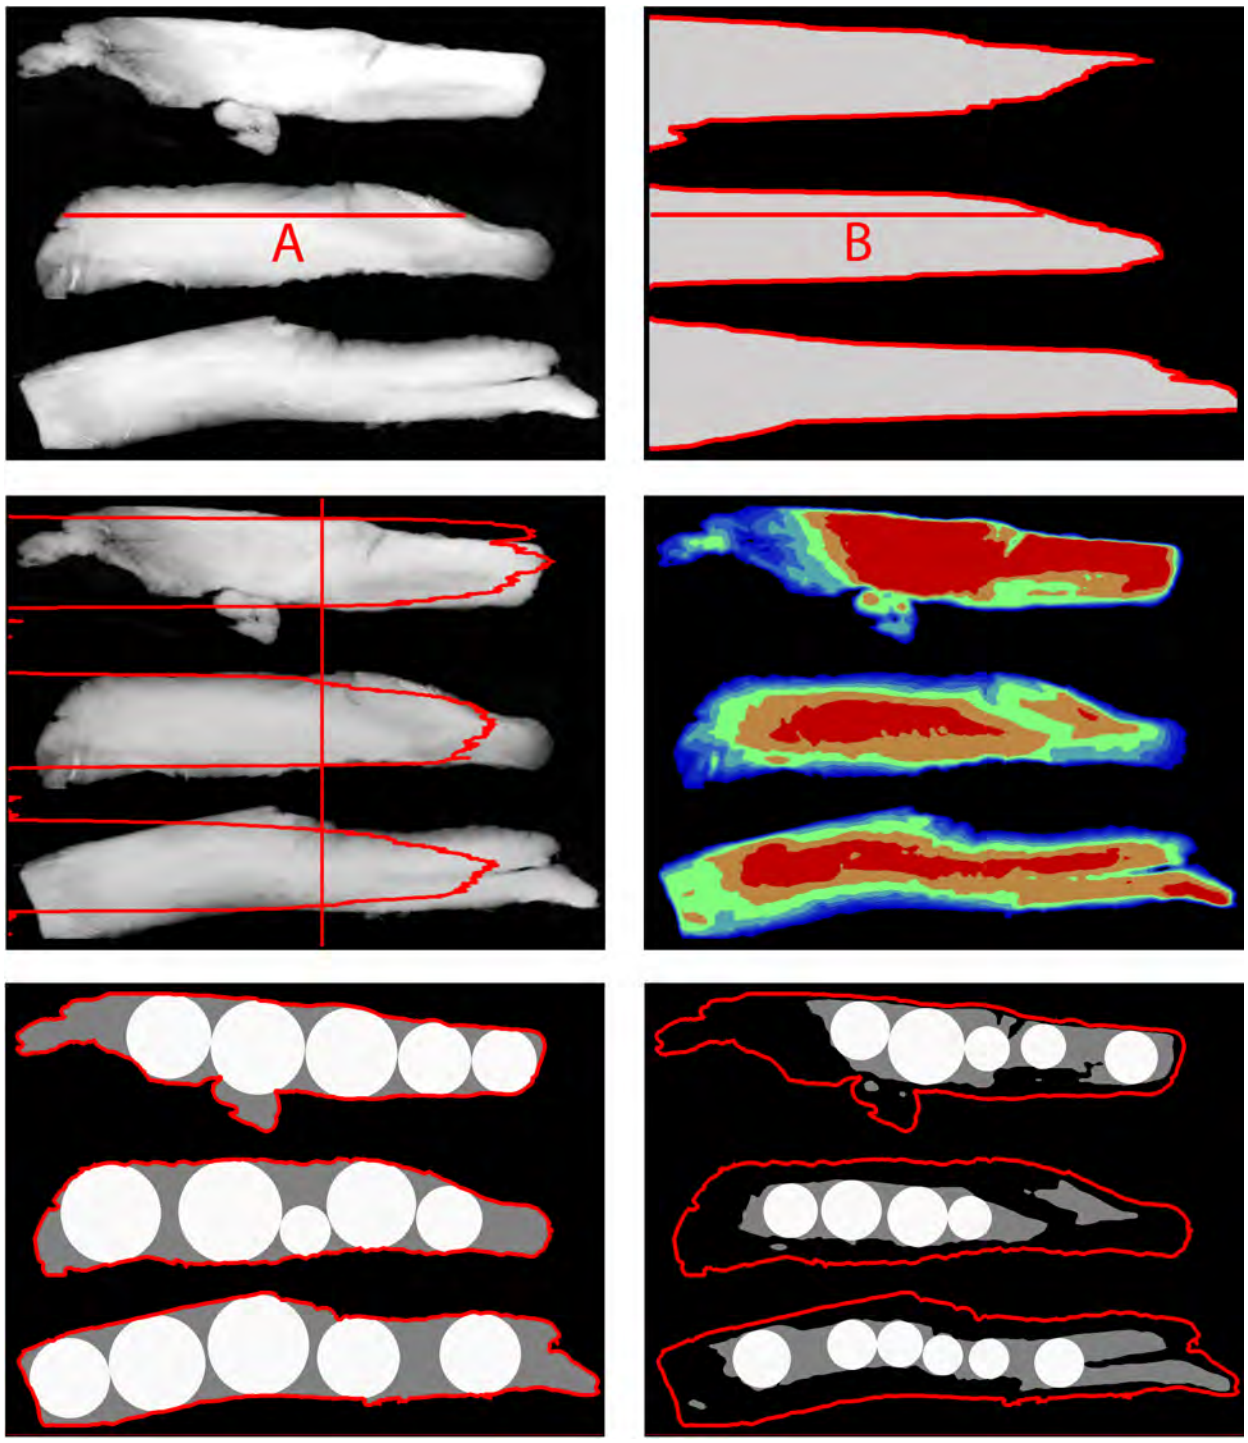

(Top-left) Original image. [

(Top-right) Calculating the weight. The length of line B is determined by the sum of all the intensities in row A in [TL]). The total weight is given by the area of the grey region.

(Centre-left) Intensity profile, indicates the thickness of the flesh.

(Centre-right) Imagine the fish being sliced progressively, like bacon. Colours indicate different levels of cut.

(**Bottom-left)** Simulation of robotic cutting with no regard for thickness.

(Bottom- right) Simulation of robotic cutting, limited to thick parts of the fillet.

MEAT & VEGETABLE PIE (Cornish pasty)

This sample application was inspired by a recent BBC Television programme: *"Inside the Factory"* (Series 5, Episode 6), showing a visually guided robot packing Cornish pasties. No details were given about how the system works but it is possible to make some reasonable deductions. Clearly, a light-stripe scanner was used. This does not depend on reflectance changes and produces a binary image, directly, without ever generating either a colour or grey-scale image. For this illustration, it was necessary to begin with a colour image

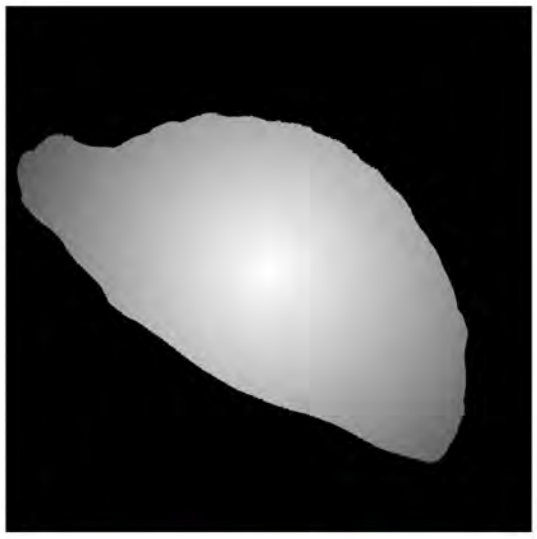

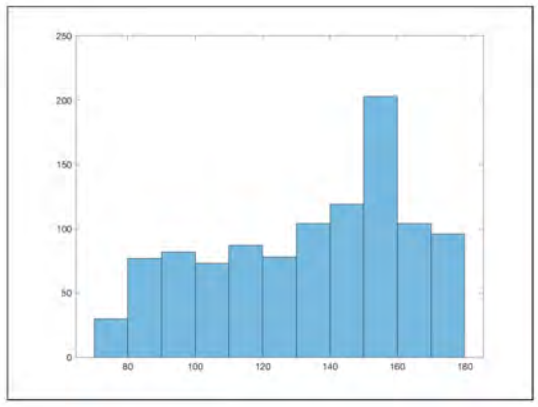

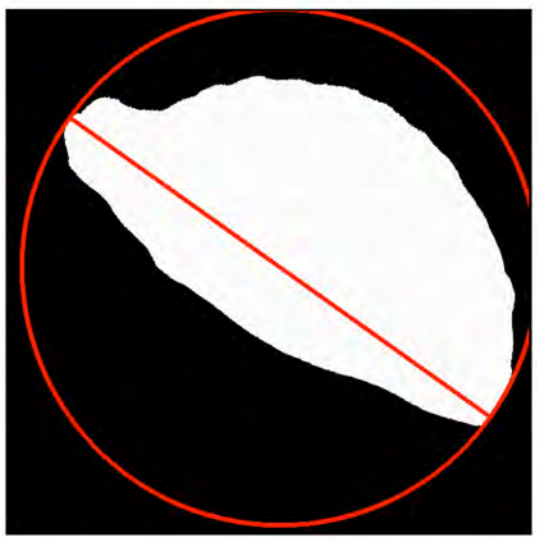

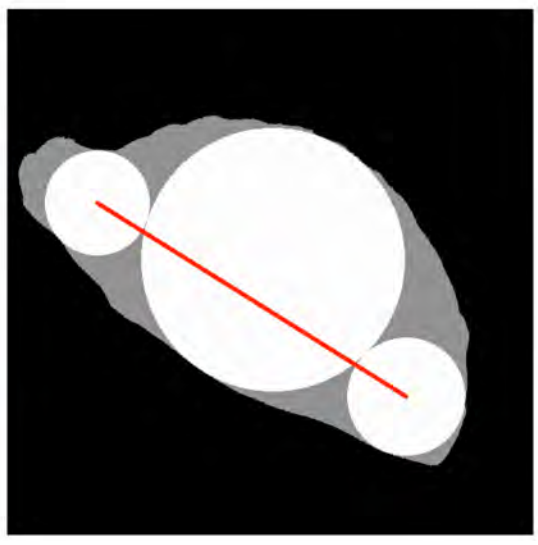

(Top-left) Original image

(Top-right) Binary image, derived from [TL].

(Centre-left) Intensity indicates distance from the centroid of [TR].

(Centre-right) Histogram of distances of edge pixels from the centroid. This is a description of the blob shape that is independent of orientation.

(**Bottom-left)** Line joining the two edge pixels that are furthest apart. This can be used to measure orientation.

(Bottom- right) Simple model of the blob using 3 largest circles drawn inside the blob. The red line can also define orientation.

X-RAY IMAGES

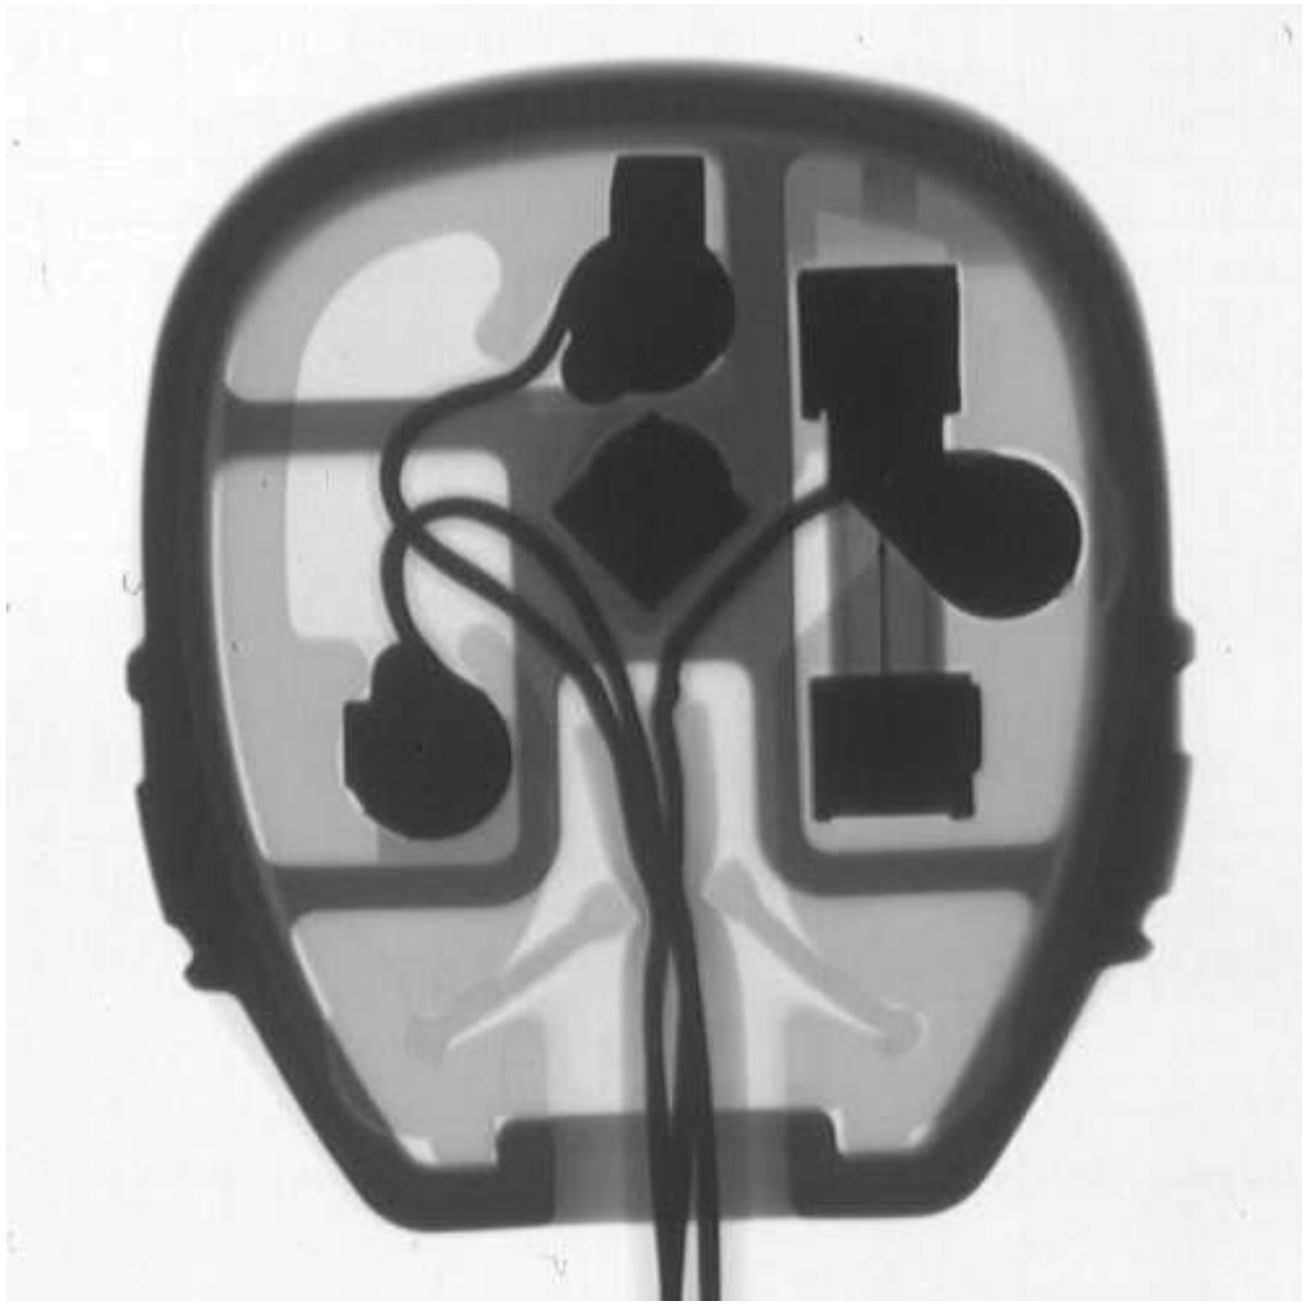

X-RAYS (badly wired UK plug)

(Left) Original image.

(Right) Loose strands, detected using a “crack detector” filter (grey-scale closing).

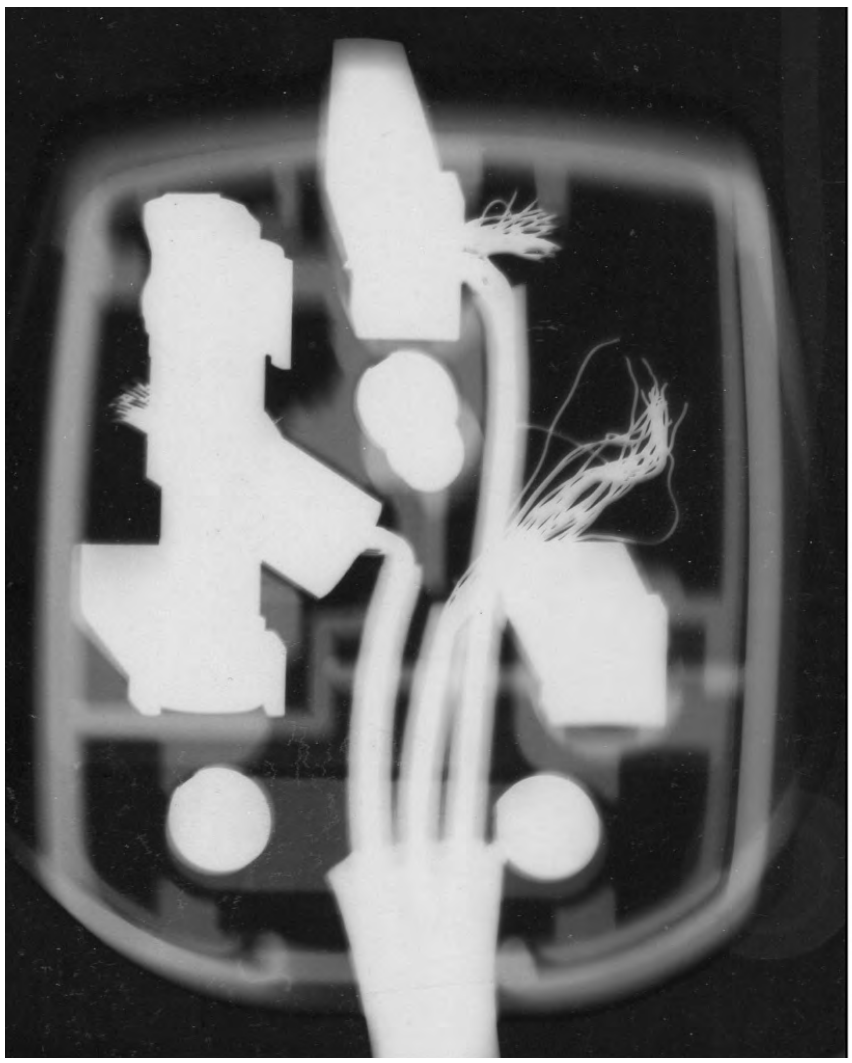

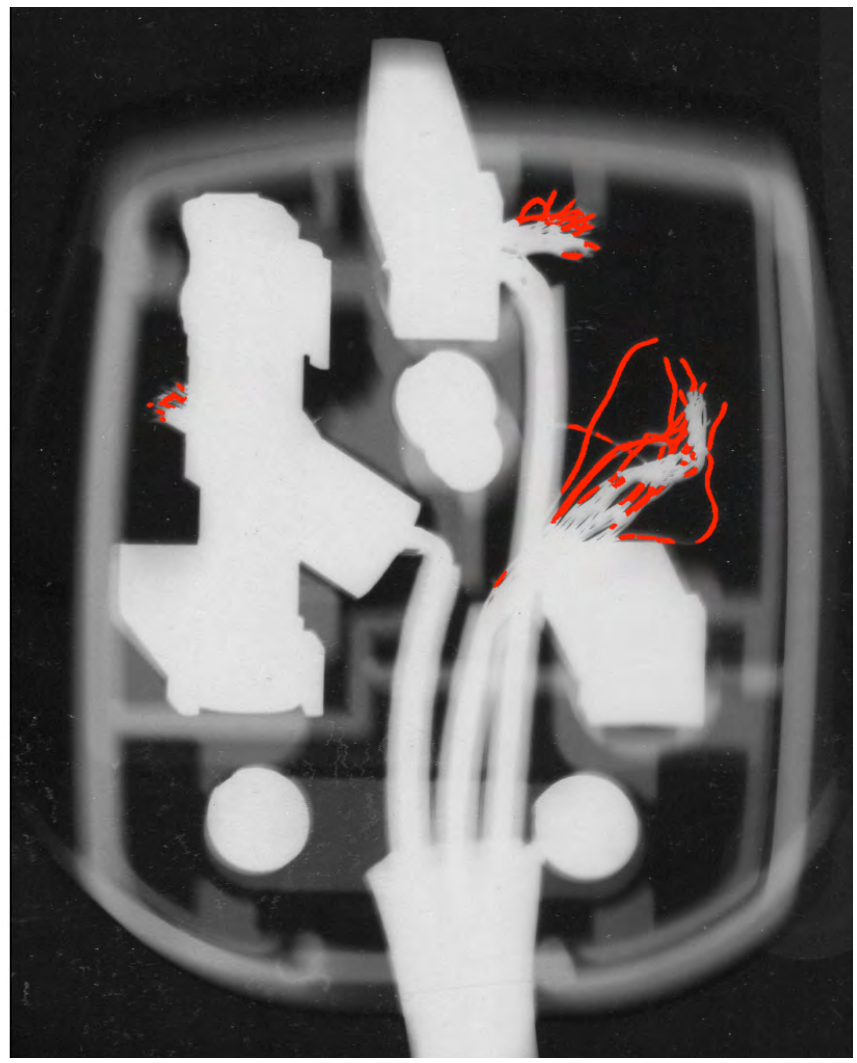

X-RAYS (tooth brush)

(Top-left) Original image. The dark rectangles are metal clips retaining the bristles.

(Top-right) Intensity histogram.

(Centre-left) Threshold between zero (black) and A, identifies the retaining clips.

(Centre-right) Threshold between A and B, reveals the retaining clips and the bristles.

(Bottom-left) Threshold between B and C.

(Bottom-right) Threshold between C and 255 (white).

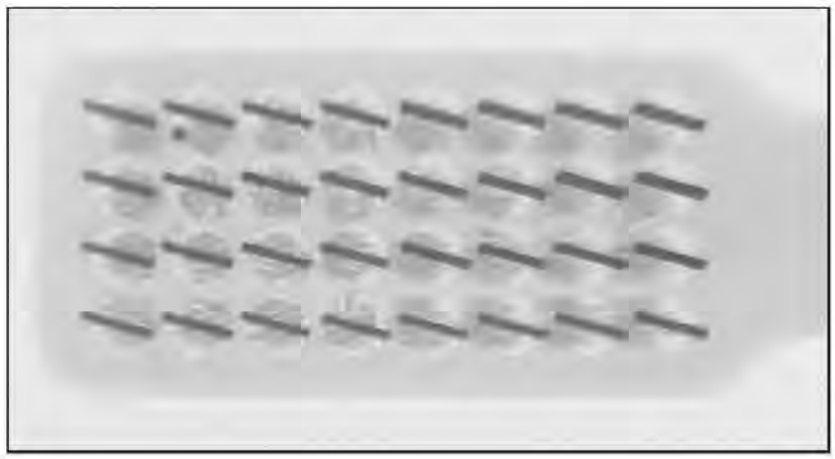

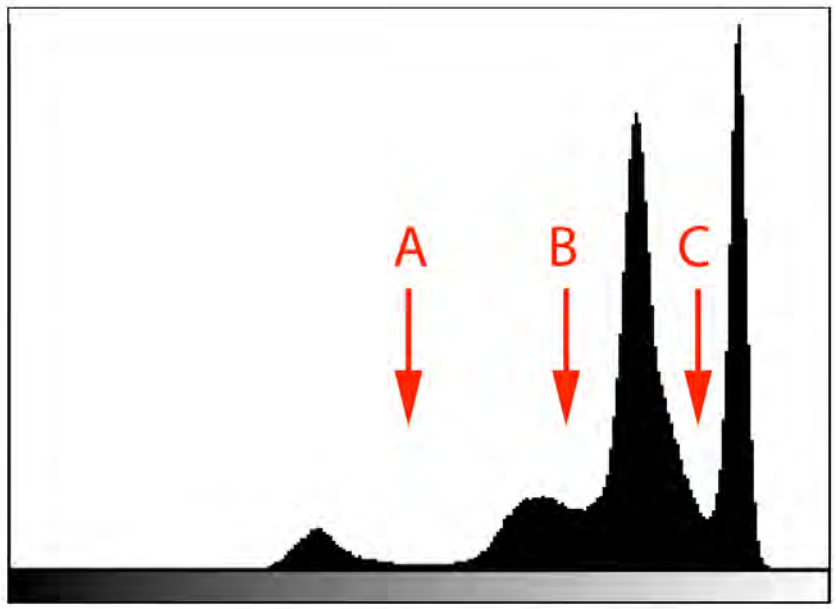

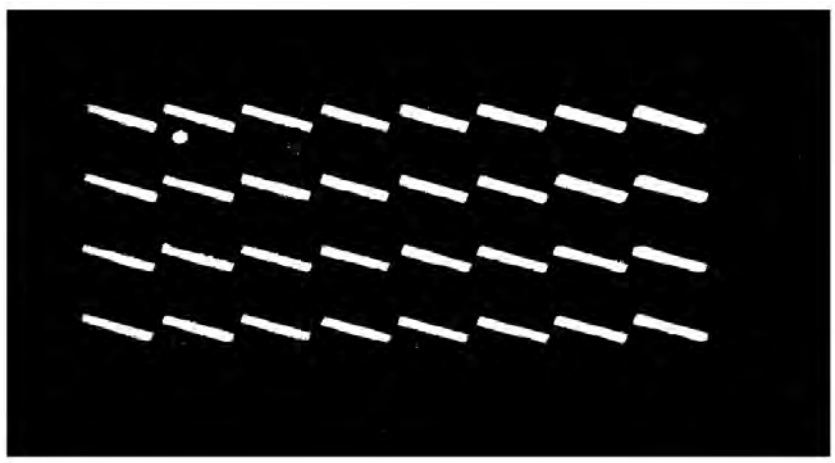

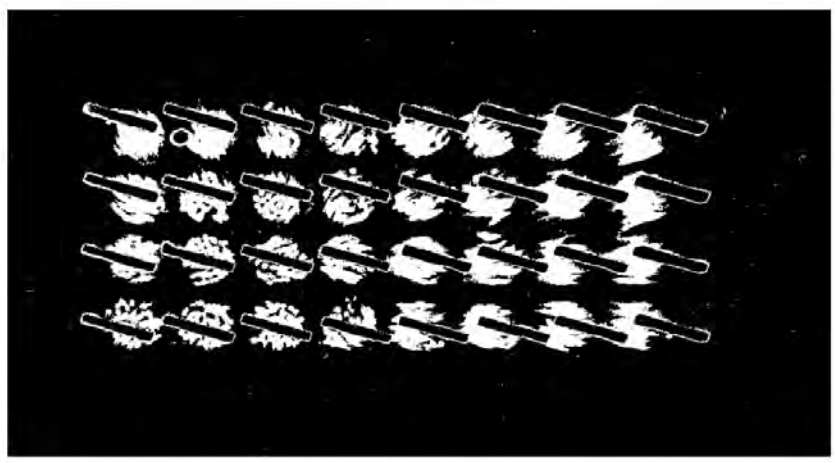

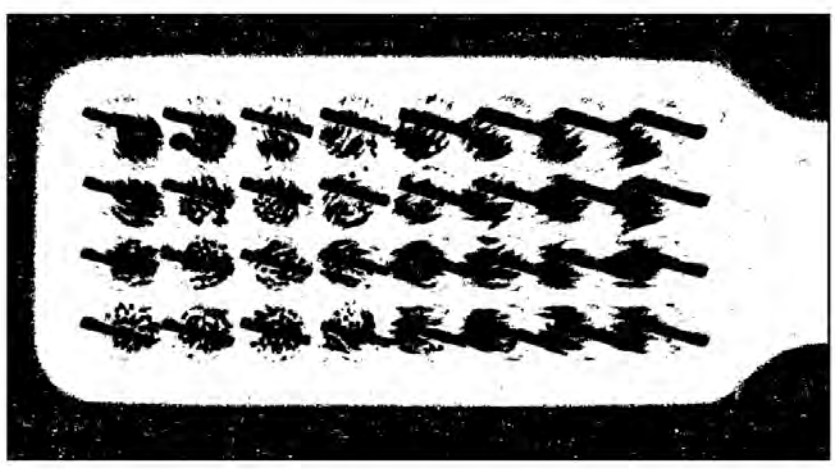

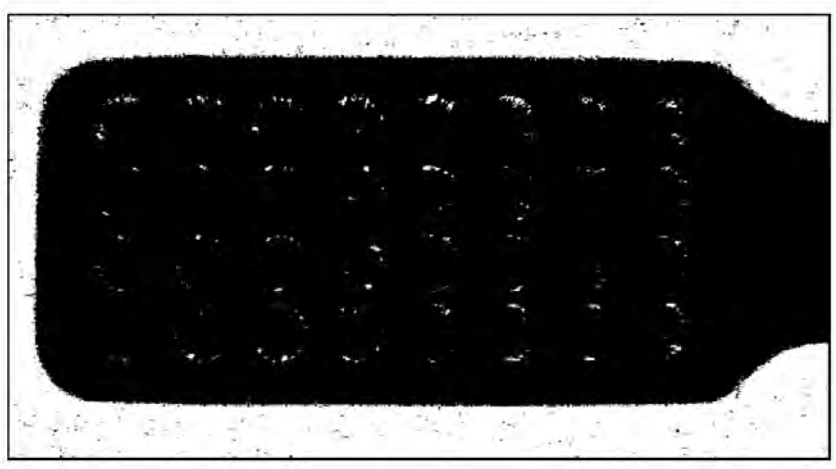

X-RAYS (micro-electronic circuit)

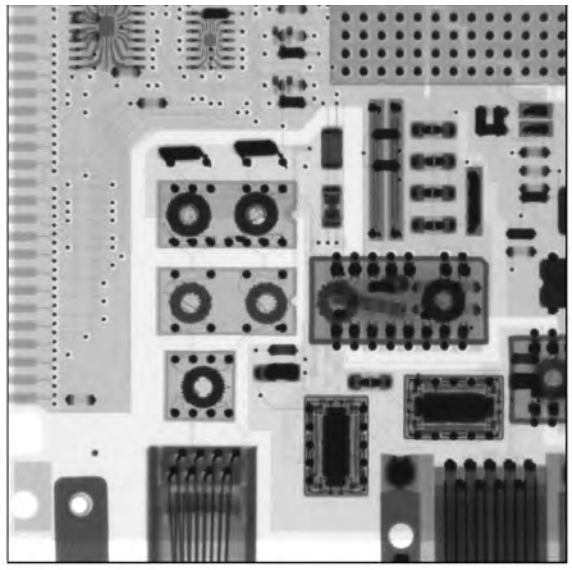

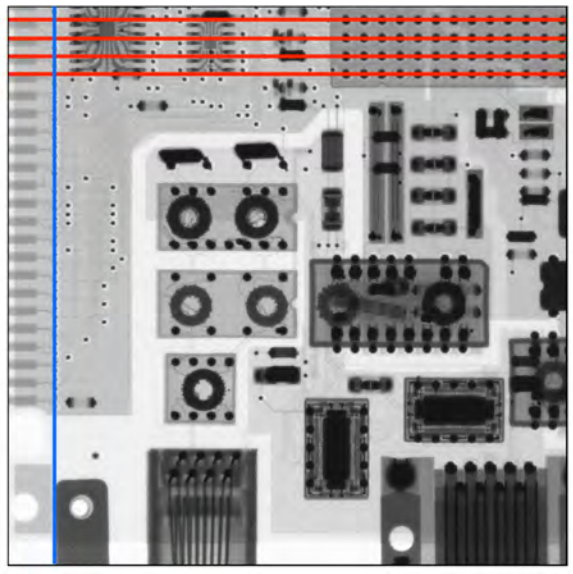

(Top-left) Original image.

(Top right) Detecting linear groups of dark spots, aligned vertically (blue) and horizontally (red).

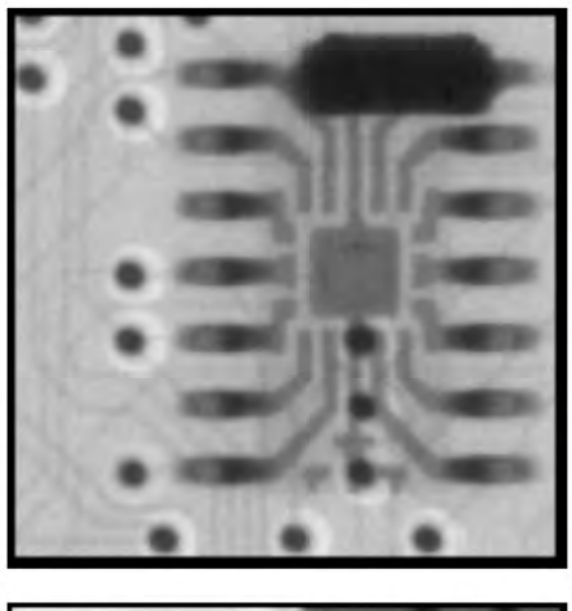

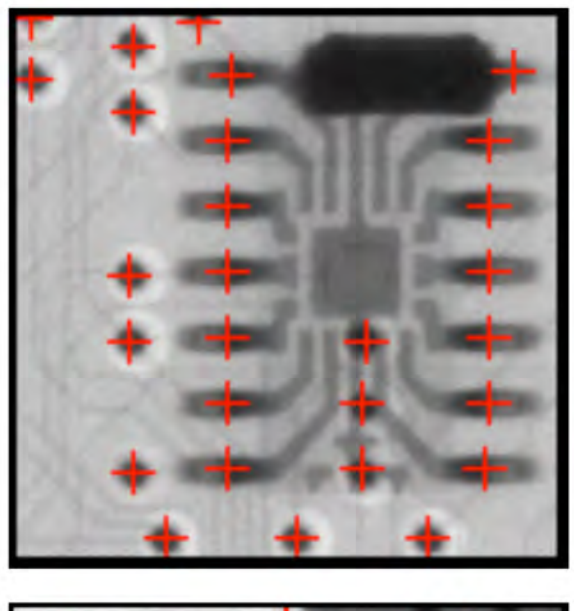

(Centre-left) Sub-image.

(Centre-right) Large dark spots (pads and solder joints).

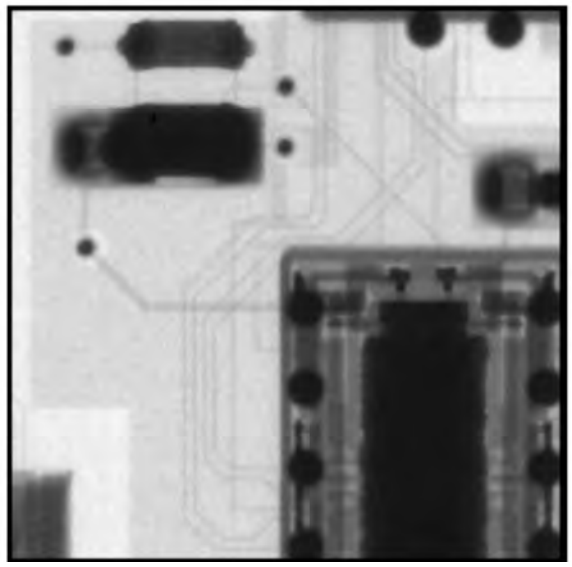

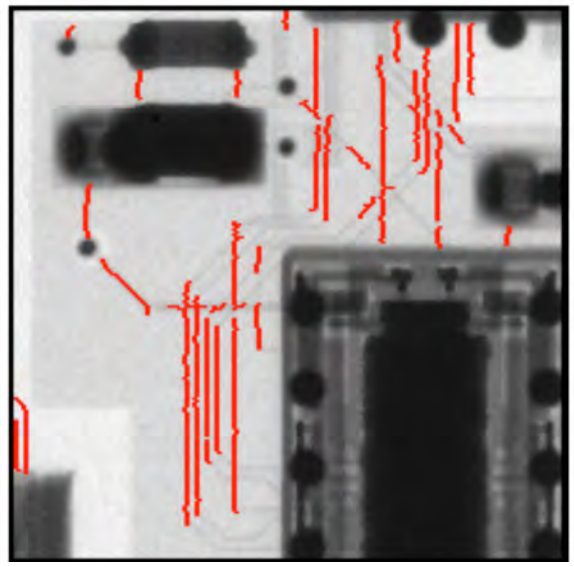

(Bottom left) Another sub-image.

(Bottom-right) Faint vertical grey lines (printed conductors).

X-RAYS (foreign bodies in a filled glass food container)

X-ray inspection of food products is, of course, a hugely important and diverse class of potential applications. This simple example is included here merely to highlight the possibility of Machine Vision being applied in this area. The x-ray behaviour of the container/wrapping, the food item and the likely contaminants all contribute to the formation of the image. Contrast can be very low. To improve matters, dual energy x-rays sensors are sometimes used to generate two images are then co-processed. Modelling the container and filling can, also produce a useful reference guide for interpreting the raw x-ray image.

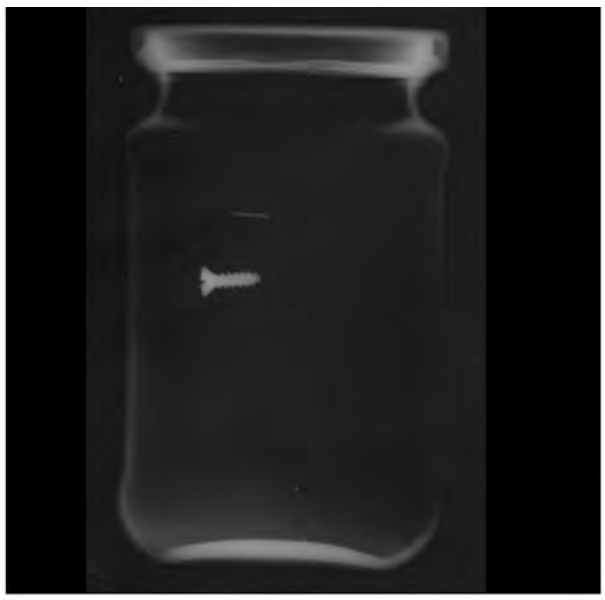

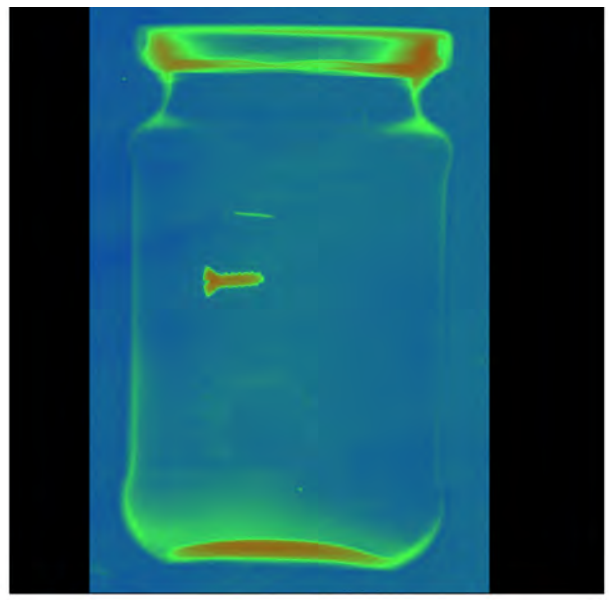

(Top-left) Original image [Glass jar, containing tomato sauce. Metal contaminants were added deliberately for demonstration purposes.]

(Top-right) Pseudo-colouring, with some smoothing first.

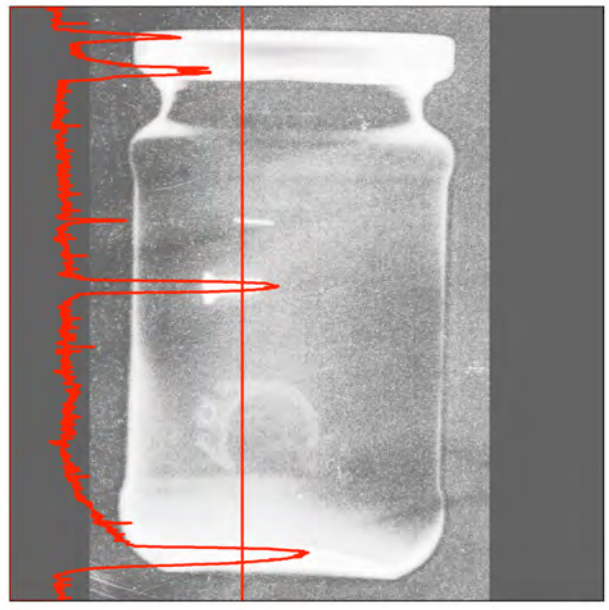

(Centre-left) Histogram equalisation on [TL].

(Centre-right) Intensity profile along a vertical line. Notice the shadow of the thick base of the jar.

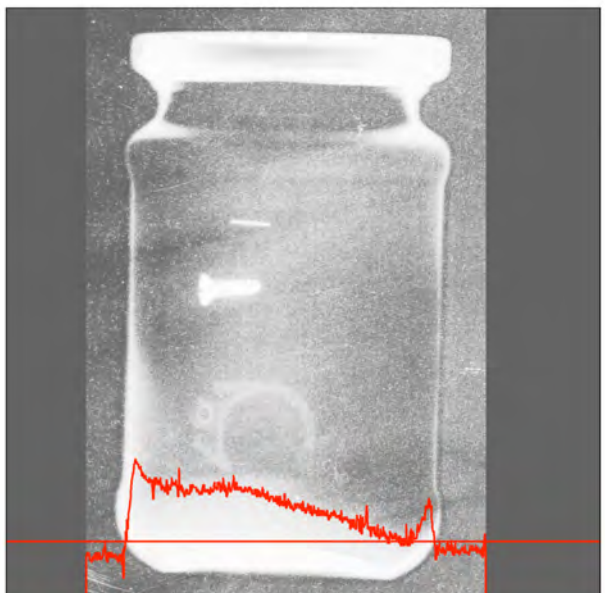

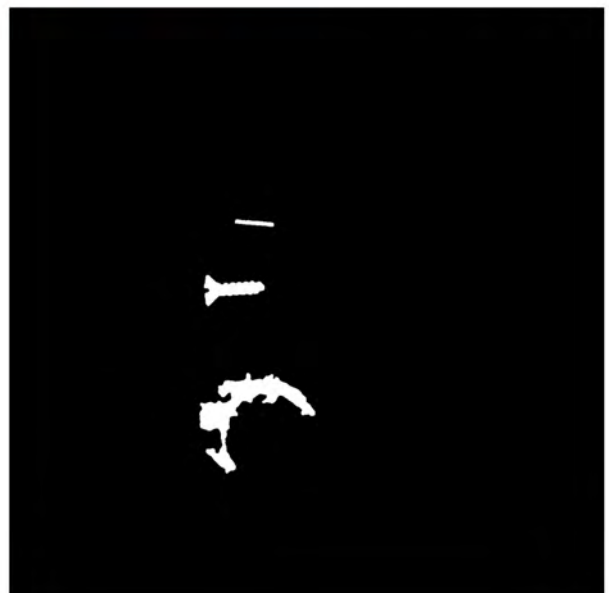

(Bottom-left) Intensity profile along a horizontal line, through the base of the jar

(Bottom-right) Contaminants detected.

X-RAYS (aerosol spray nozzle)

X-ray inspection of manufactured items has huge potential but is often rather more complicated than visual examination for three reasons: (i1) Object shadows are superimposed; (2) X-ray sources do not offer the same flexibility of beam shape, optics, and spectrum; (3) X-rays are more hazardous than visible light. X-rays are emitted in a cone-shaped beam. As a result, they distort the geometry of objects and their apparent size varies with position along along the beam axis. X-ray inspection is made easier if we know exactly what to expect; modelling the image formation is very useful.

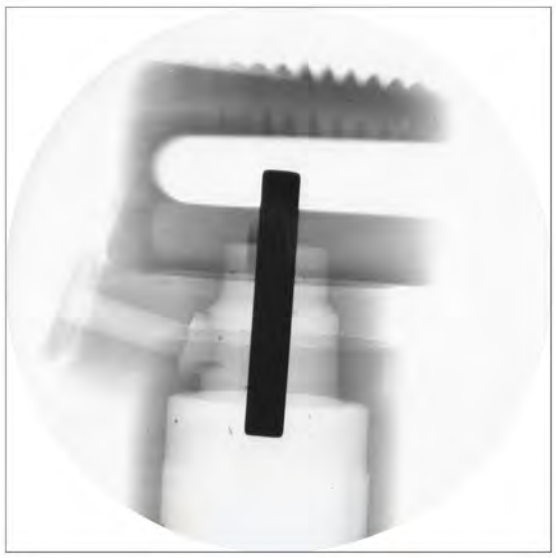

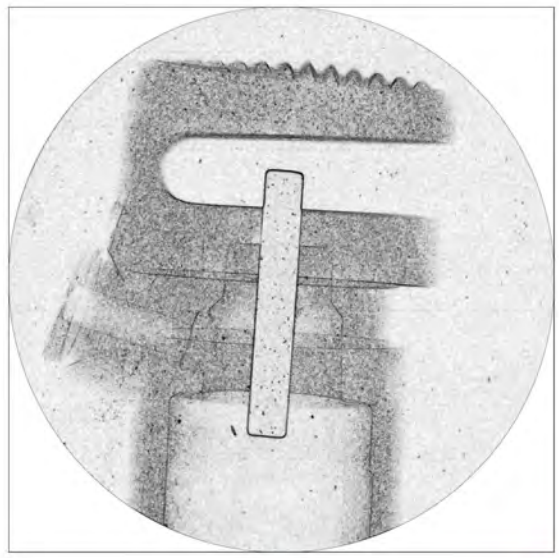

(Top-left) Original image.

(Top-right) High-pass filtering, emphasises noise.

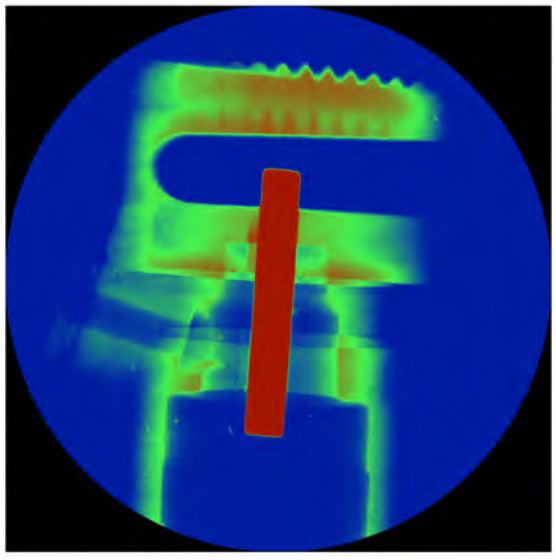

(Centre-left) Filtering to reduce noise, then false colouing.

(Centre-right) Intensity profile, along the vertical red line.

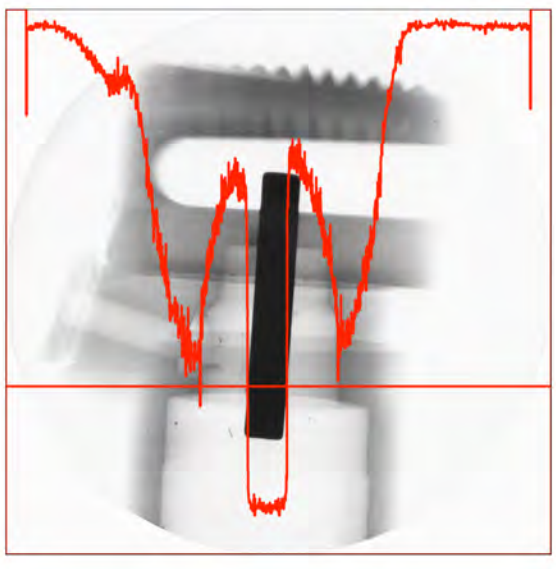

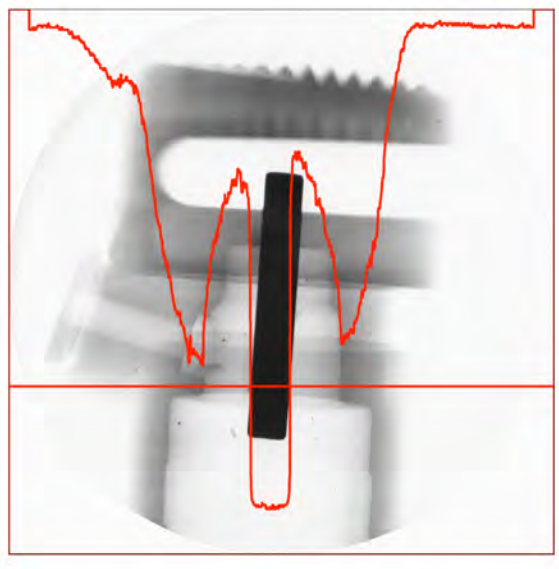

(Bottom-left) Intensity profile, along the horizontal red line..

(Bottom-right) Median filter, reduces noise level.

X-RAYS (9V battery)

(Top-left) Original image.

(Top-right) Intensity profile along a horizontal line.

(Centre-left) Intensity profile along a vertical line.

(Centre-right) Edge detector, emphasises the strong linear features.

(Bottom-left) Horizontal and vertical linear features.

(Bottom-right) [BL] superimposed on the original.

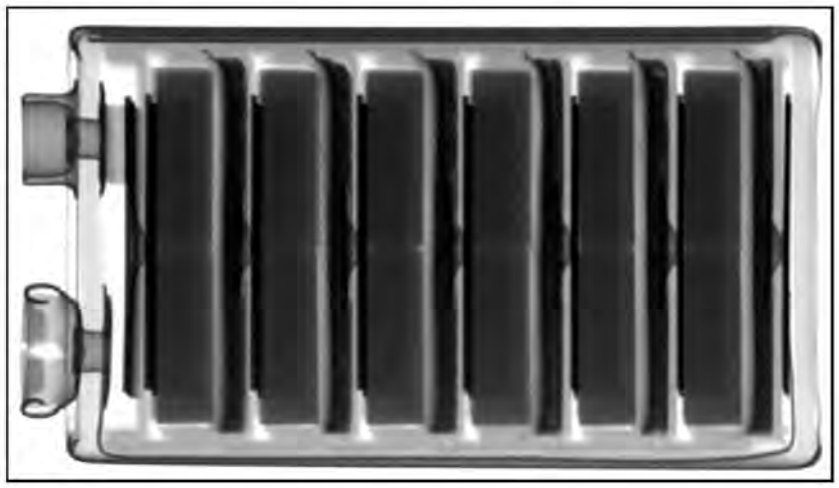

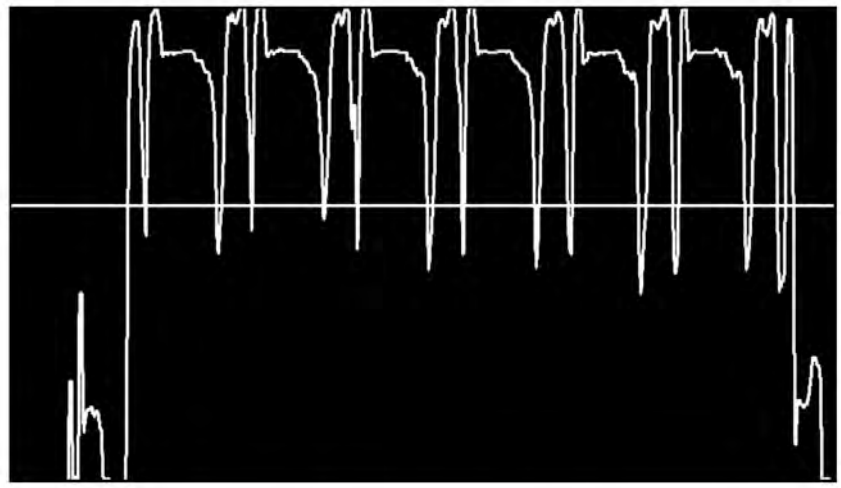

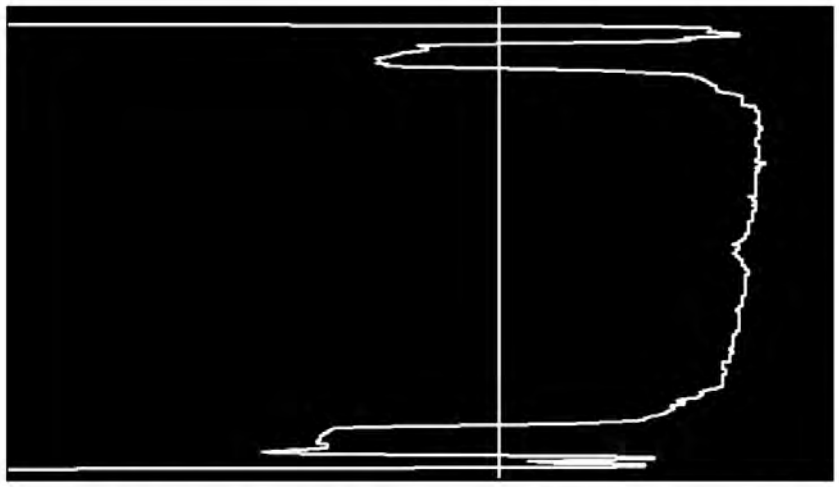

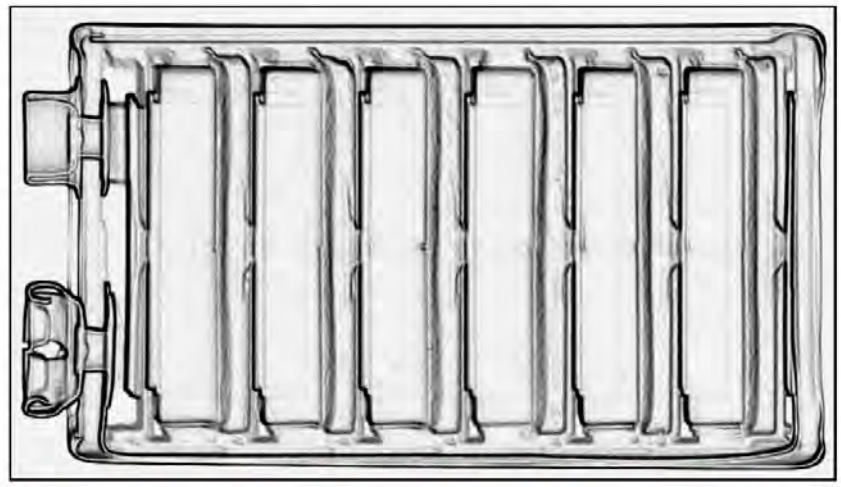

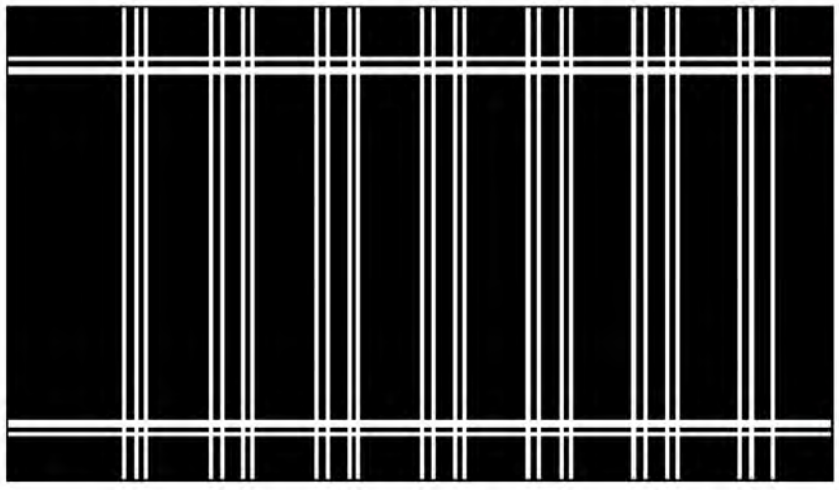

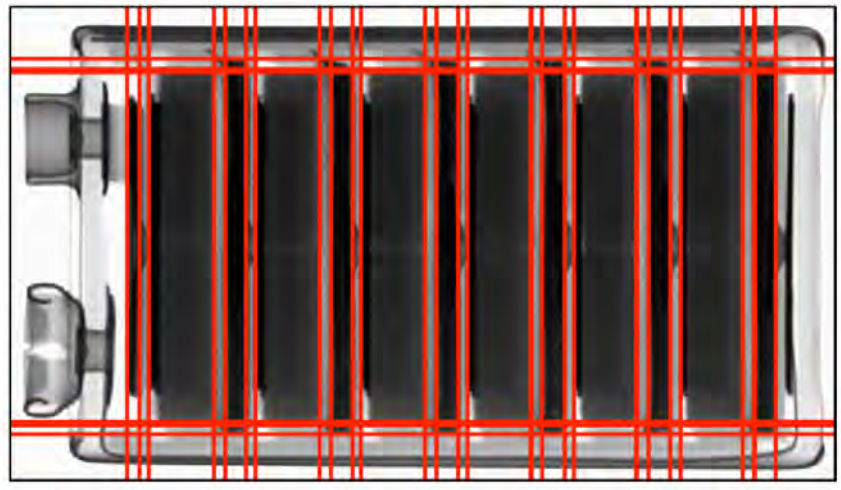

X-RAYS (aircraft turbine blade)

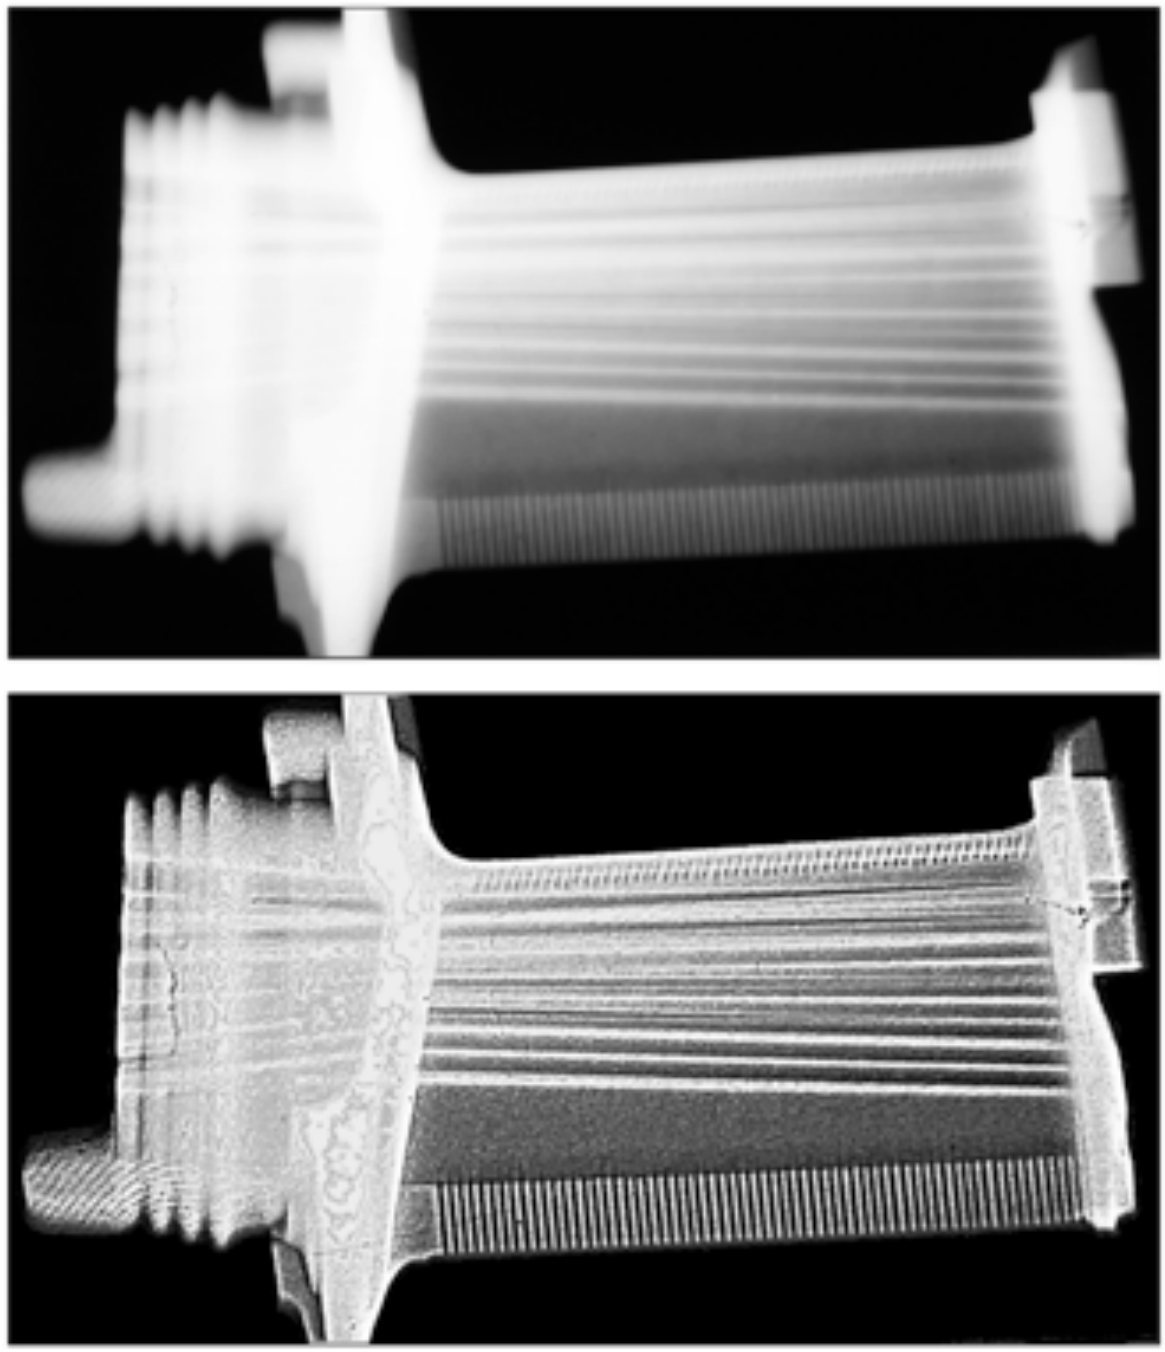

(Top) Original image.

(Centre) Contrast enhanced

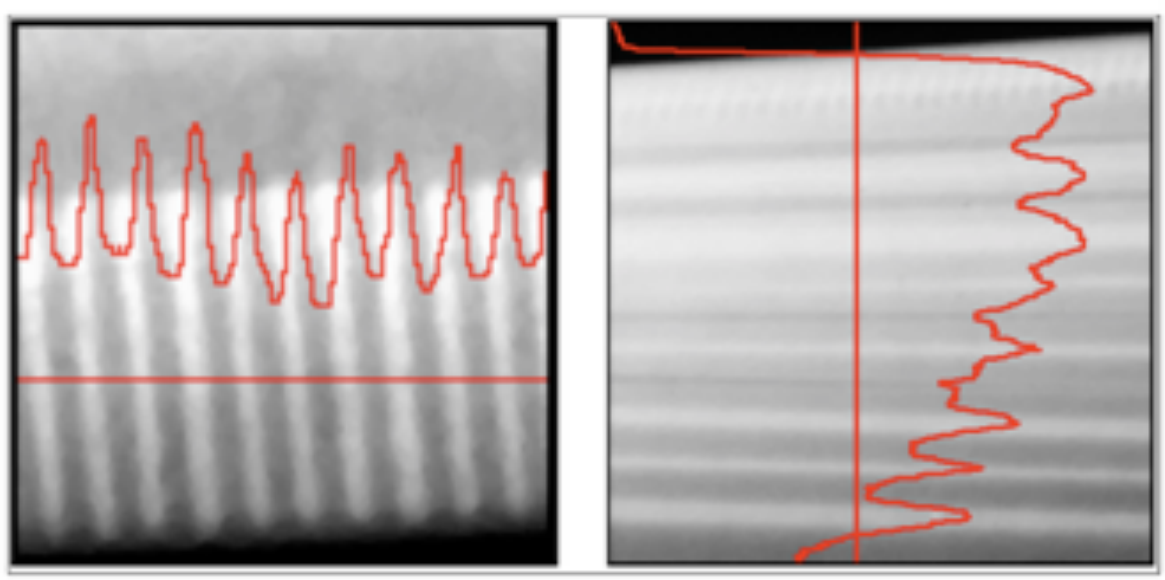

(Bottom-left) Expanded view of the stripes on the bottom edge (fine air ducts) Intensity profile.

(Bottom-right) Expanded view of the stripes along the body of the blade (air ducts) Intensity profile.

CRACK DETECTION USING X-RAYS

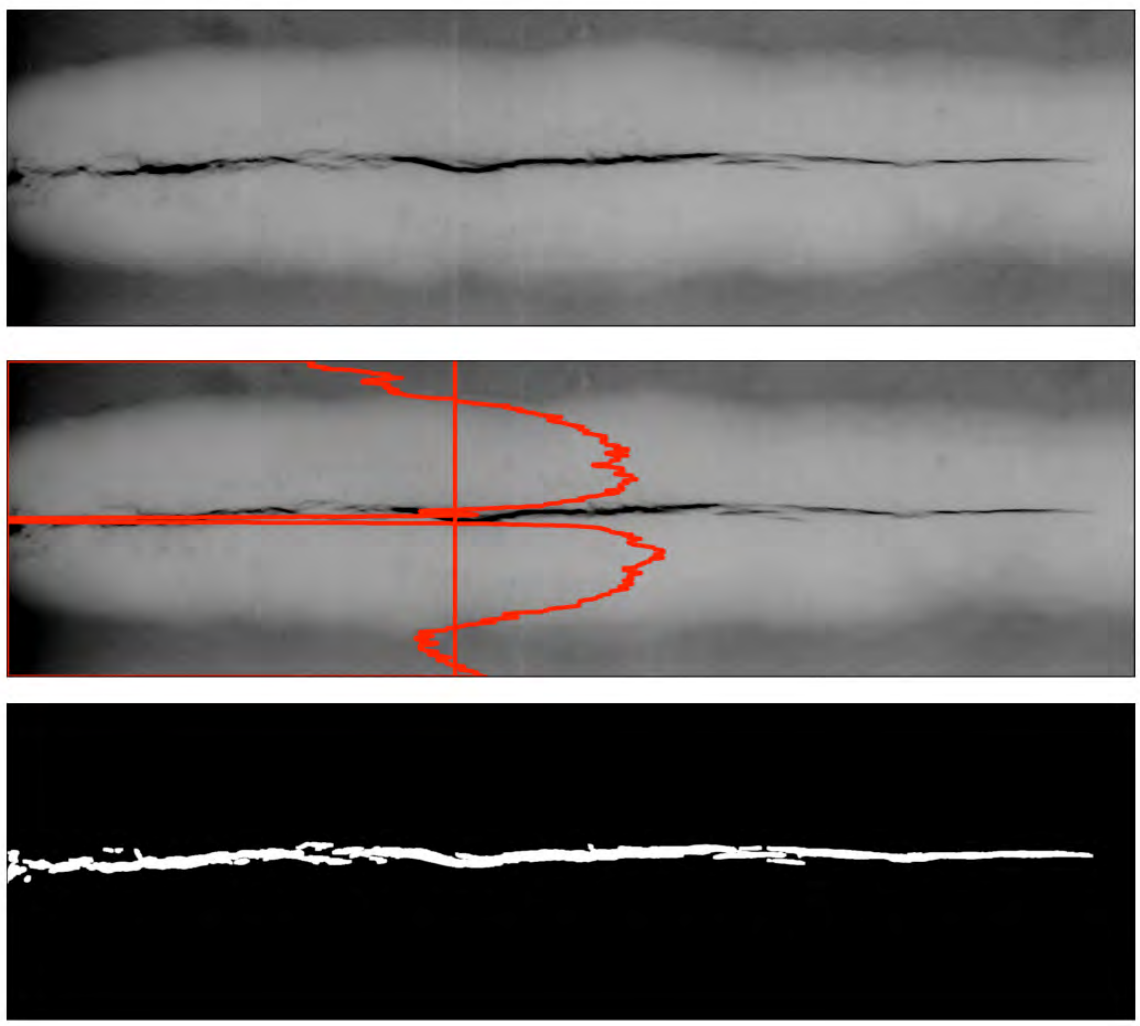

(Top) Original image of a welded joint.

(Row 2) Intensity profile.

(Row 3) “Crack detector” filter.

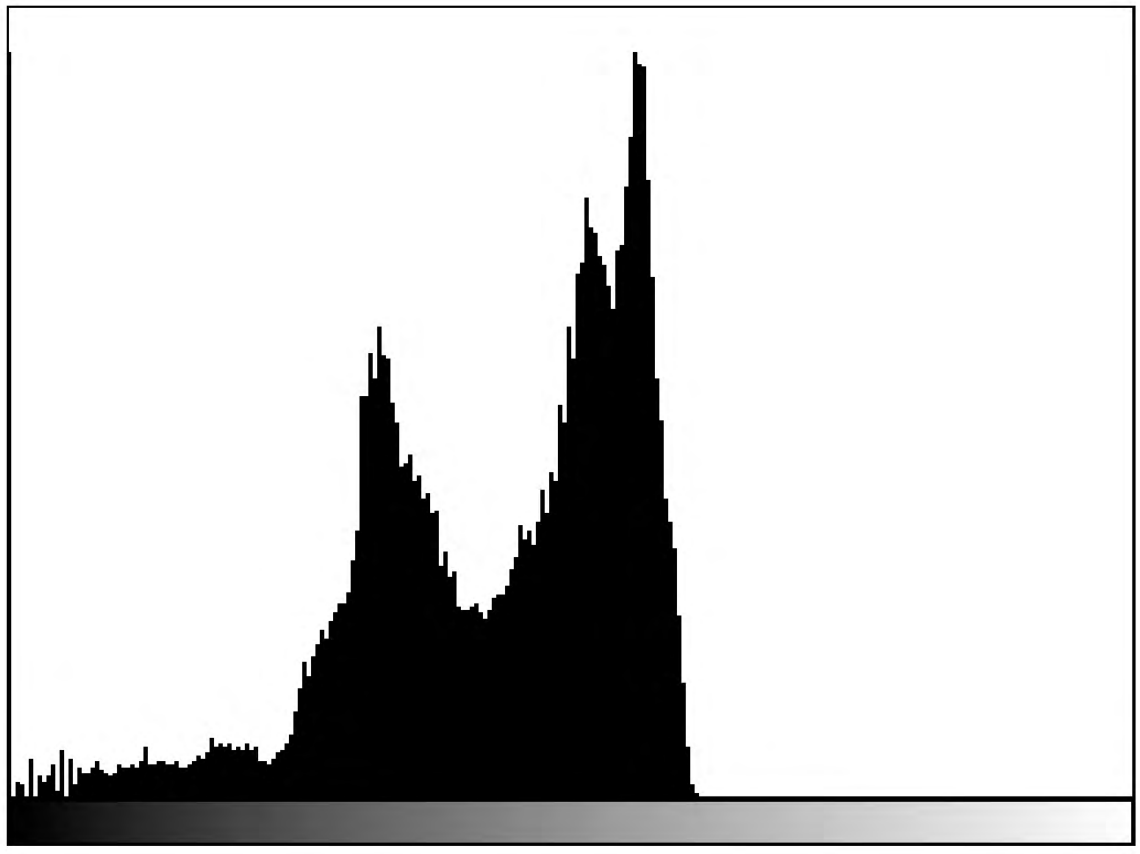

(Bottom) Intensity histogram of [T]. In this instance, thresholding also identifies the crack reasonably well.

Section 5

THERMAL IMAGING

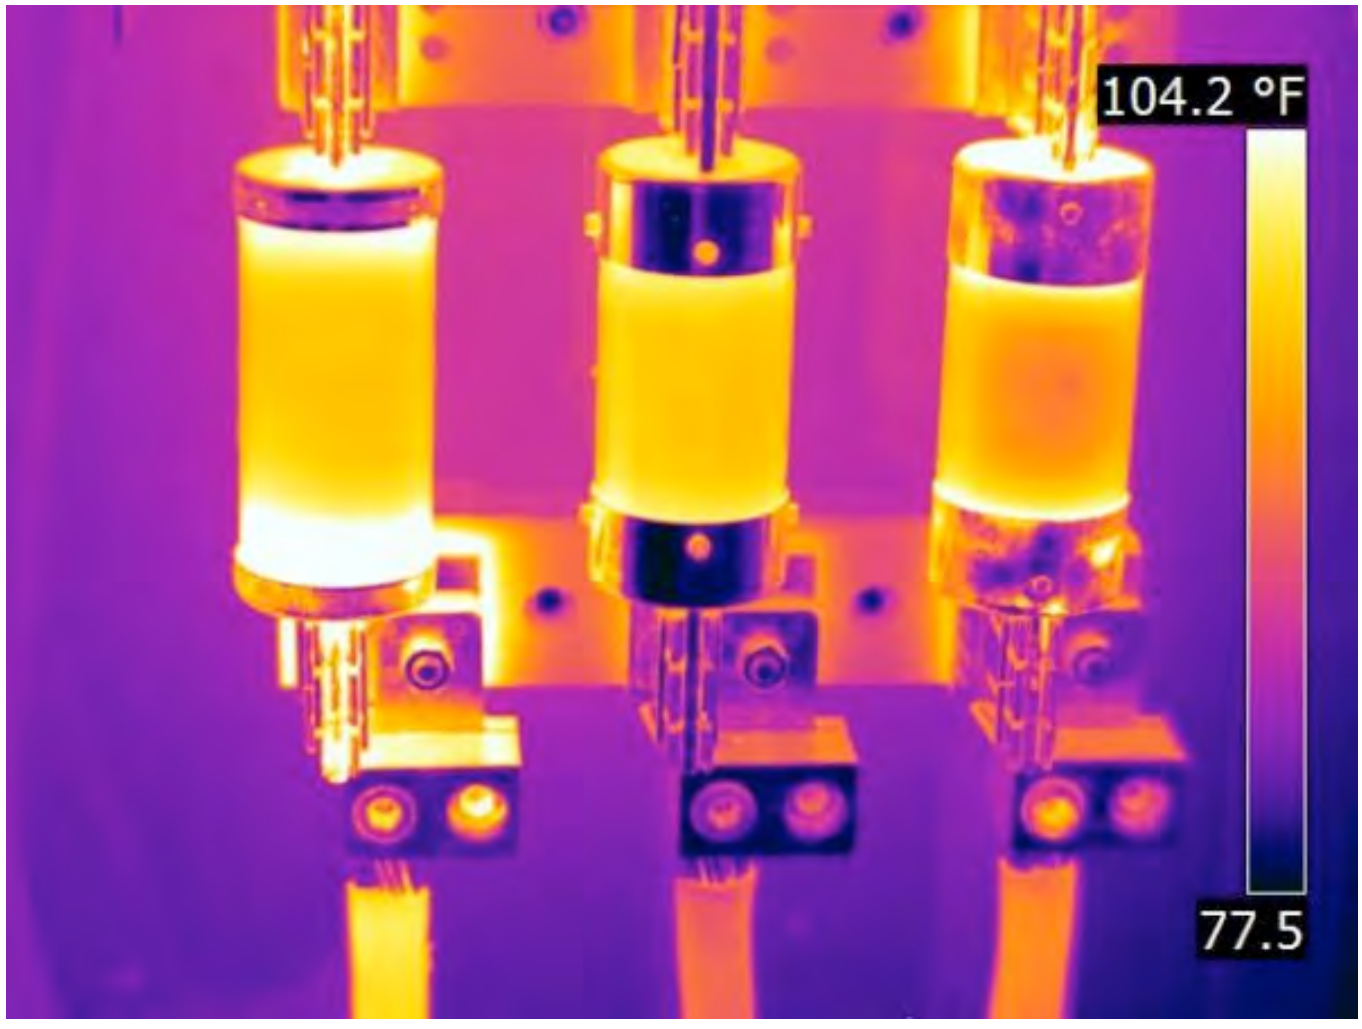

THERMAL IMAGING (sample images and pseudo-colour mapping)

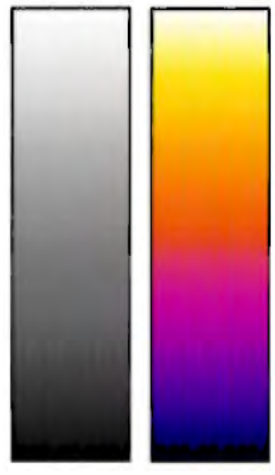

Thermal images are very often shown using a pseudo-colour mapping shown shown opposite. However, the grey-scale version is normally preferred when processing thermal images.

(Top) Electrical wiring. Hot spots are bright.

(Centre-left) Electrical power line.

(Centre-centre) Industrial pipe-work.

(Centre-right) Aero-engine.

(Bottom-left) Damp floor. The wet areas are cooler (darker) due to evaporation

(Bottom-centre) Heat loss at the corner of a room. Notice the "shadow" of the battens.

(Bottom-right) Heat loss in a building.

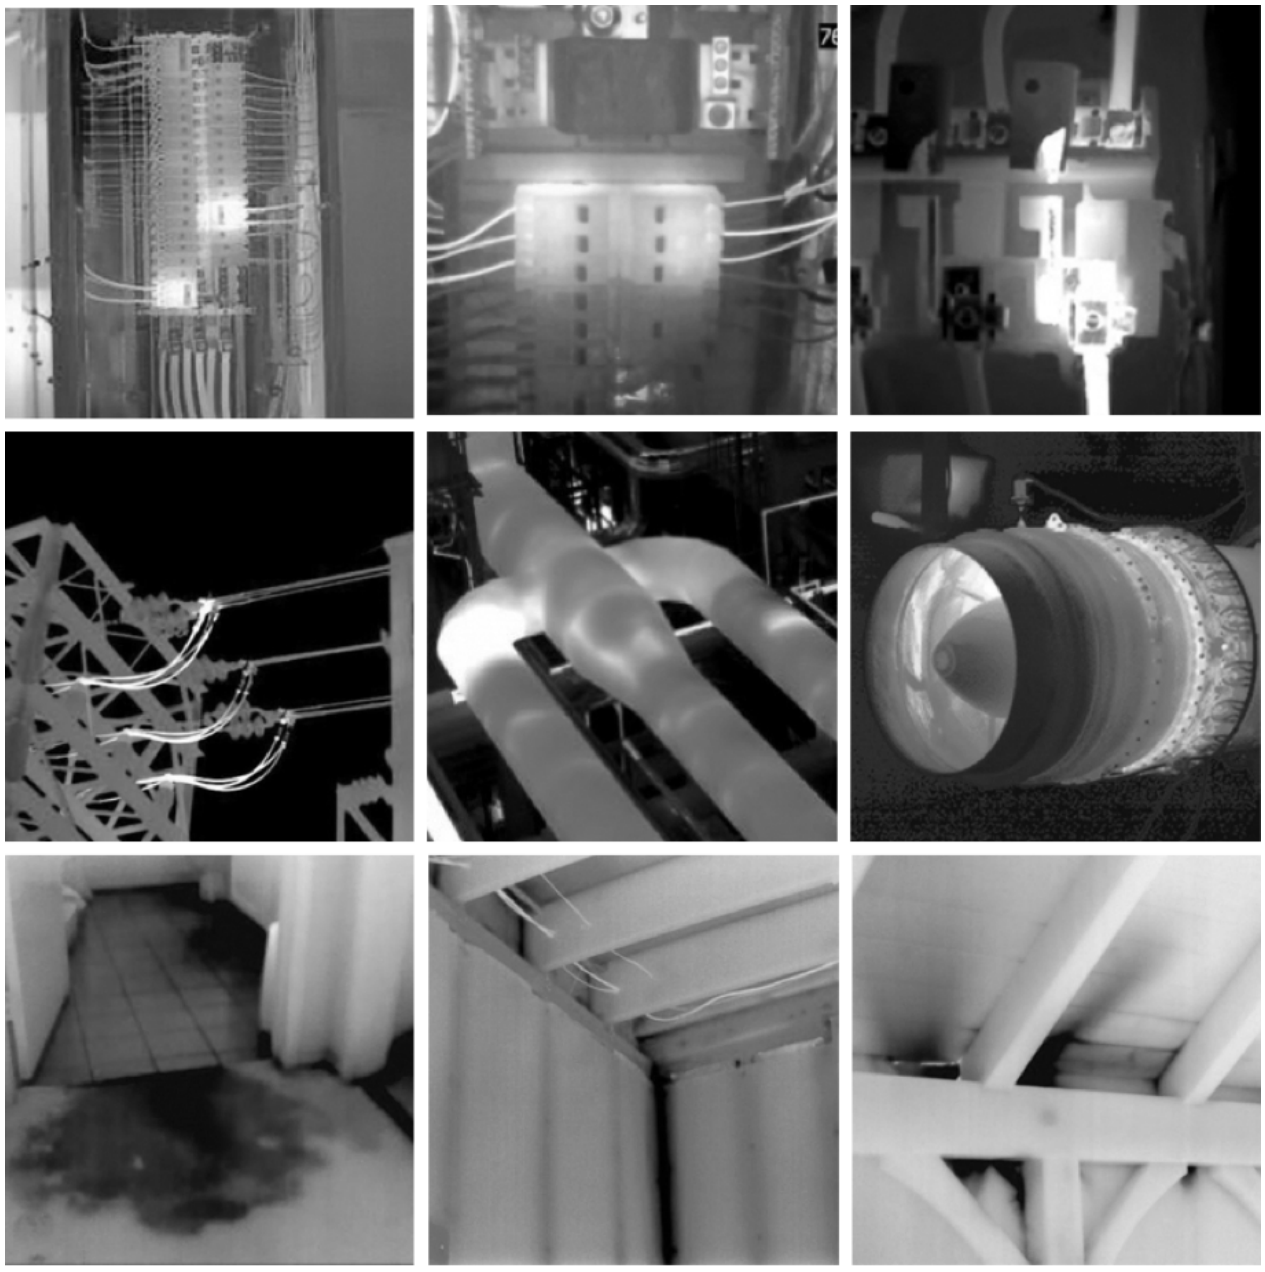

THERMAL (electric fuses)

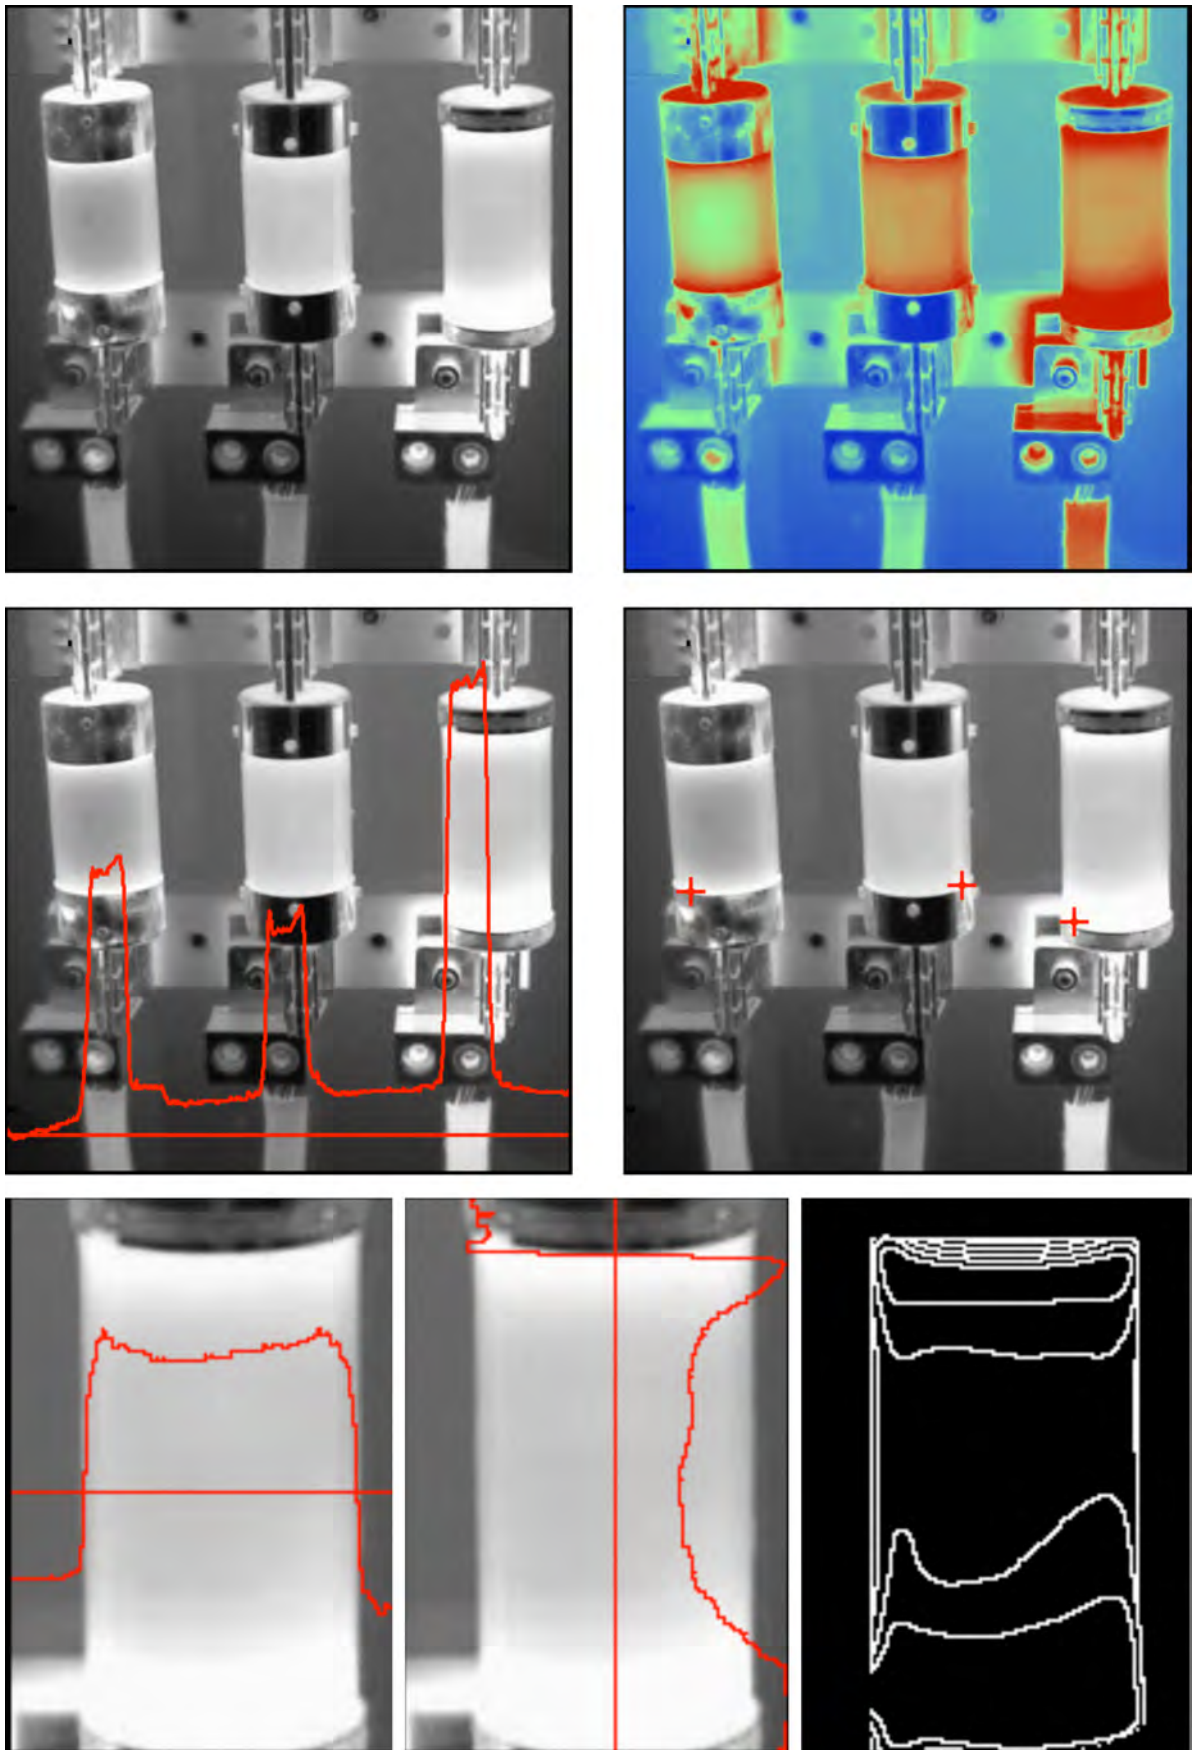

(Top-left) Thermal image of electrical connectors and fuses. (bright rectangles). Original image. Intensity is approximately proportional to temperature over the range 25-40°C.

(Top-right) Pseudo-colour rendering of [TL]. This uses a different mapping from that shown on the previous page.)

(Centre-left) Temperature profile across the cables. (Horizontal red line)

(Centre-right) Crosses indicate points of maximum temperature for each fuse.

(Bottom-left & Bottom-centre) Row & column temperature profiles for the right-most fuse.

(Bottom-right) Isotherms (Isophotes) for the same fuse.

HEATING BY AIR FLOW

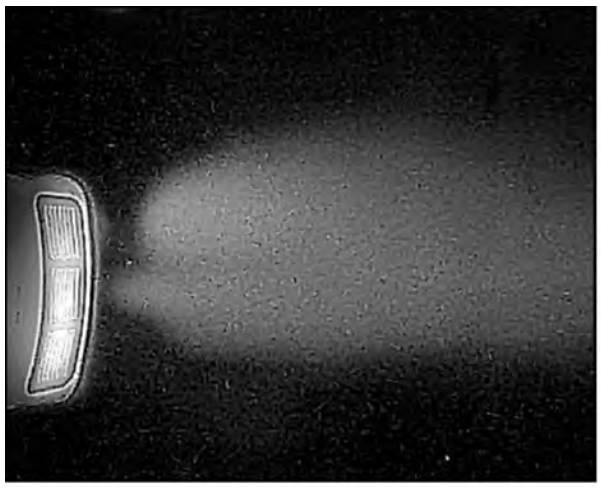

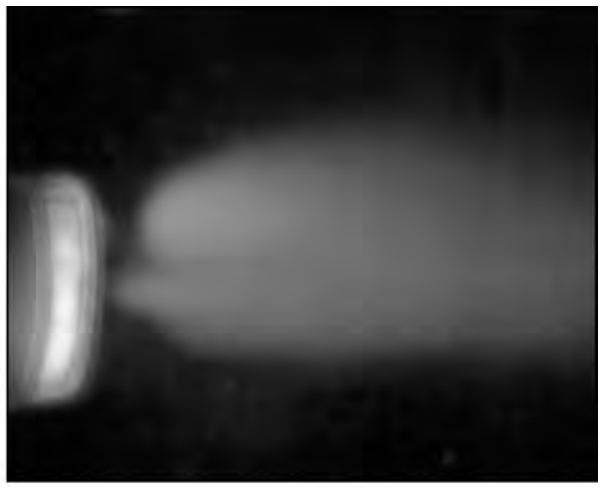

(Top-left) Original image. Domestic electric fan-heater blowing hot air onto a carpeted floor.

(Top-right) Median filter reduces noise.

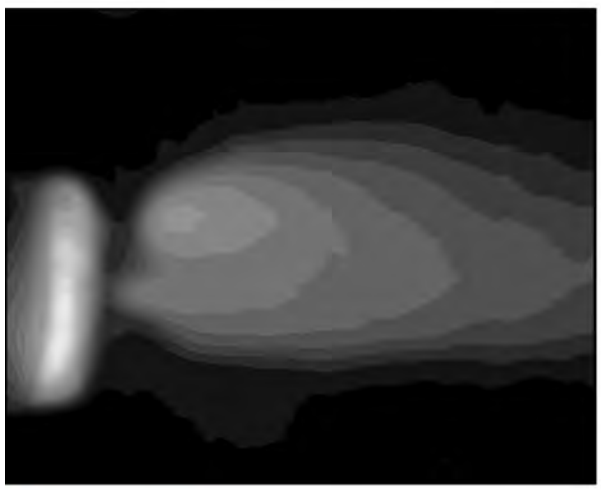

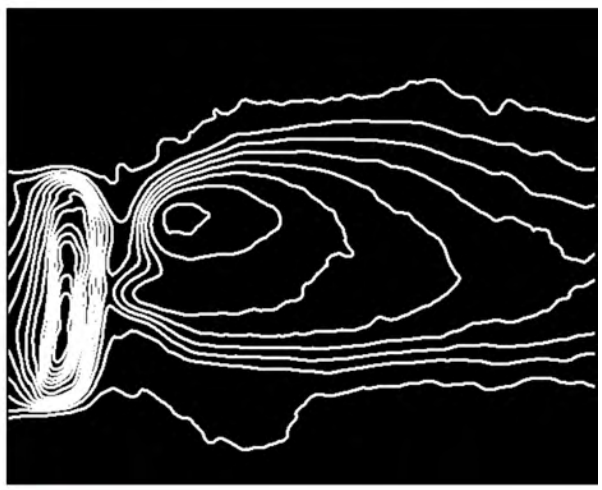

(Centre-left) Posterising, 16 levels.

(Centre-right) Edges in [CR]. Isophotes, equivalent to isotherms.

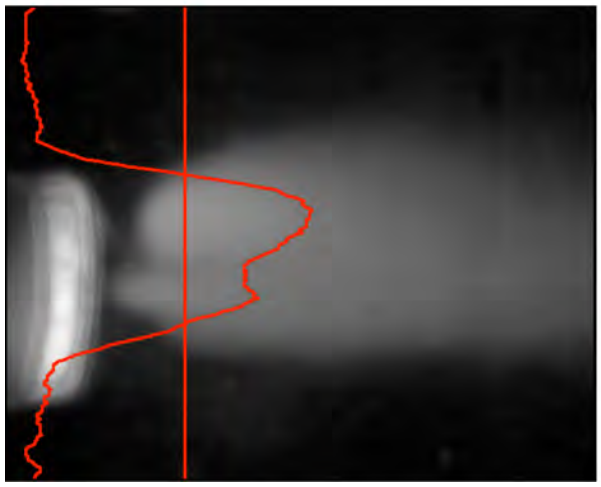

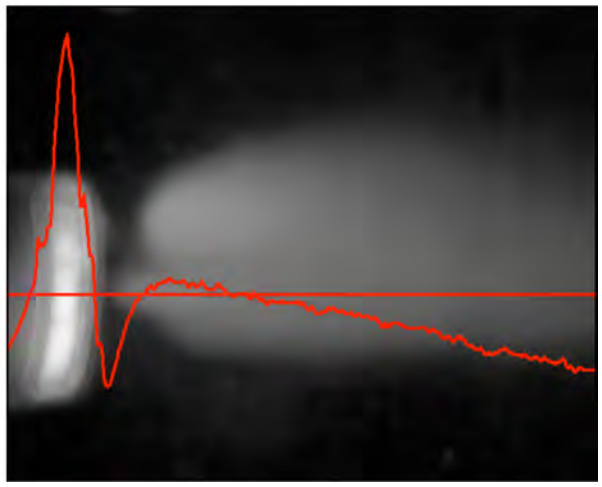

(Bottom-left) Intensity profile, column.

(Bottom-right) Intensity profile, row.

HOT BOTTLES (IR self-luminance)

The photograph shows red-hot bottles that are just a few seconds old. Automated inspection is invaluable because a person cannot work close to the hot moulding machines for more than a few minutes.

The mouth and upper-neck region of the bottle is cooler (darker) than the body. The base of the bottle is hottest, where the glass is thickest and therefore retains heat longest.

Solid-state cameras are highly sensitive to IR. It is therefore possible to use either visible-light, IR , or a combination of the two for inspection. Examining the thermal emission can yield information about body shape, estimate the body-wall thickness and detect the presence of certain defects. An MV inspection system can monitor changes in the IR-emission profile and provide an early alert of malfunctioning of the manufacturing system.

THERMAL IMAGING (U-bend water trap)

(Top-left) Original image Notice that it contains hot water

(Top-right) Tops of the water columns.

(Centre-left) Maximum intensity points (hot spots)

(Centre-right) Intensity contours/isotherms.

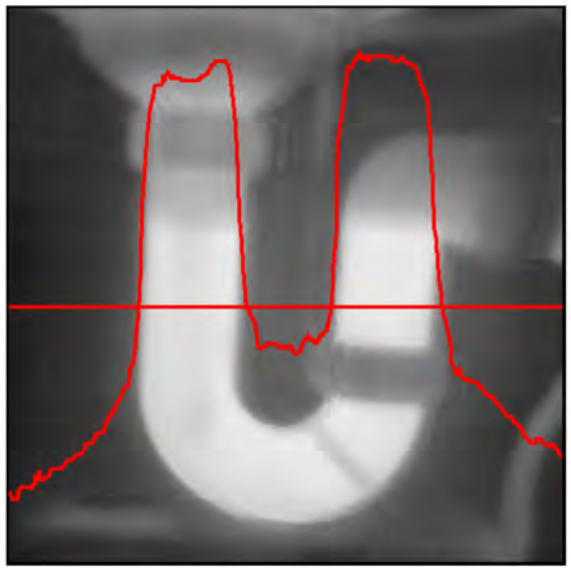

(Bottom-left) Intensity profile, row.

(Bottom-right) Intensity profile, column.

ENERGY MONITORING (electrical power strip)

This application shows visible and IR images being combined.

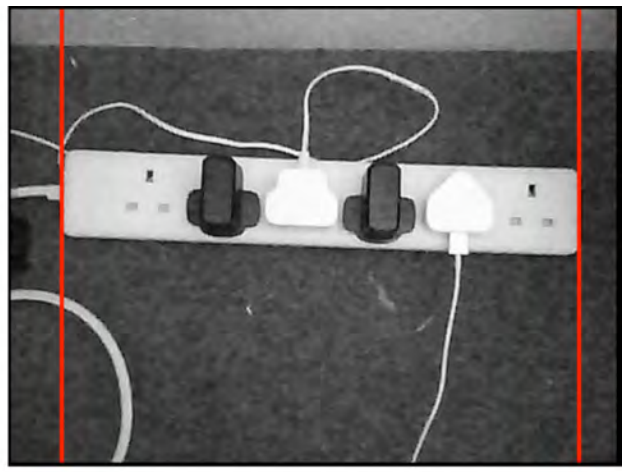

(Top) Original image, visible-light (VIS) camera

(Top-right) Power strip end limits, calculated from [TL].

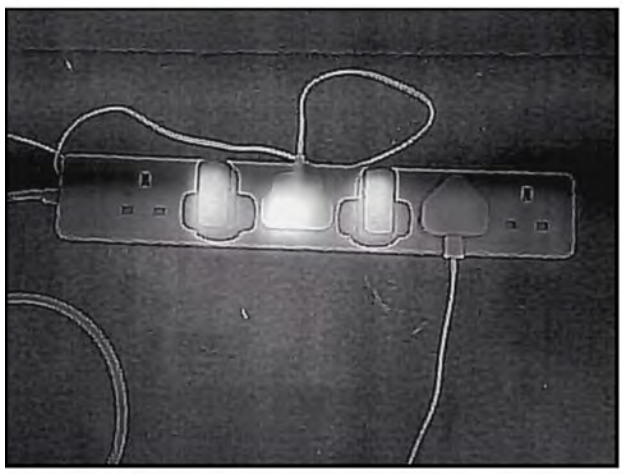

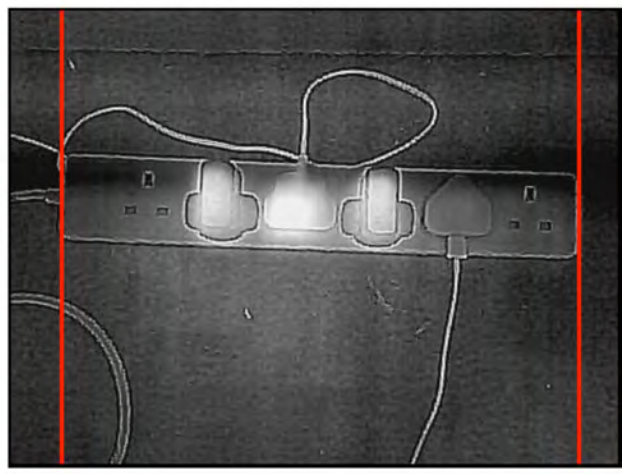

(Centre-left) IR image. (FLIR One camera

(Centre-right) End-edges superimposed on [CL].

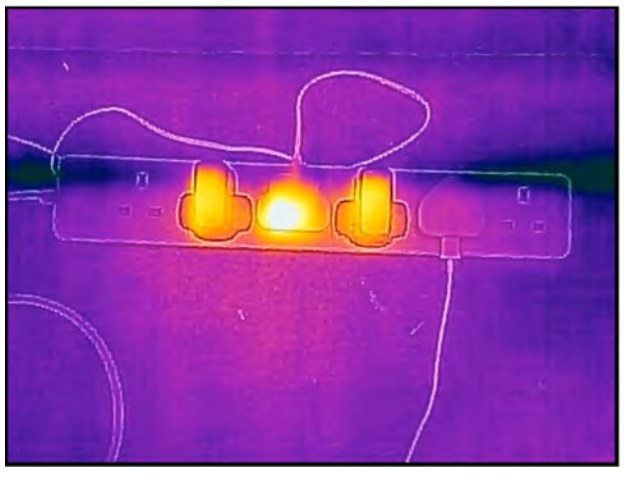

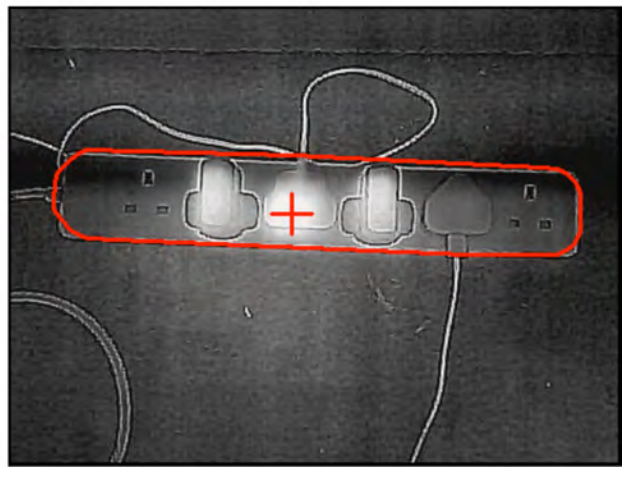

(Bottom-left) [CL] in pseudo-colour.

(Bottom-right) Edge contour, calculated from [TL], superimposed on the IR image. Cross indicates the brightest point (hottest).

THERMAL IMAGING, RAYNAUD'S SYNDROME

Raynaud's Syndrome is characterised by cold hands, feet, ears, nose, lips or nipples. It is therefore ideally suited to examination by thermal imaging. The image shown here was generated by a low-cost thermal imager (FLIR One. Cost about £200). The hand used in this study is my own as are the ideas for processing the camera image.

(Top-left) Original image.

(Top-right) Median filter & blurring filter.

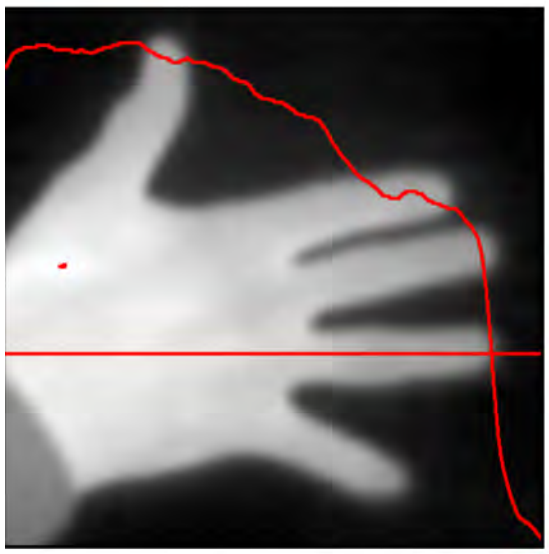

(Centre-left) Intensity profile, row.

(Centre-right) Intensity profile, column.

(Bottom-left) Intensity contours/ isotherms.

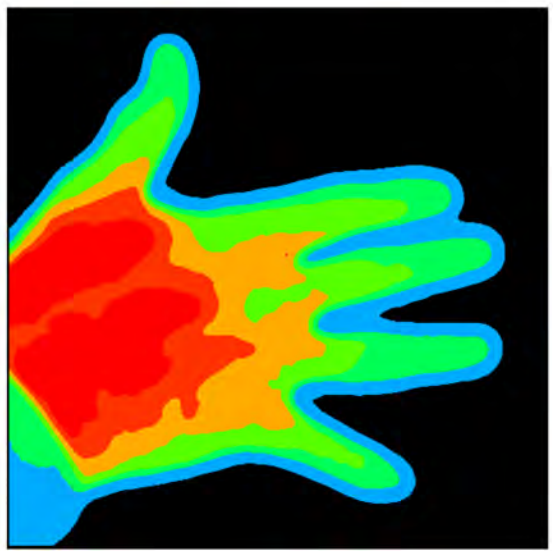

(Bottom-right) Posterising & pseudo colour applied to [TR].

THERMAL IMAGING (footprints)

Visualising the pressure areas on the soles of feet when standing was suggested to me as a potential application of image processing around 1978. Deformities can be observed by looking at the contact areas when standing bare-foot on a sheet of glass. Thermal imaging provides an alternative method of sensing. (Camera: FLIR One)

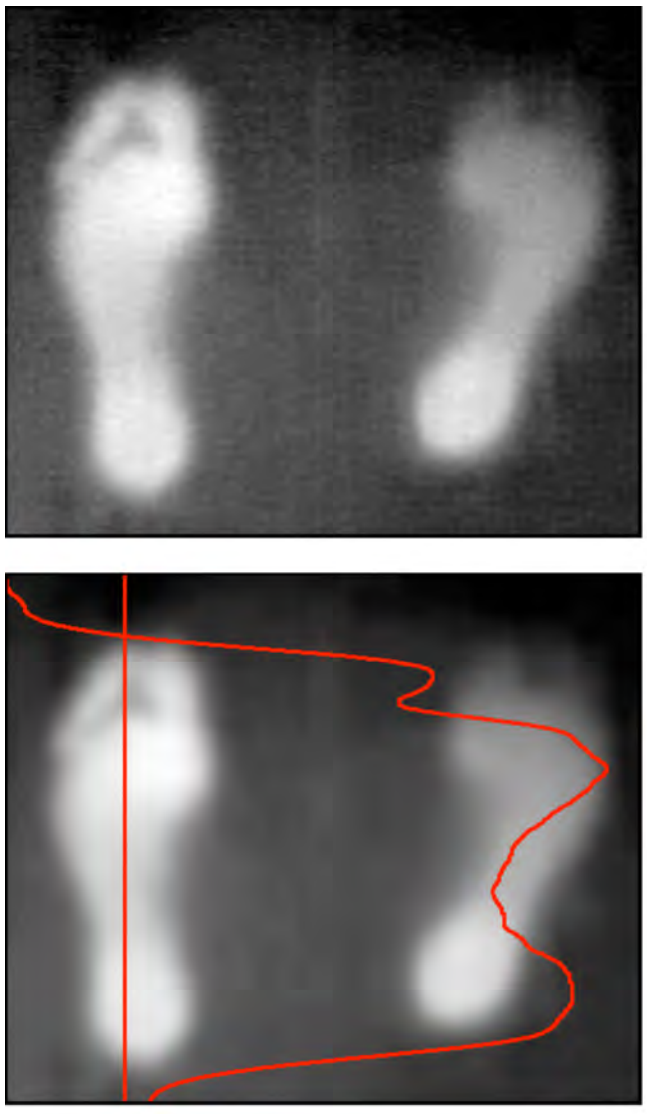

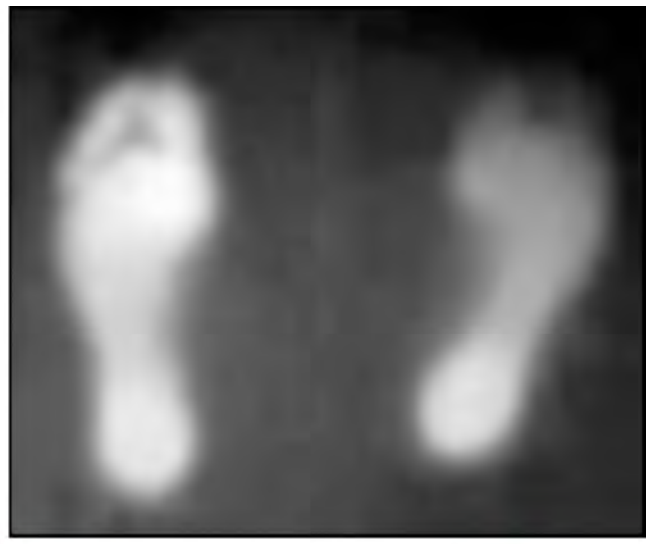

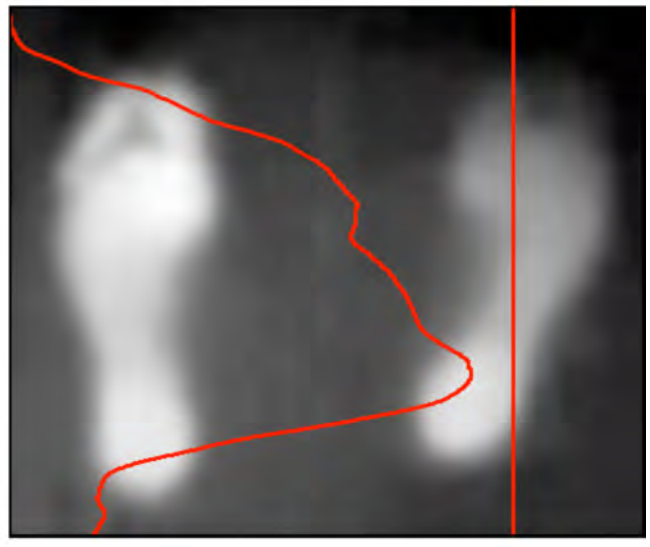

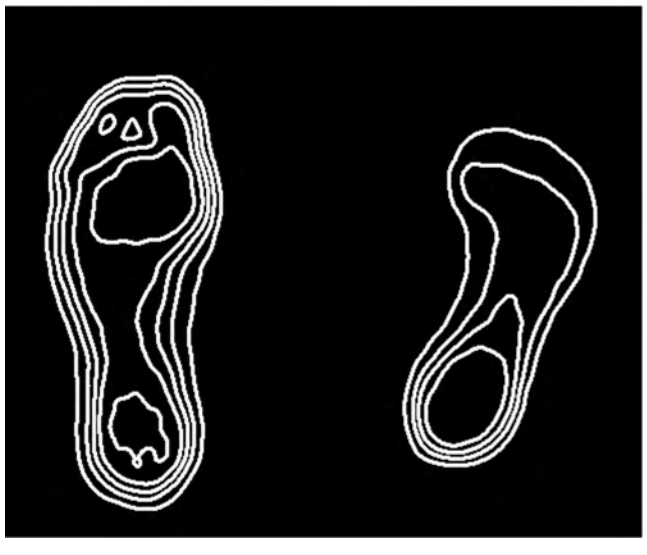

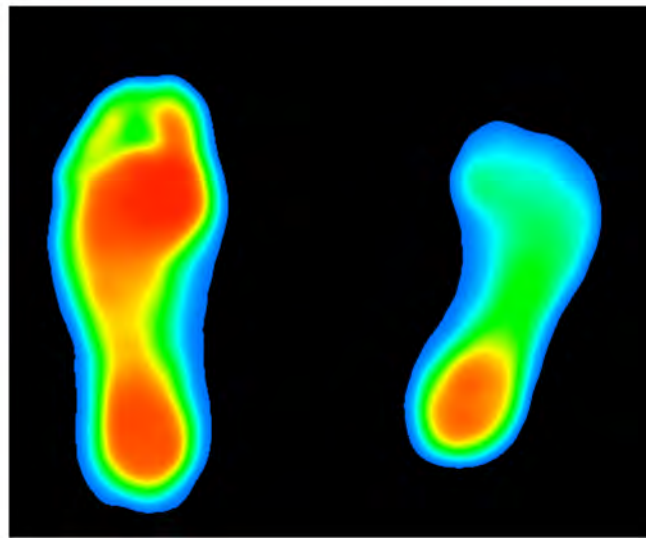

(Top) Original image. [Heat “shadow” left on a carpeted floor after standing barefoot for a few seconds]

(Top-right) Median & smoothing filters.

(Centre-left) Intensity profile, column.

(Centre-tight) Intensity profile, column.

(Bottom-left) Isophotes/isotherms.

(Bottom-right) Pseudo-colour rendering of [TR].

Section 6

MANUFACTURING INDUSTRY

CATALYTIC CONVERTER SUBSTRATE

Ceramic substrate for an automobile catalytic converter.

(Left) Original image. This poor image was obtained using a standard macro lens and front lighting. The camera's angle of view varies across the image. In the dark area, near the top, the camera is able to see deep into individual exhaust-gas channels. Elsewhere, the angle of view is slightly oblique, so the brighter top parts of the side walls are evident.

(Right) There are 3449 black spots; one channel was missed. (Near the edge, 11 o’clock position.)

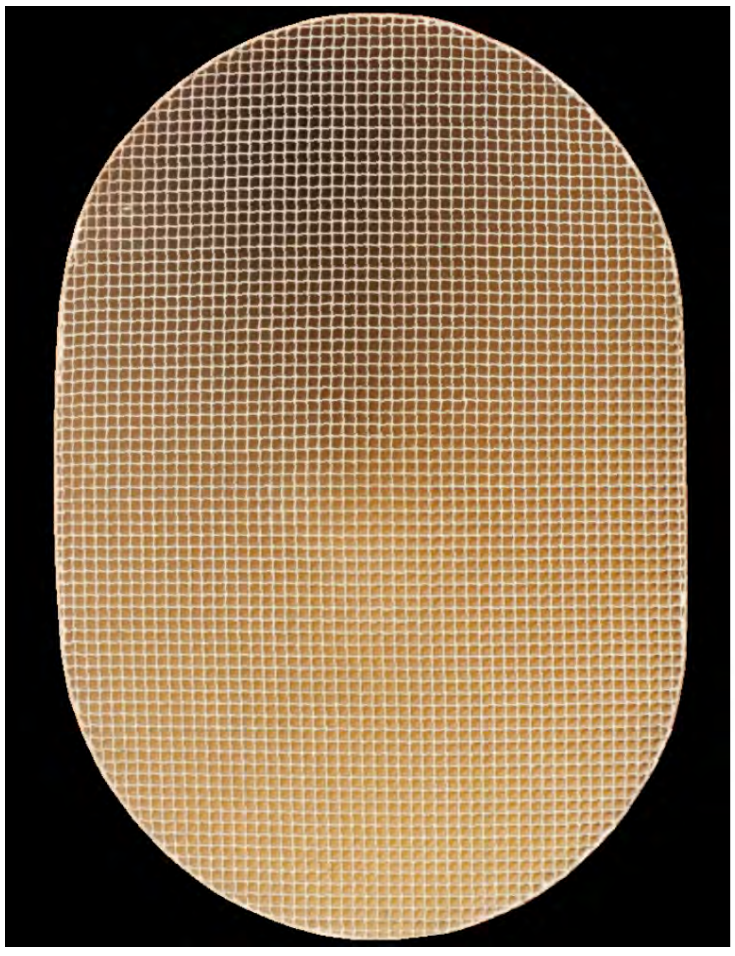

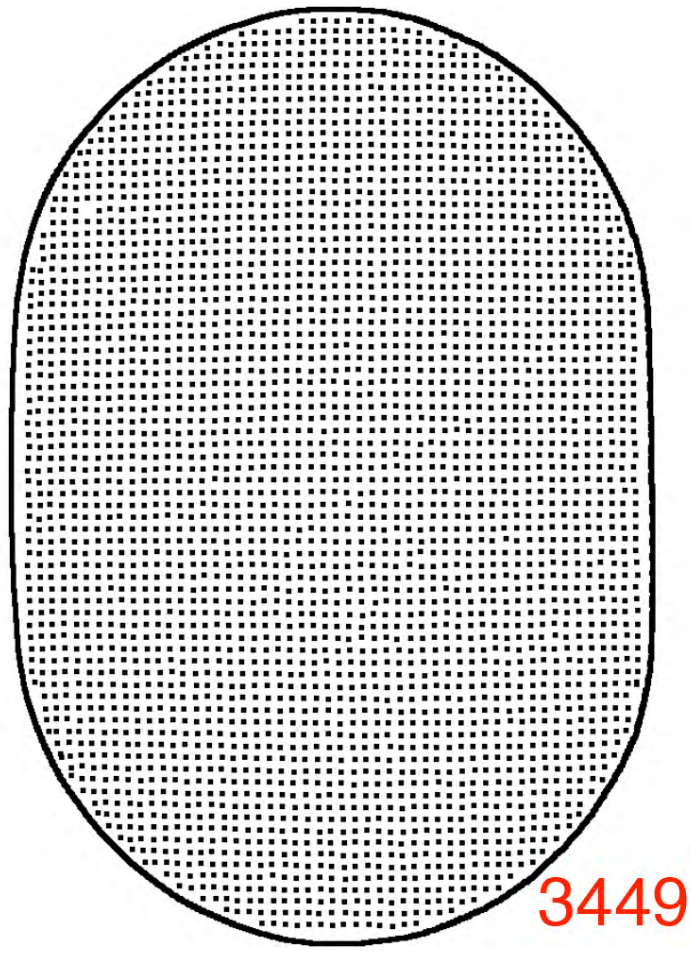

PRINTED CIRCUIT BOARD (locating component lead holes)

(Top-left) Original image.

(Top-right) R-image. (Red channel)

(Bottom-left) [TR] after thresholding and noise reduction.

(Bottom-right) Centroids of white areas in [BL]..

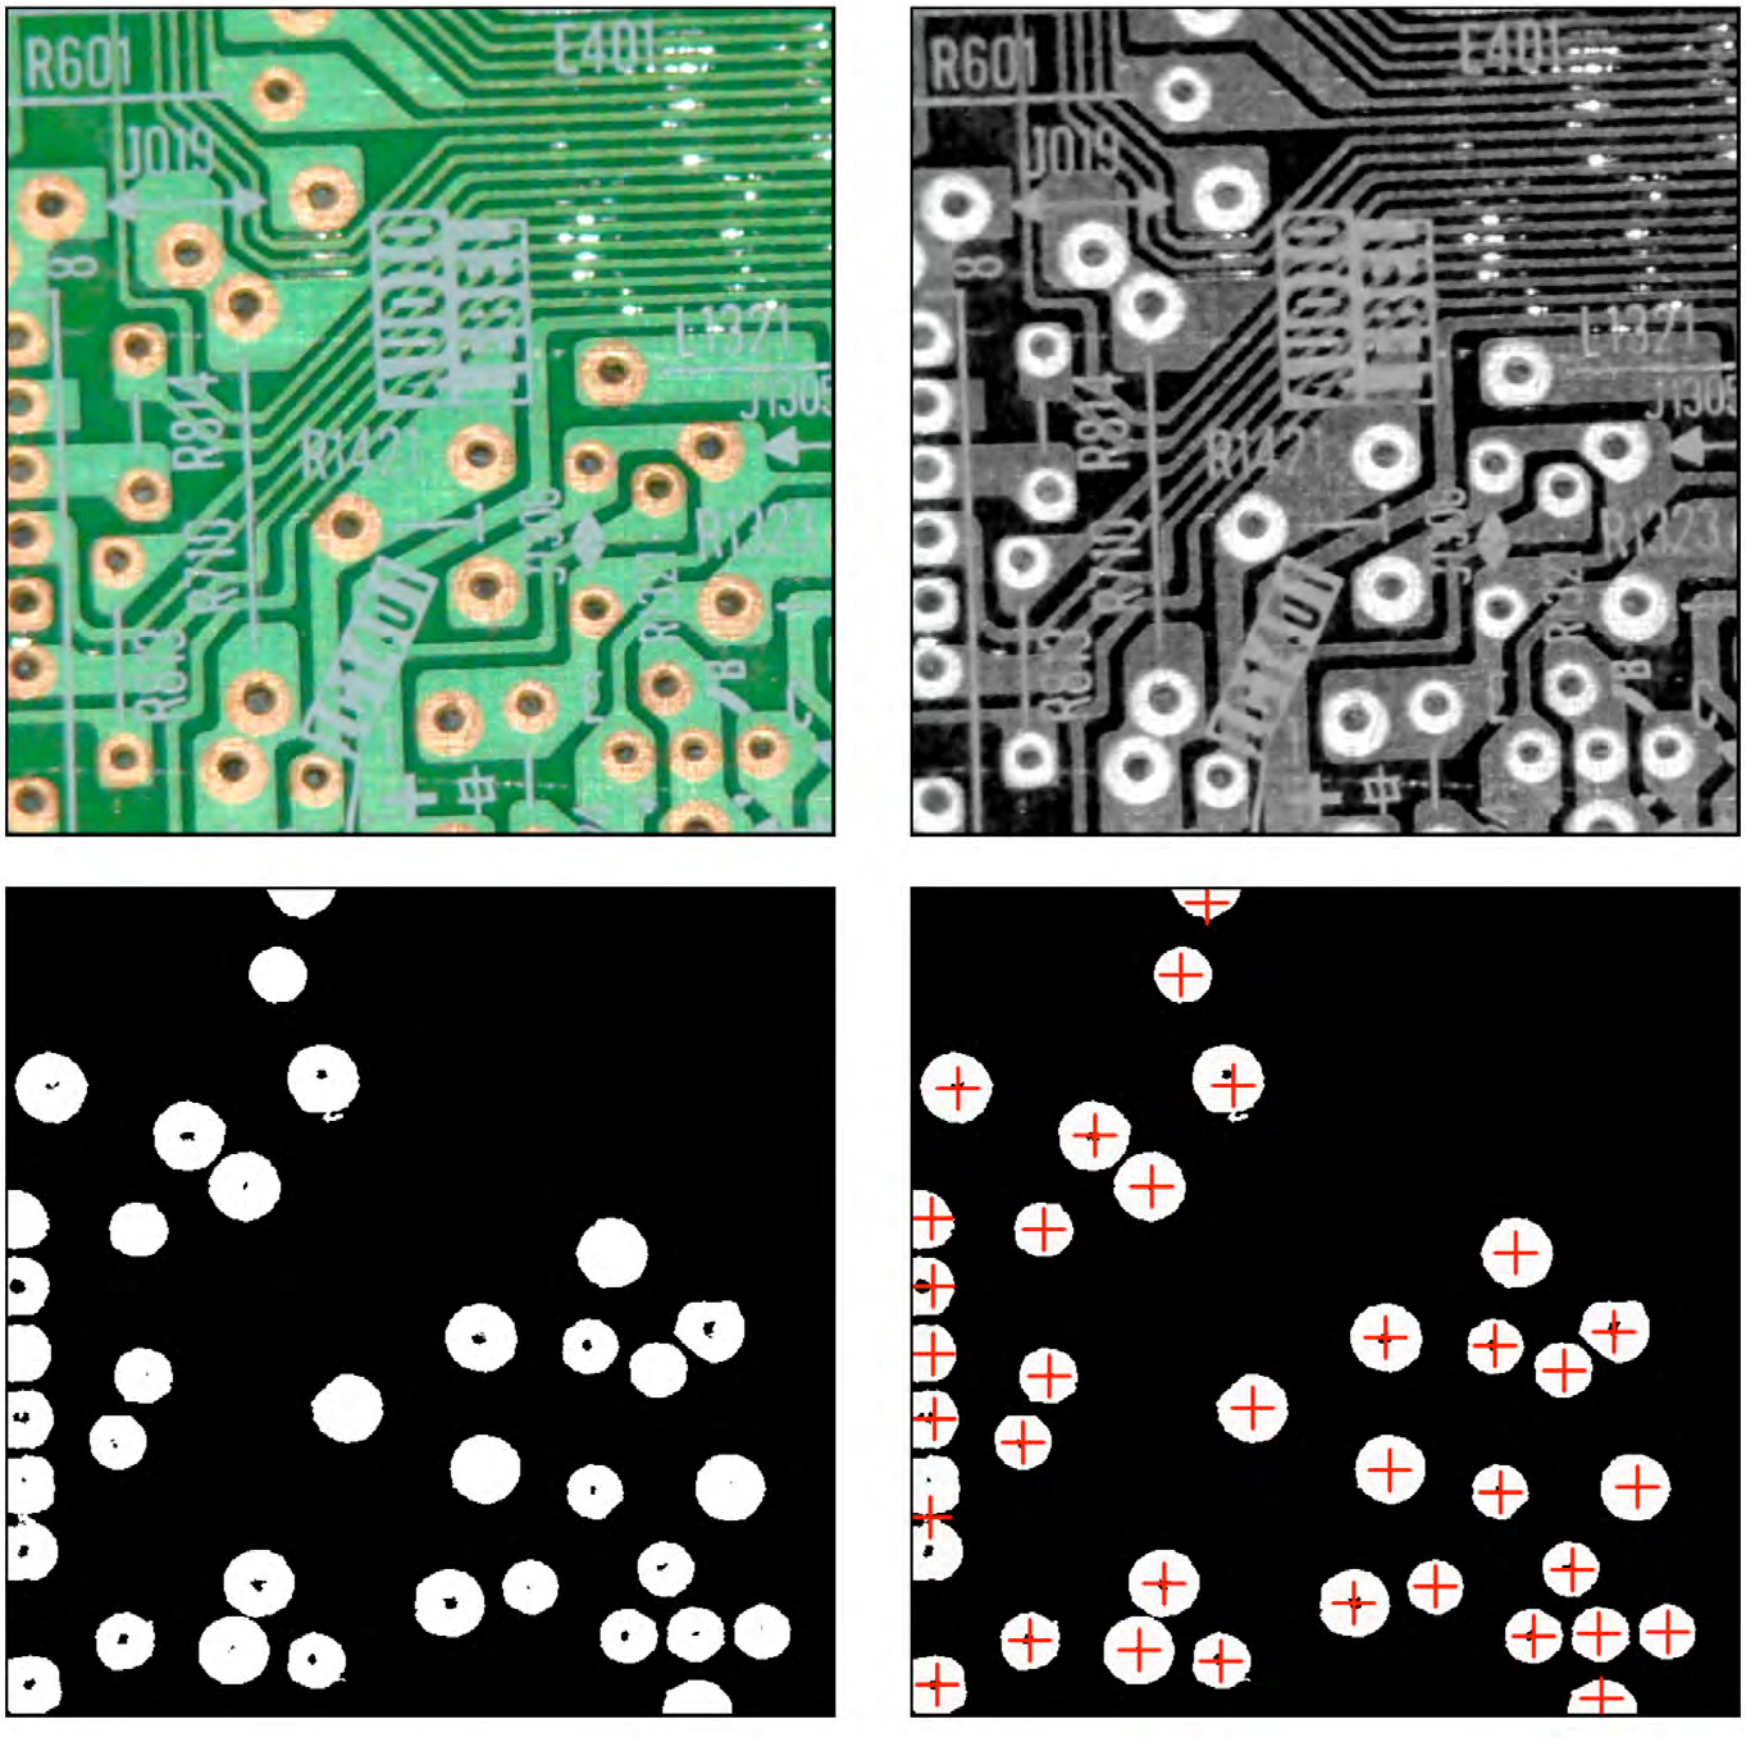

SCRATCHES (internal surface of an hydraulics cylinder)

(Top) Photograph. The bore (about 100 mm long) to be inspected is facing the camera and has a smooth surface with a near mirror-quality finish.

(Bottom-left) Image captured from a laser flying-spot scanner.

(Bottom-right) Scratches, detected using grey-scale morphology.

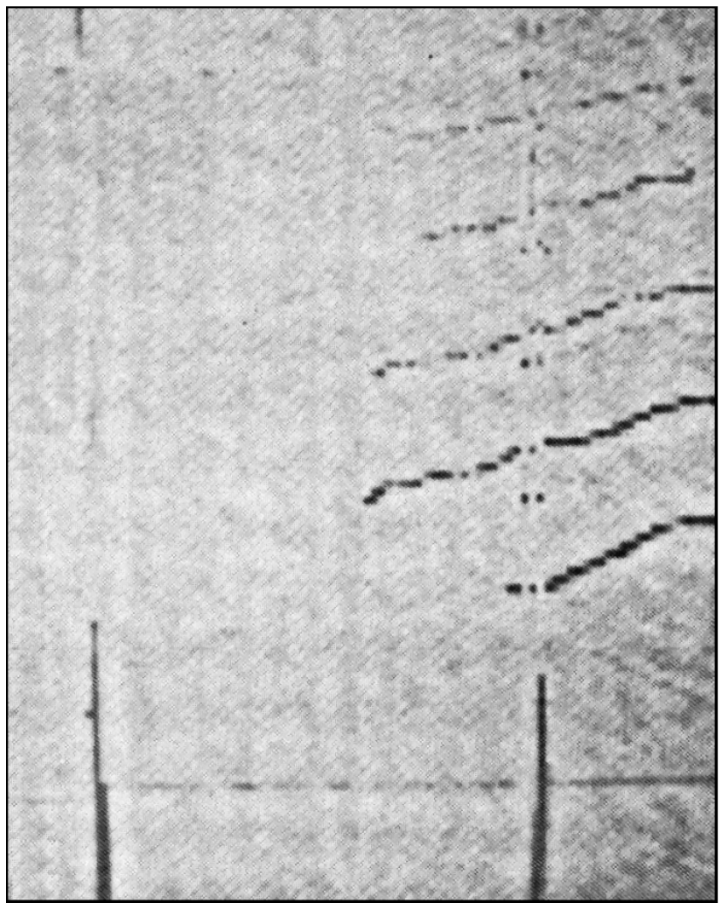

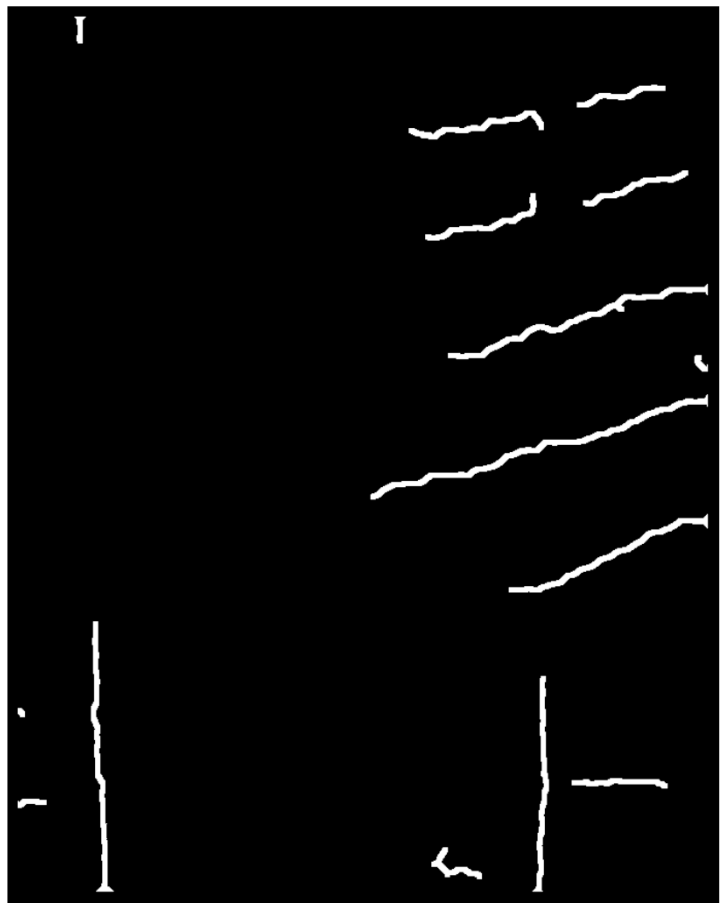

SAFETY CRITICAL INSPECTION (piston for car brake hydraulics system)

(Top) The assembly consists of four parts, including two plastic washers.

(Bottom) Enlarged views of the washers. The washer on the left has a straight profile, while the one on the right has a “kink” (arrowed). The correctness of assembly can be checked with simple linear scans along four lines (A - D).

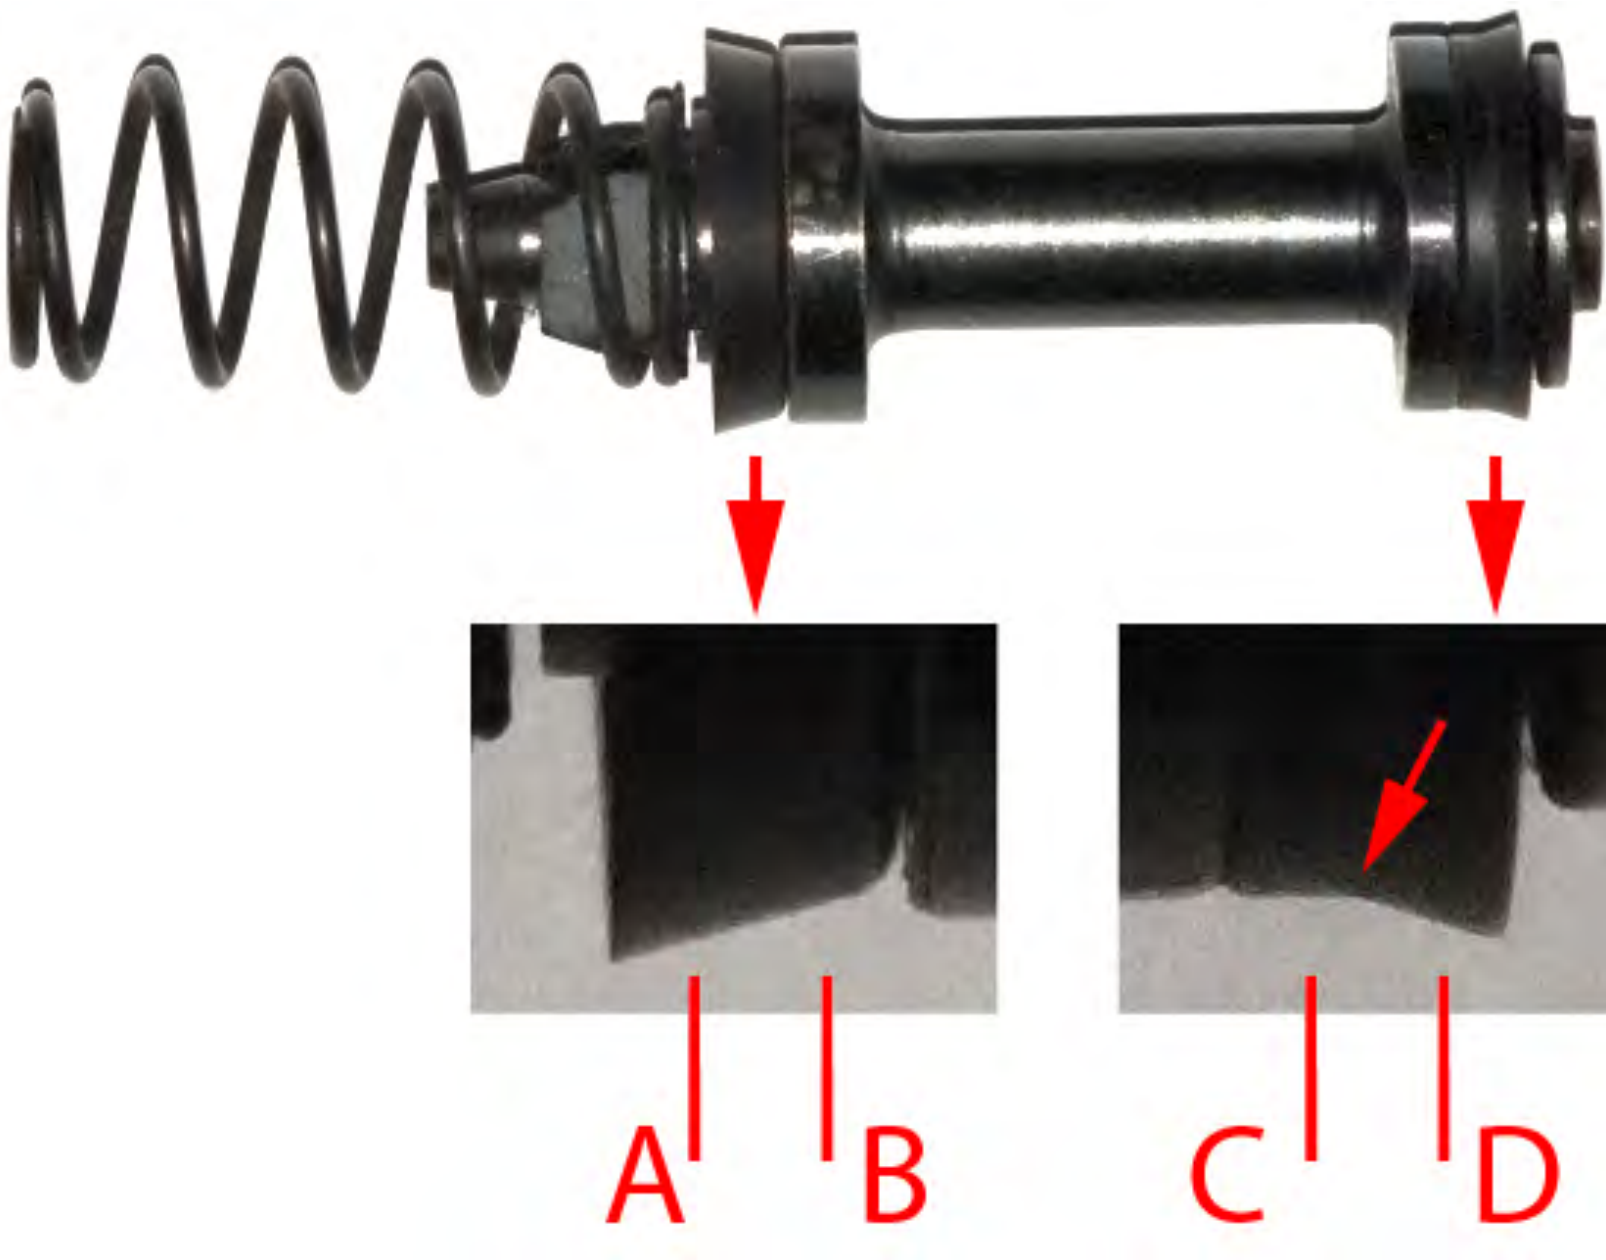

LIGHT BULB (loose wires`)

Occasionally, the internal wires to the "eyelet" surface contacts are not cropped properly, making the light-bulb potentially dangerous.

(Left) Original image.

(Right) The loose wire has been detected using grey-scale morphology

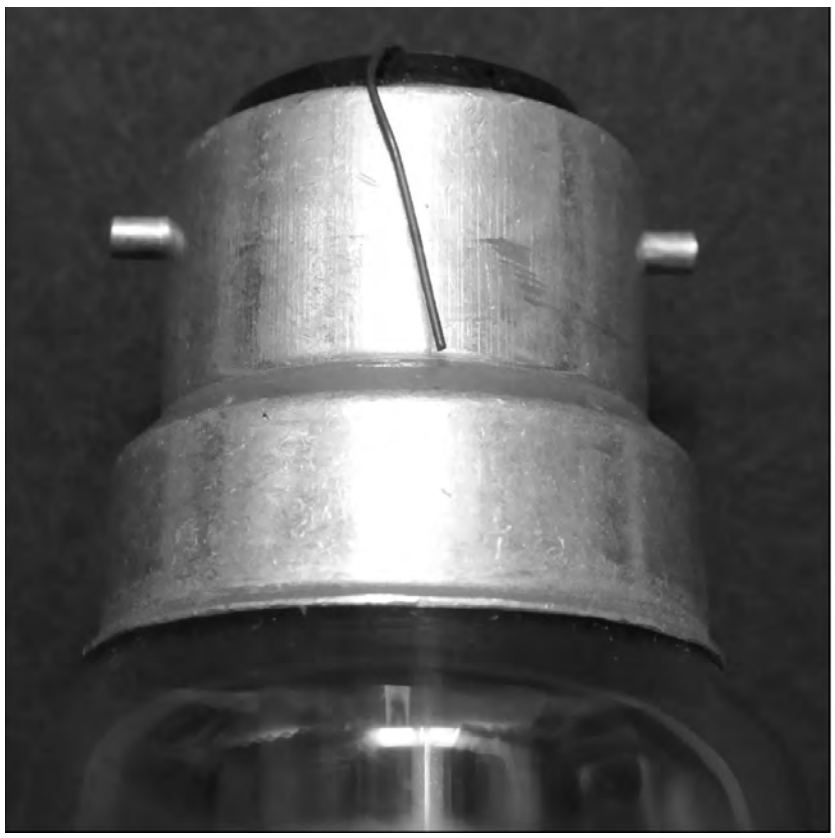

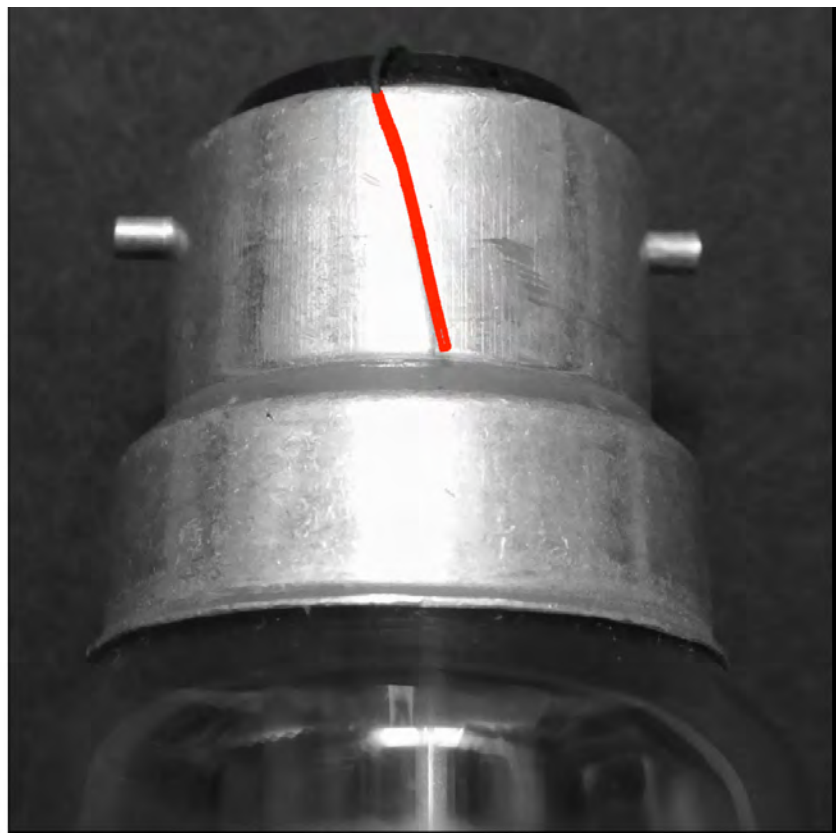

ASSISTING AUTOMATED PARTS -HANDLING (electric motor brushes)

Components, such as these carbon brushes can often be orientated roughly, using guide rails, placed just above a conveyor belt. However, to assemble them into a motor we must be sure that they are travelling “heads first”.

that they are not interlocked. Check that the area within the red contour is within the expected limits.

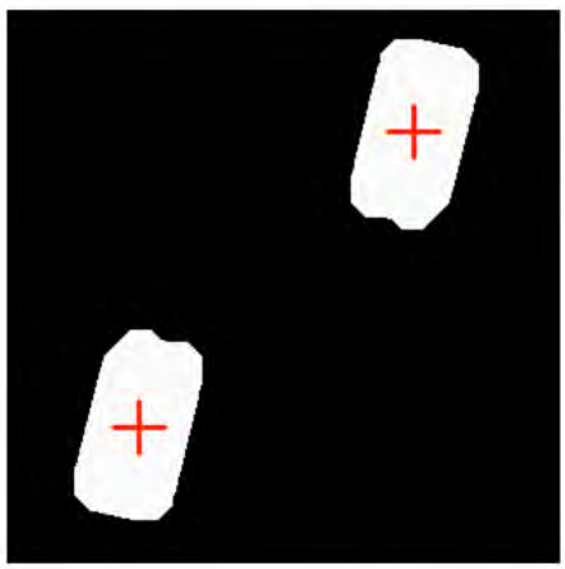

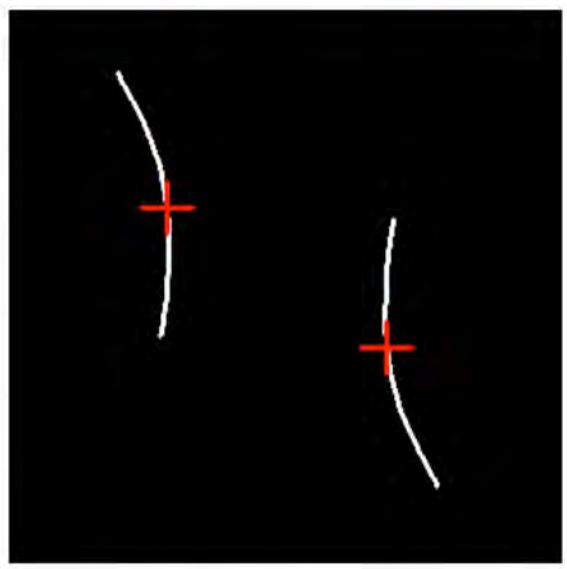

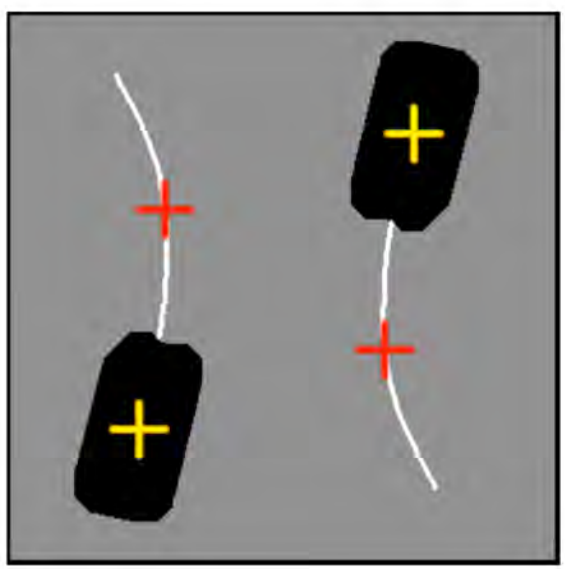

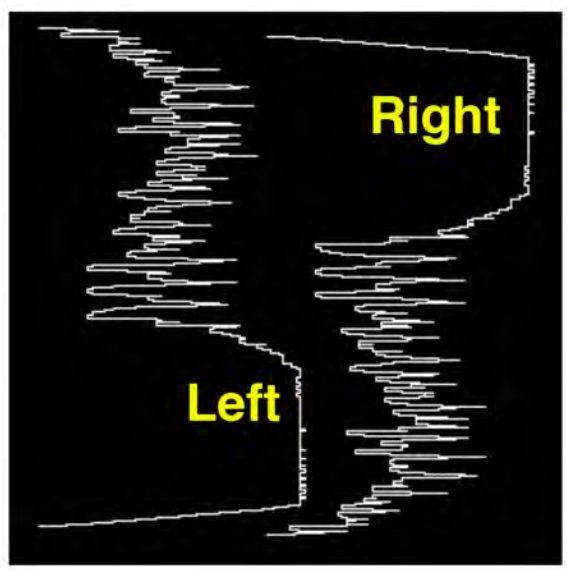

(Top-left) Original image.The brushes have been oriented roughly by guide rails above a conveyor belt.

(Top-right) Contours around each brush separately. To detect interlocking, measure the area within the red contour.

(Centre-left) The carbon blocks have bee isolated and the centroid of each one has been calculated.

(Centre-right) The springs have been isolated and points near the centre of each one has been calculated.

(Bottom-left) Merging [CL] and [CR] for ease of comparison.

(Bottom-right) Row-integration profiles.

COMPLEX OBJECT, MULTI-FEATURE INSPECTION (car brake assembly)

(Top-left) Original image.

(Centre) The area within the outer edge and its centroid.

(Bottom-left) A sub-image, centred on the cross in [TR], defines the first region-of-interest. (Call this **ROI-start**.)

(Bottom-right) Centroids of the holes in **ROI-start**. These allow us to calculate the position more precisely and determine the orientation..

The next page continues the story.

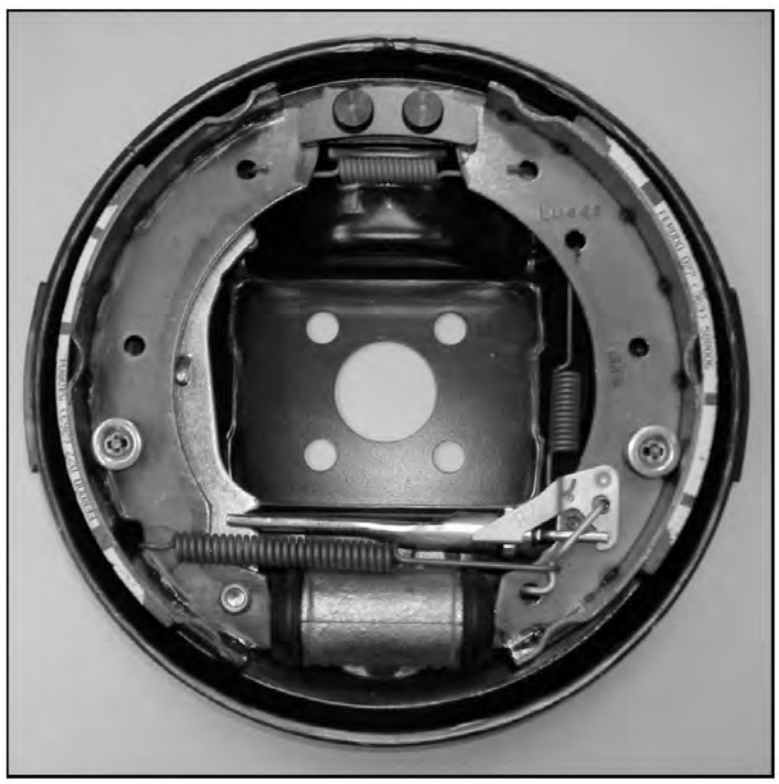

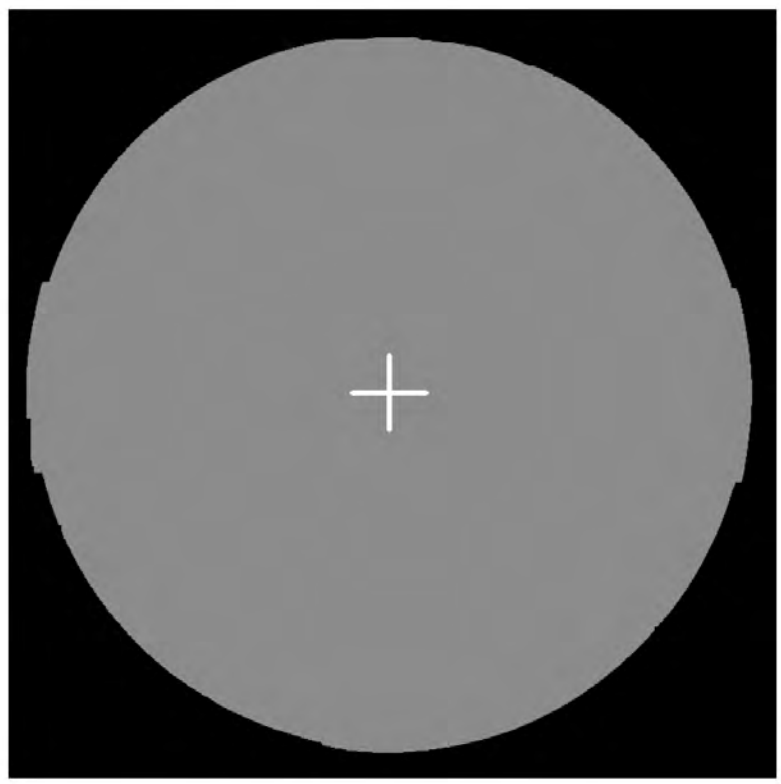

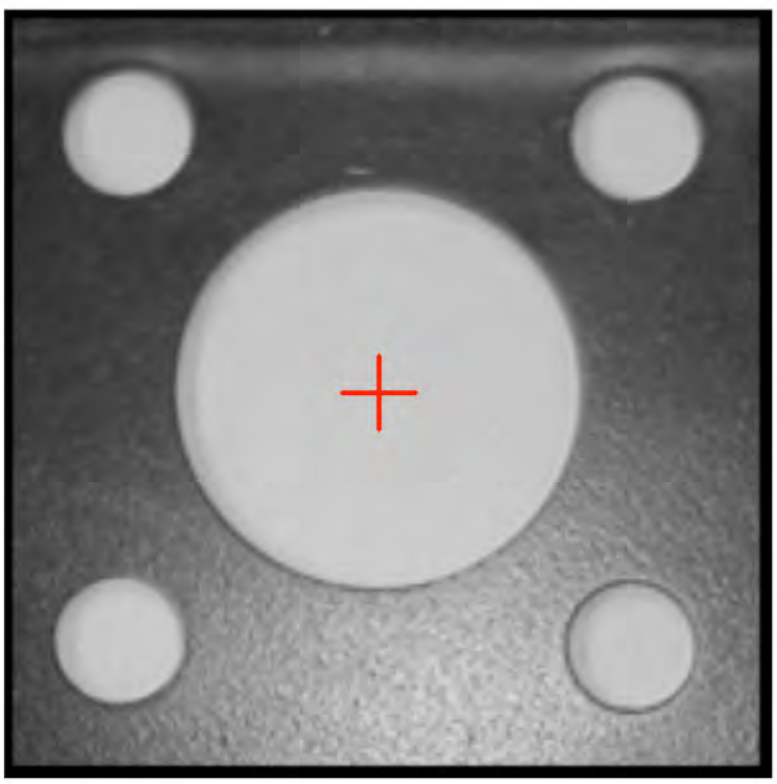

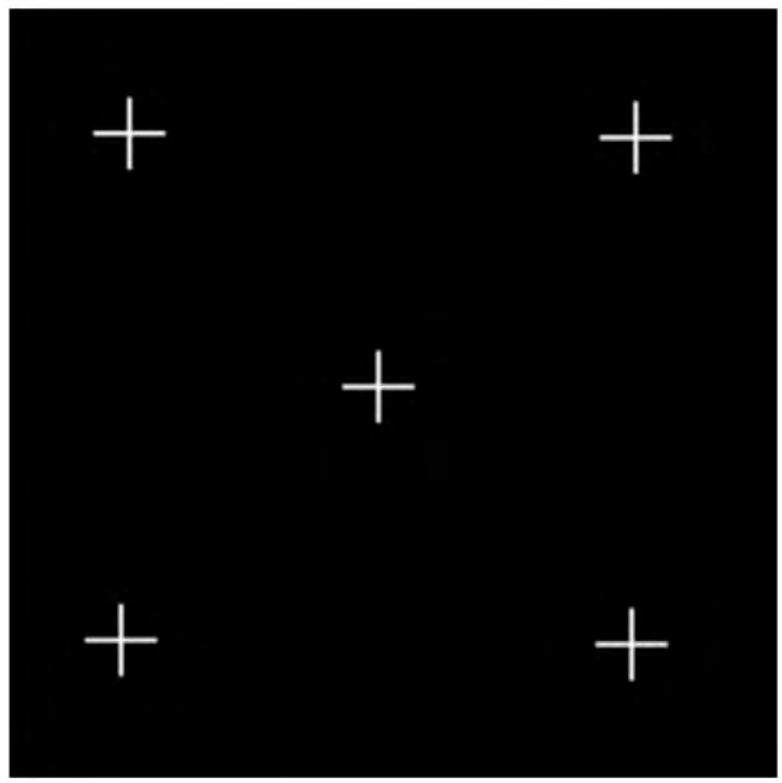

COMPLEX OBJECT, MULTI-FEATURE INSPECTION (continued)

Four additional regions of interest (red, **ROI 1** - **ROI 4**) have been identified following the initial placement of **ROI-start**.(yelllow) Each ROI is small and can be processed quickly. The processing within each ROI is usually very simple, because we know what to expect. We can reasonably assume that the assembly is held in place during inspection.

ROI 1 Spring. Each coil has been identified (vertical white stripes at the bottom of the lower image)

ROI 2 Rivets. Both have been checked and their positions calculated.

ROI 3 Cap of a spring (not visible). Checked and its position calculated.

ROI 4 A complex collection of levers and springs. These are checked by looking for their edges. Such a simple test is sufficient because we know what to expect.

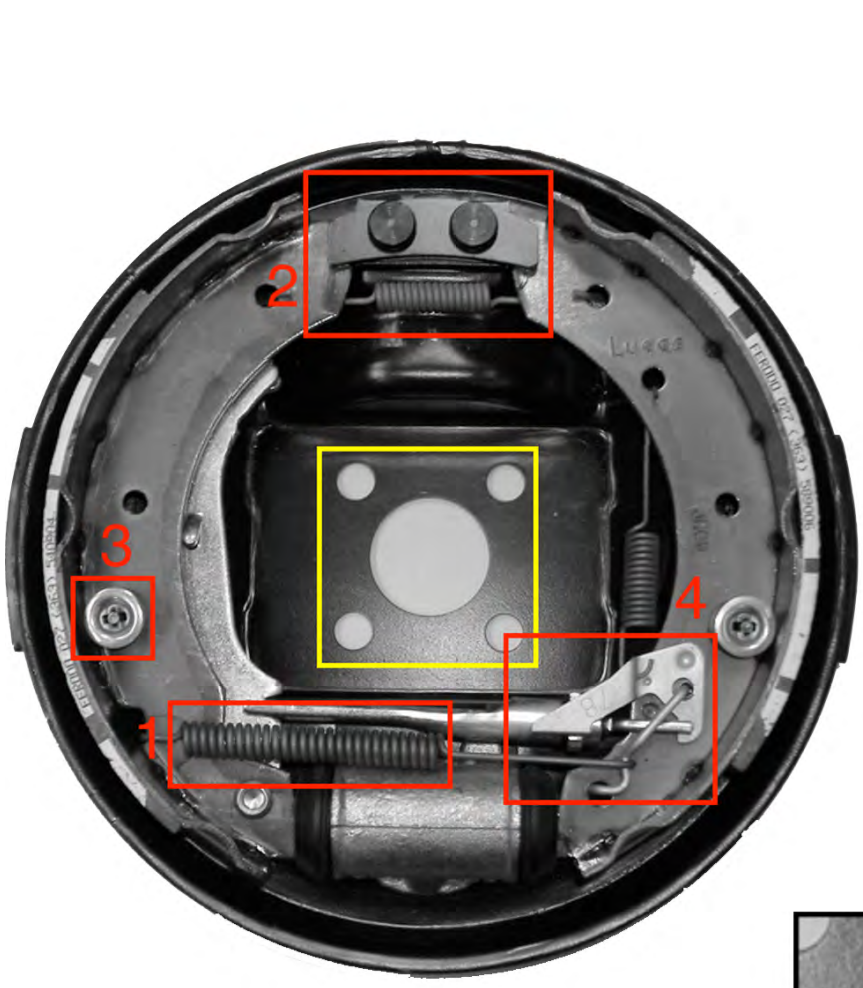

ROI 1

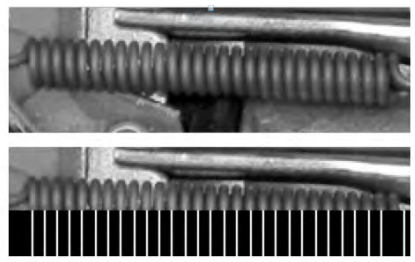

ROI 2

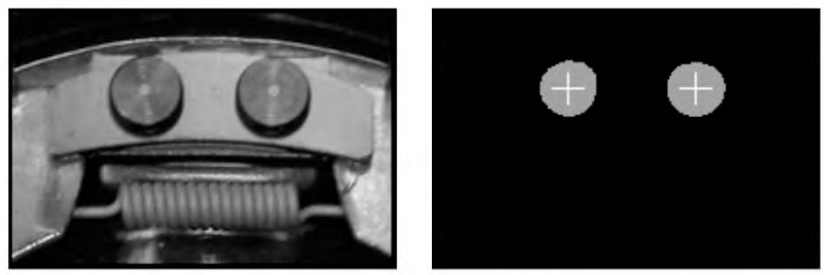

ROI 3

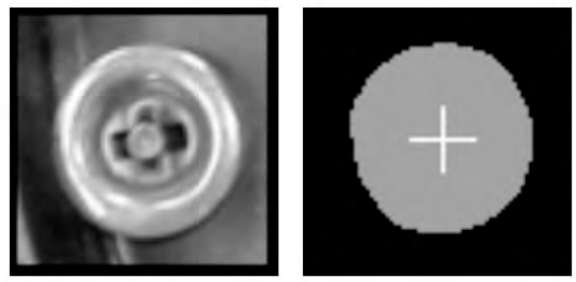

ROI 4

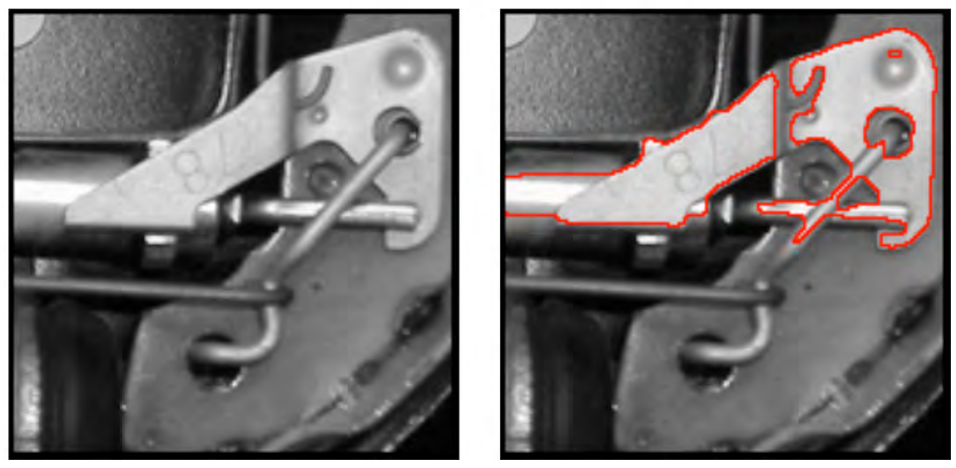

GUIDING A ROBOT (3½ degrees of freedom)

A SCARA robot (or a pick-and-place robot positioned over an (X,Y,θ)-table) could be used to pick up the con-rod. Three measurements ([X,Y]-position and orientation) are sufficient for this task. Two different gripper types are simulated: *surface-grasp* (i.e suction or magnetic, [CR]), or 2-finger *edge-grasp*. [BL]

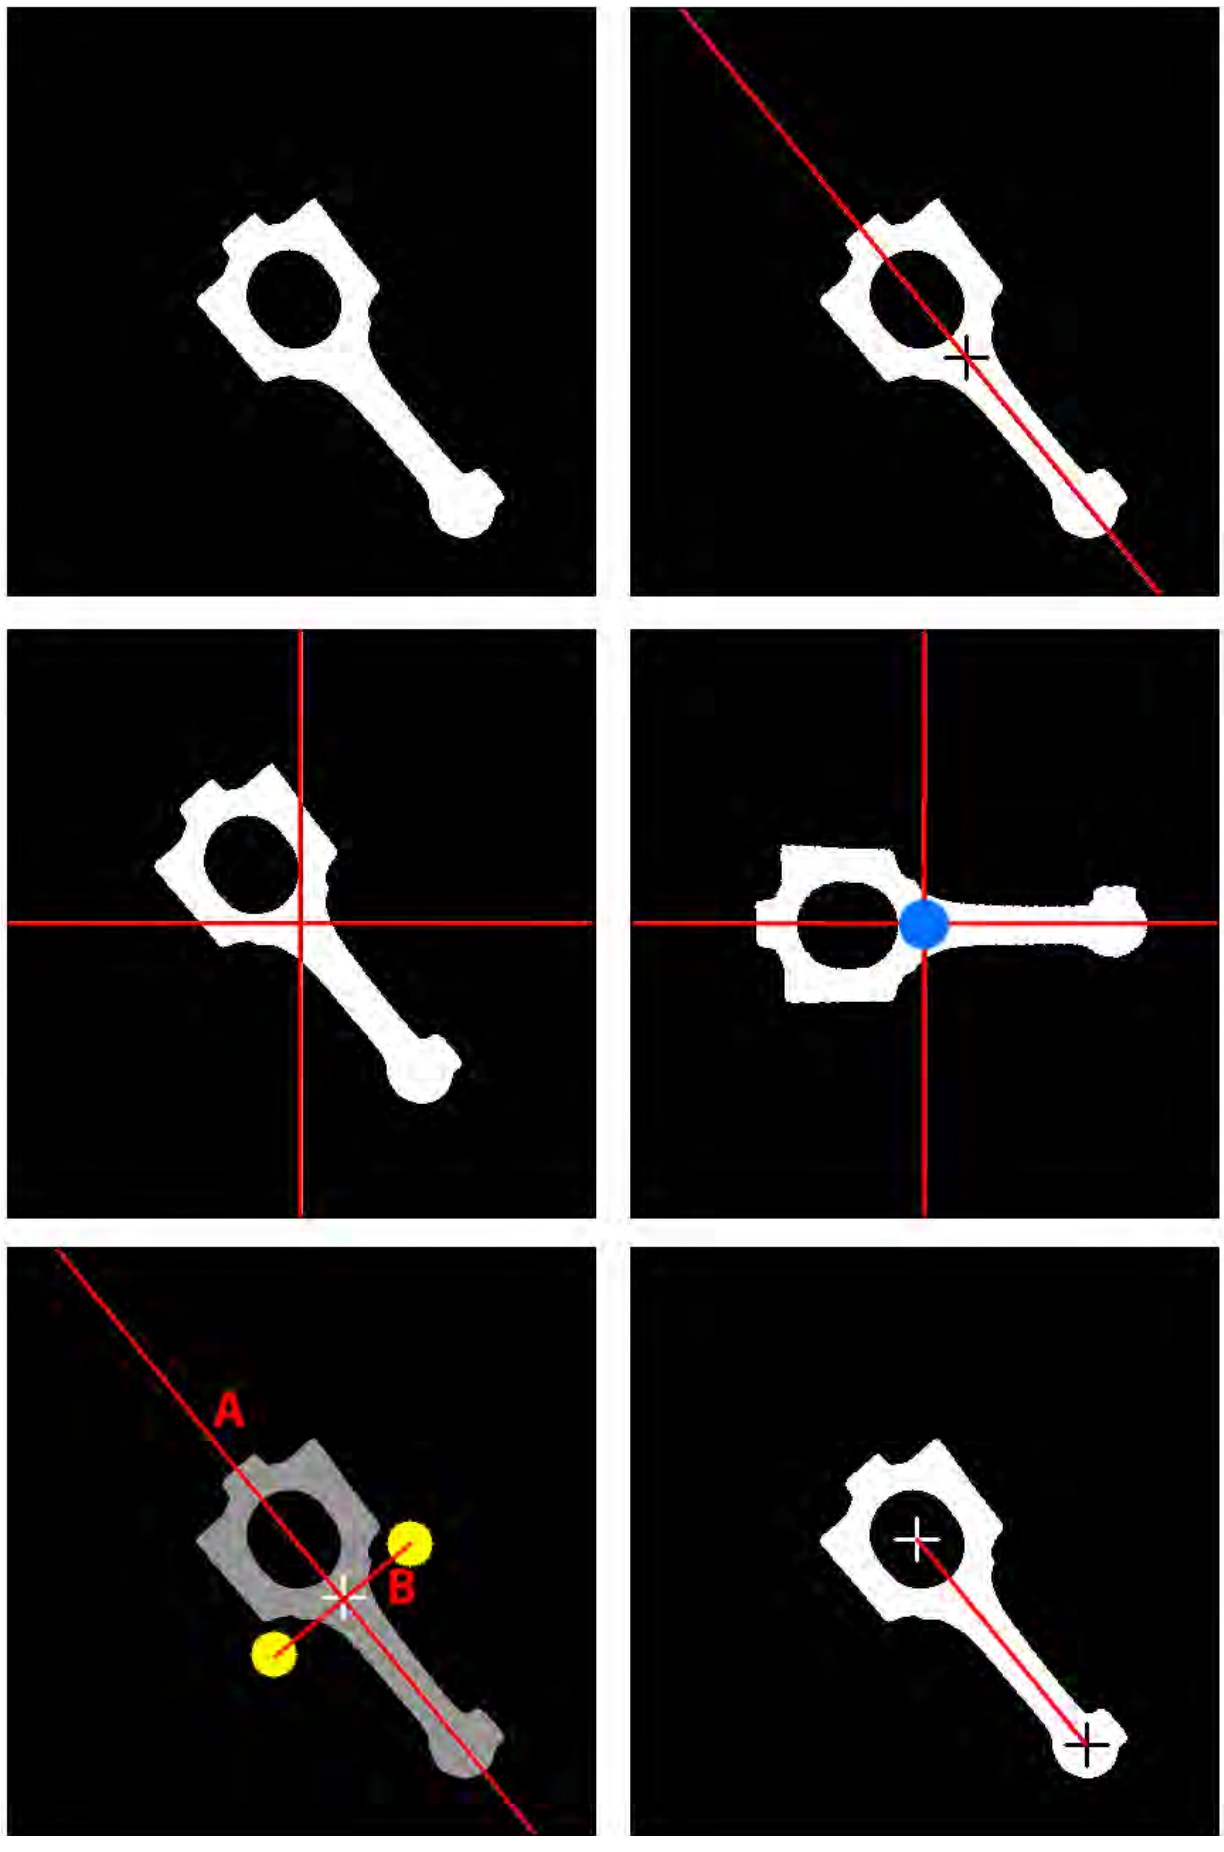

(Top-left) Original image (binary).

(Top-right) Centroid and principal axis.

(Centre-left) Normalising position so that the centroid is at the middle of the image.

(Centre-right) Orientation and position have both been normalised. Blue disc shows the position of a magnetic/ suction gripper. (If the gripper is larger than this, lifting is not safe.)

(Bottom-left) Placing a robot gripper with two fingers (yellow discs). Line B is at 90˚ to the principal axis (A) and passes through the centroid.

(Bottom-right) Alternative way to define orientation. Centroid of the largest lake and the centre of the inscribed circle at the “little end”

GUIDING A ROBOT (alternative ways of measuring position & orientation)

The centroid of a U-shaped object lies outside its boundary. For object 3 in [TL], the centroid lies close to its edge. Object 1 does not have a unique principal axis. The principal axis for object 3 is ill-defined; its behaviour is erratic, due to quantisation and camera noise. In these situations, other ways to measure position and orientation are needed.

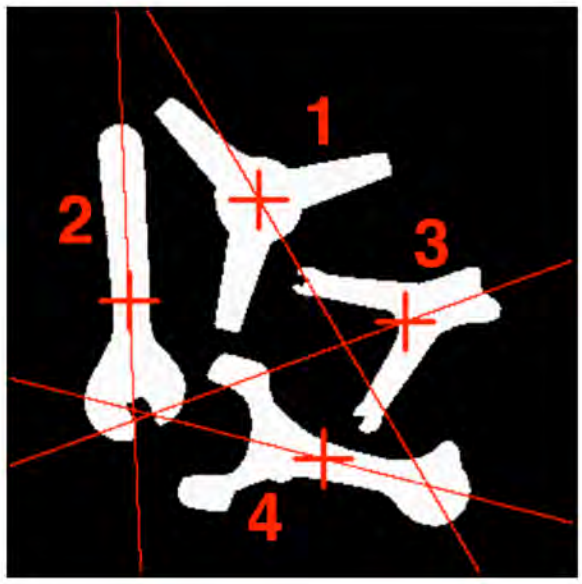

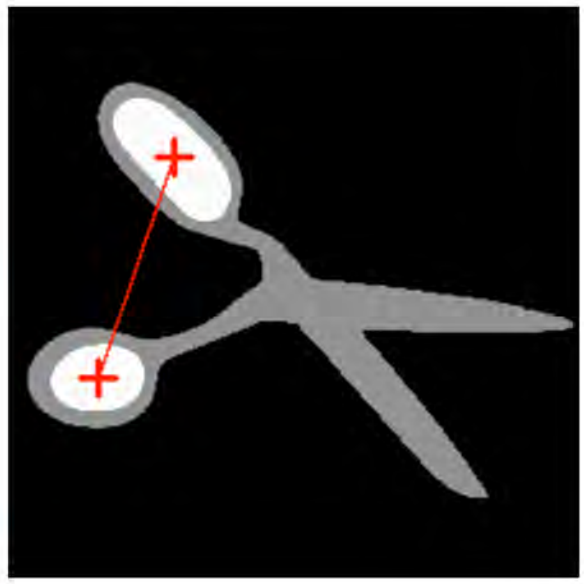

(Top-left) 4 flat objects

.**(Top-right)** Centroids of finger holes define both position & orientation

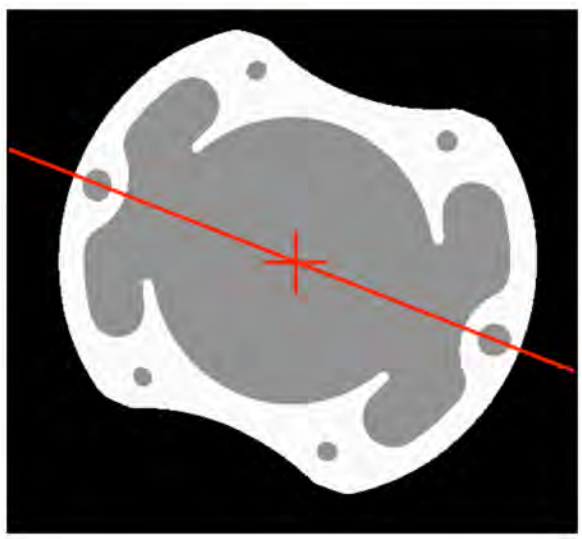

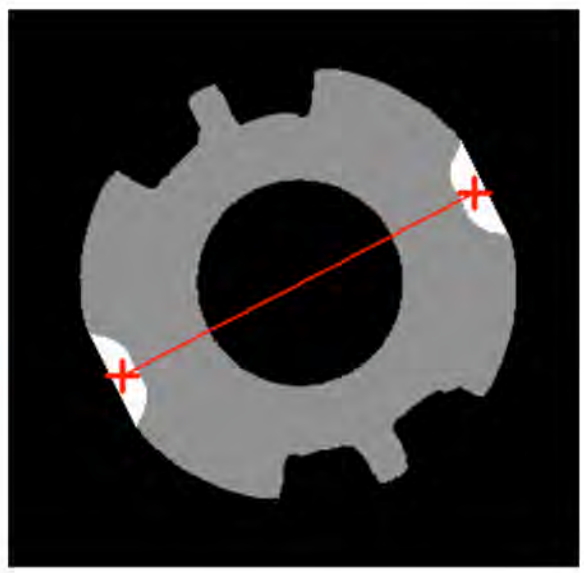

(Centre-left) Stamping for an electric motor. Position & orientation defined by centroids of 2 holes, selected by size.

(Centre-right) Ferrite core for electronic circuit. Centroids of bays define position & orientation.

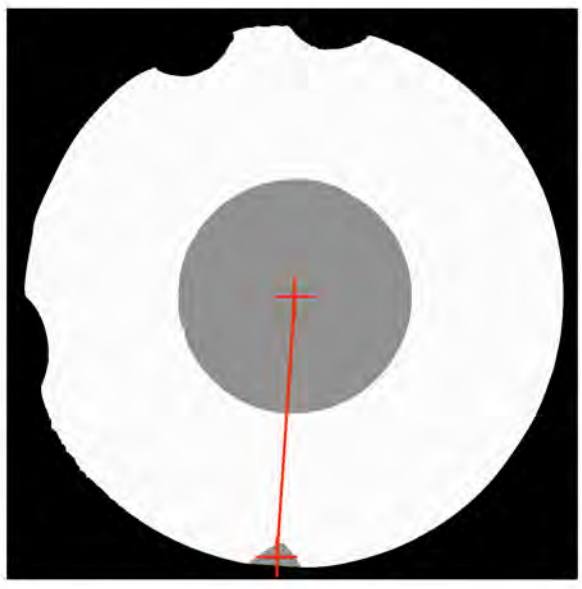

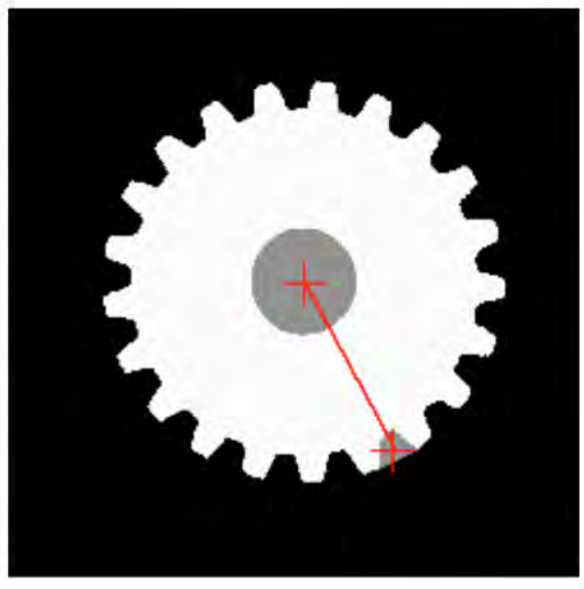

(Bottom-left) Alloy casting for a clutch. Orientation is defined by the notch-shaped bay.

(Bottom-right) Gear. Any bay will suffice to define orientation.

FINDING ORIENTATION OF CIRCULAR OBJECT

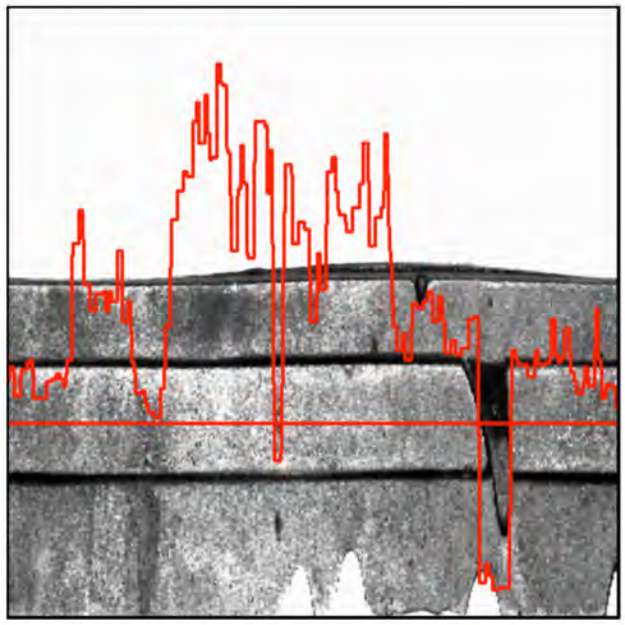

(Top-left) Alloy casting, tractor clutch component

.**(Top-right)** Transformation to polar coordinates.

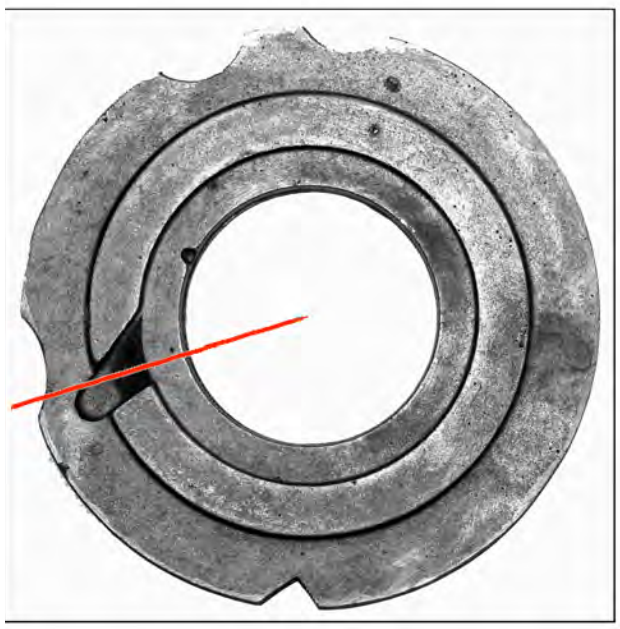

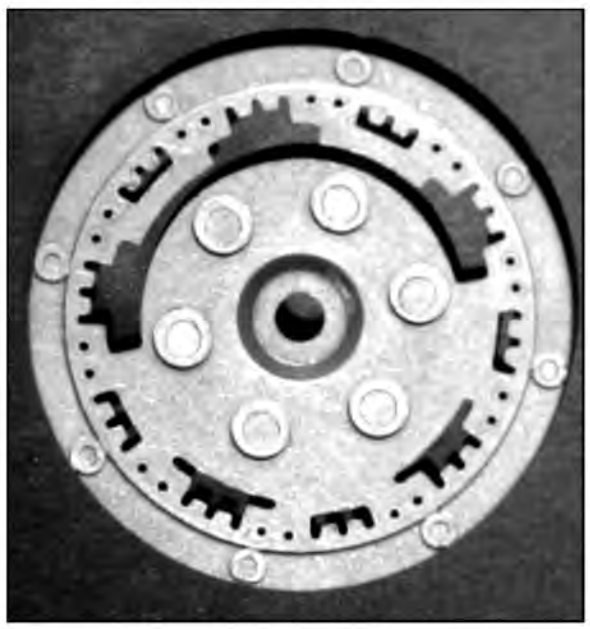

(Centre-left) Orientation defined by centroid and dark surface feature.

(Centre-right) Alloy cam.

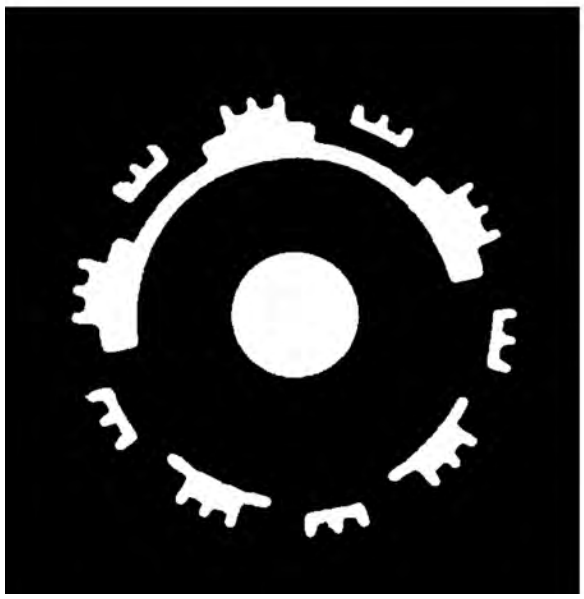

(Bottom-left) Lakes derived from binary version of [CR]..

(Bottom-right) Centroids of the lakes. The largest lake can be identified easily, to obtain a unique measure of orientation.

ROBOT HANDLING COIL WITH LOOSE WIRES

One plan-view and four side-cameras are are required. Their relative positions must be known exactly. A robot with a 2-finger gripper is used.

(Top-left) Original image. Plan view of the coil.

(Top-right) Ends of the wires have been located. (A and B) Another point (C), near to A and lying along the wire, has also been found. The line AC defines both position and orientation of the end of the wire. The finger-tips of a robot, as it is about to grasp the other wire (near B), are indicated by mauve bars.

(Bottom-left) Every side-view camera has blind-spots, where the wire ends are either hidden or obscured. A camera looking along line P can see A but B is hidden /obscured.

(Bottom-right) Four side-view cameras are needed to locate wire ends, wherever they are. Their blind regions are indicated by the colour stripes. Deciding which side-view camera provides the best view(s) requires the plan-view image.

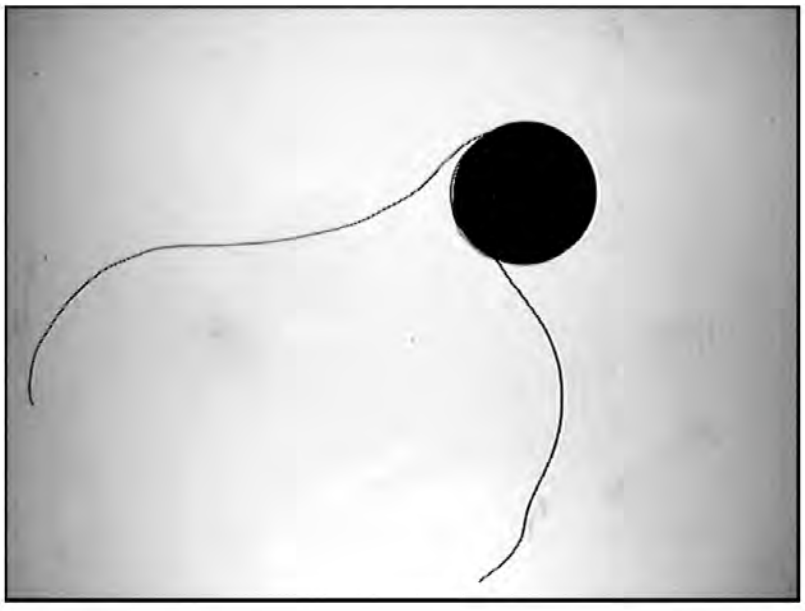

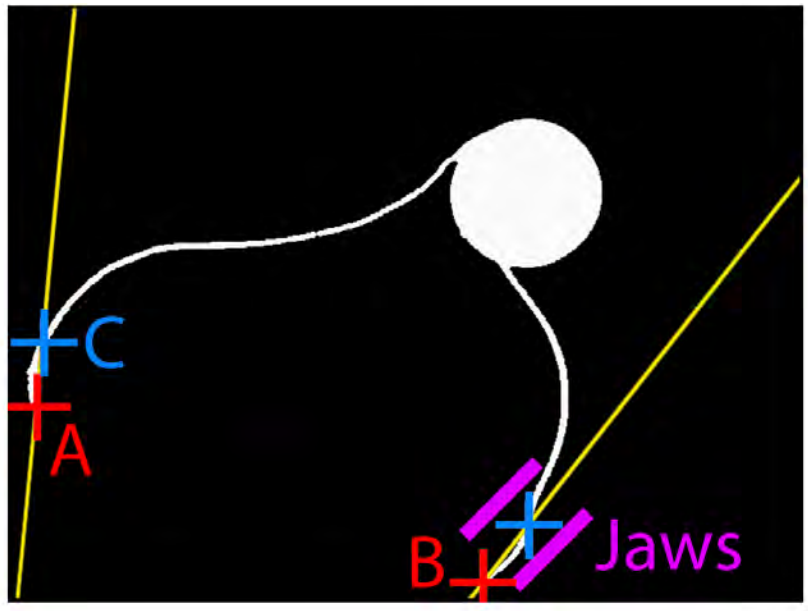

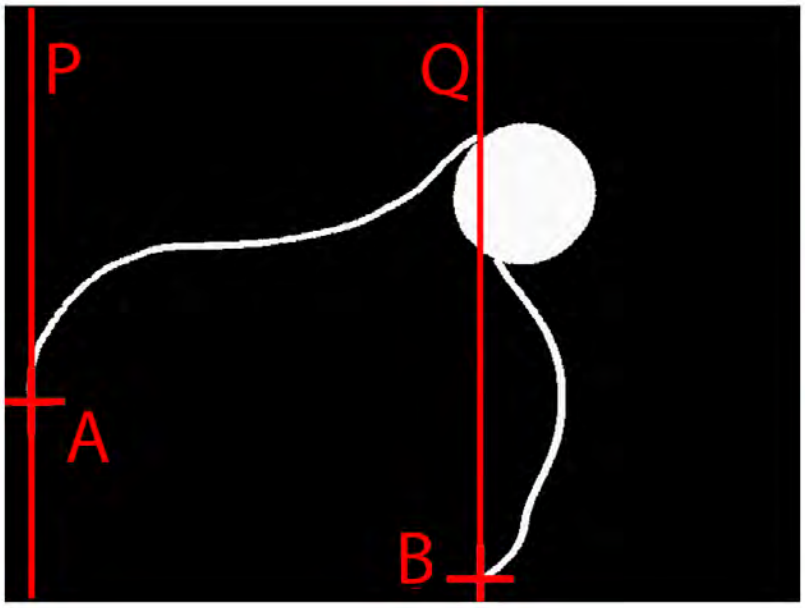

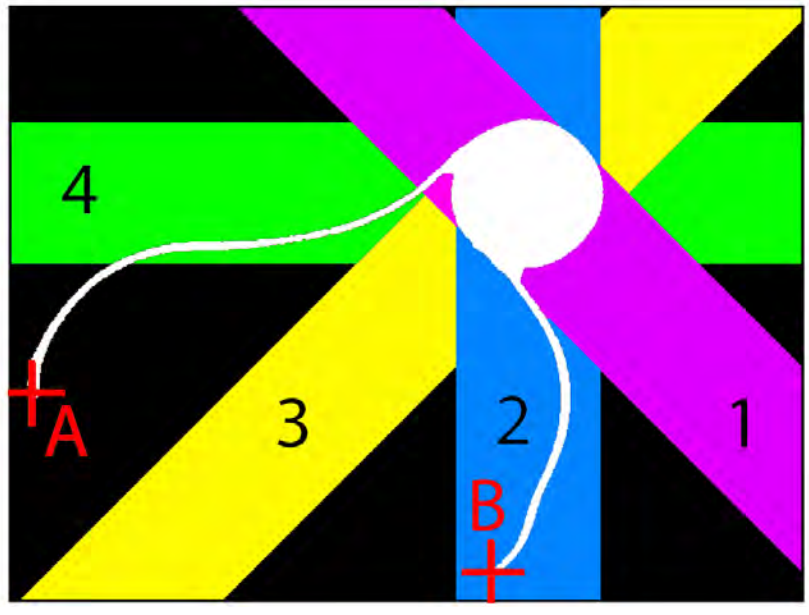

ROBOT HANDLING (reconciling views from different cameras)

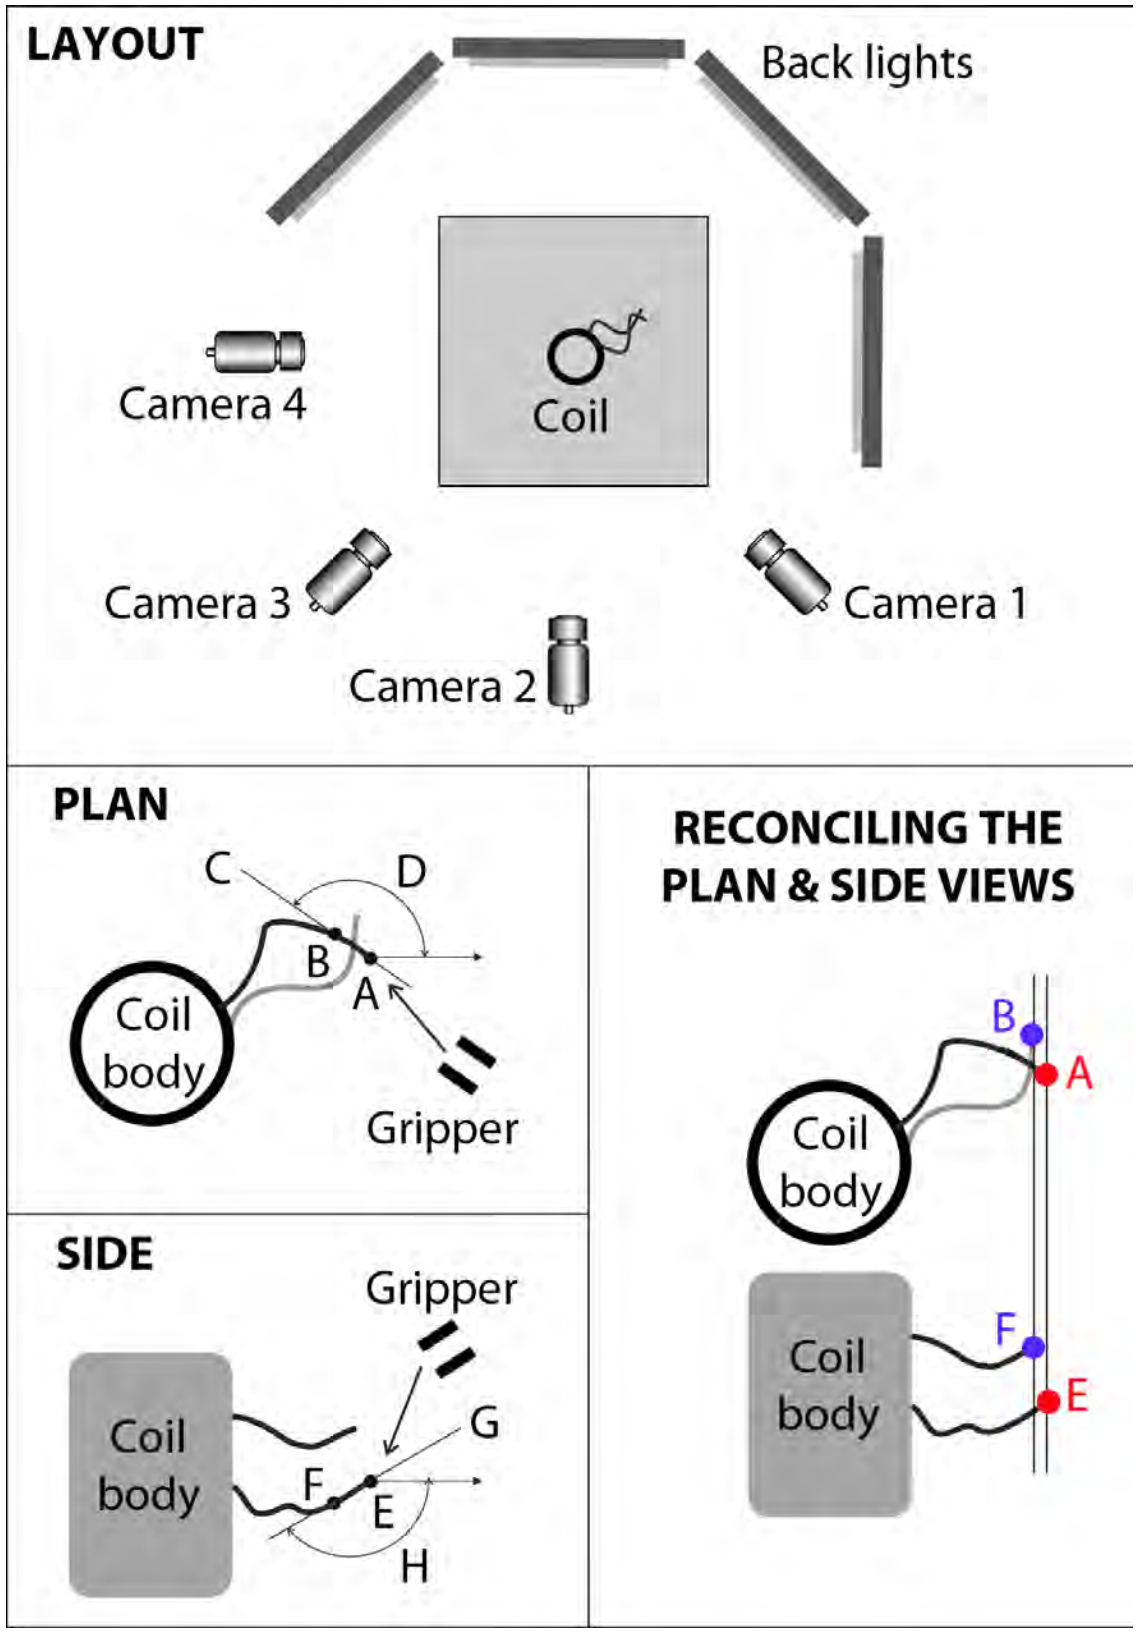

(Top) Layout, view, as seen from the (unseen) plan camera. The coil is viewed against a back-lighting unit (grey box).

(Centre-left) Gripper orientation (D) in the horizontal plane is calculated from the positions of A (end of lead) and B (another point near A and also lying along the wire).

(Bottom-left) The best side-view camera has been selected by analysing the plan view. The gripper orientation (H) in the vertical plane is calculated from positions of E (end of lead) and F (another point on the wire near E)

(Bottom-right) Aligning individual wire ends visible in the plan and side views.

HIGH-VOLUME, LOW-VALUE PARTS (battery-top plastic mouldings)

This and the three following inspection applications demand simple, high-speed, low-cost processing. Although they are simple components, they may be are critical within a larger, more complex and valuable assembly**.** Inspecting battery tops requires only simple measurements, such as area, and Euler number. (This is a fast calculation and counts the number of blobs minus the number of lakes.)

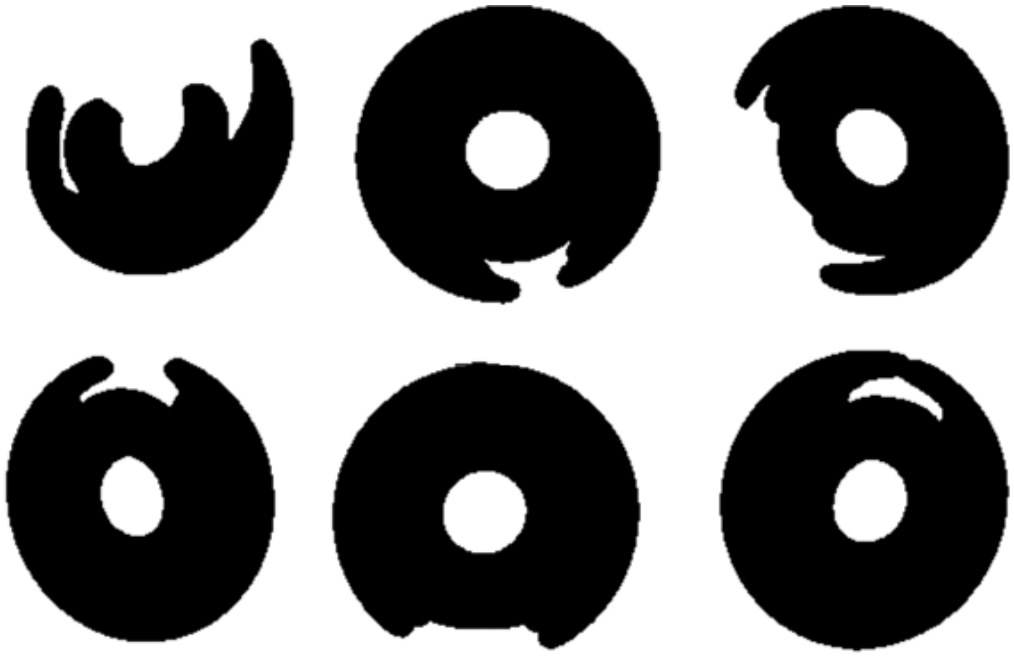

(Top) Battery tops with short-feed moulding faults.

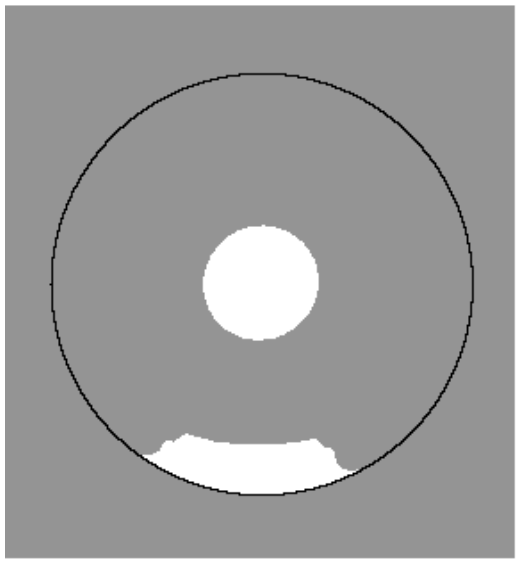

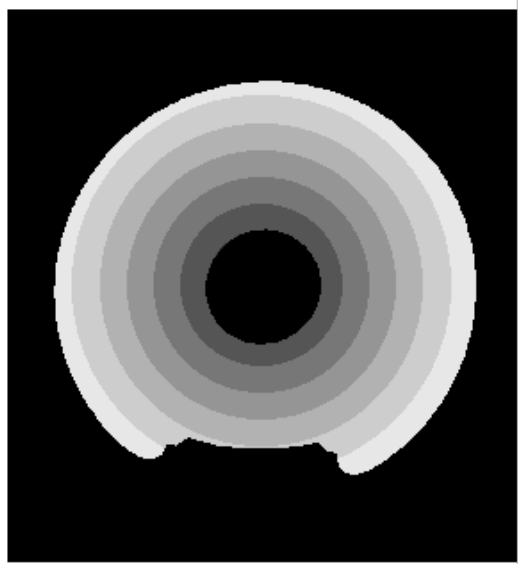

(Bottom-left) Circle fitted to the outer edge. Differences between the silhouette and the fitted circle are shown in white.

(Bottom-right) Determining the distribution of radial distances from the centroid. For clarity of illustration, the intensity is shown as

HIGH-VOLUME, LOW-VALUE PARTS (switch plastic mouldings)

Plastic mouldings for a small electrical switch. Three have moulding faults due to the presence of “flash”. (Molten plastic has flowed between the two halves of the die during the moulding process, leaving a thin “leaf” of unwanted material.) The flash is fragile and will readily break off, ;possibly jamming the switch or the swich assembly machine.

Acceptable

Convex deficiency too large

Too few lakes

Bays too large

HIGH-VOLUME, LOW-VALUE PARTS ("nail", part of a pop rivet)

A pop rivet (also called a blind rivet) has two components: the "nail" and sleeve. If the nail has not been made properly it may have a blunt or fan-tail end. This may jam the assembly machine, causing a severe loss of production.

When I studied this application 1980, the factory made 23,000,000 nails per week. The company staff knew that they were unavoidably discarding £700 worth (equivalent to £1980 in 2018) of "good" nails each week because it was not possible to detect a tiny proportion of faulty products within a large batch. Loss of production and clearing machine jams were additional costs, not included in this figure.

Inspection requires that the nails be presented in the same orientation each time. Simply measuring the chord-length across the nail close to its end is sufficient to determine when the nail is defective.

Notice that the parts-handling mechanism of the inspection system must be able to accept defective nails without jamming

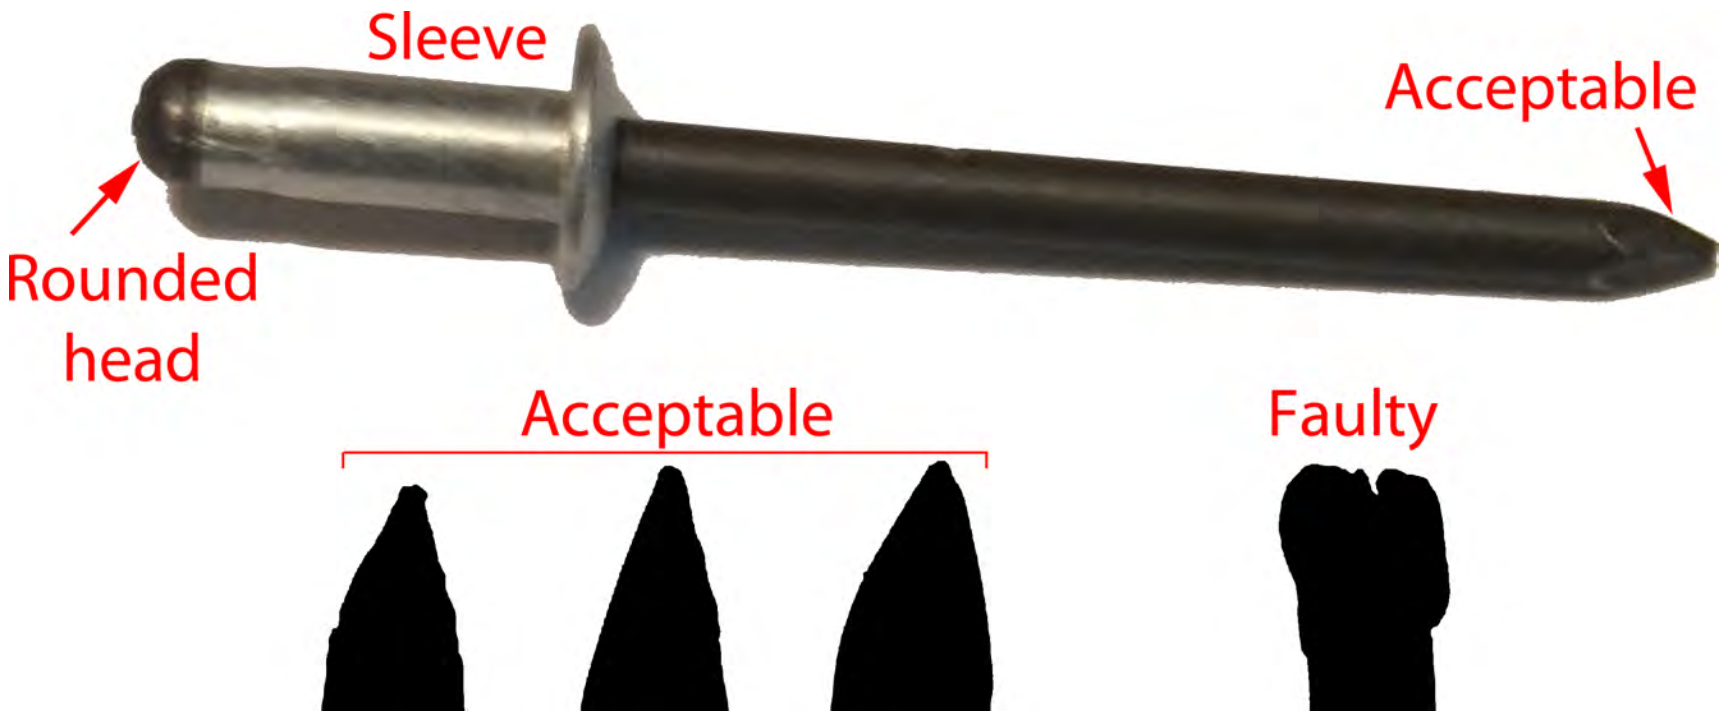

EXTERNAL SCREW THREAD

This is an exercise in computational geometry. Measurements required:

1. Flank angle.

2. Radius of curvature of the peak & trough of each thread.

3. Depth of thread.

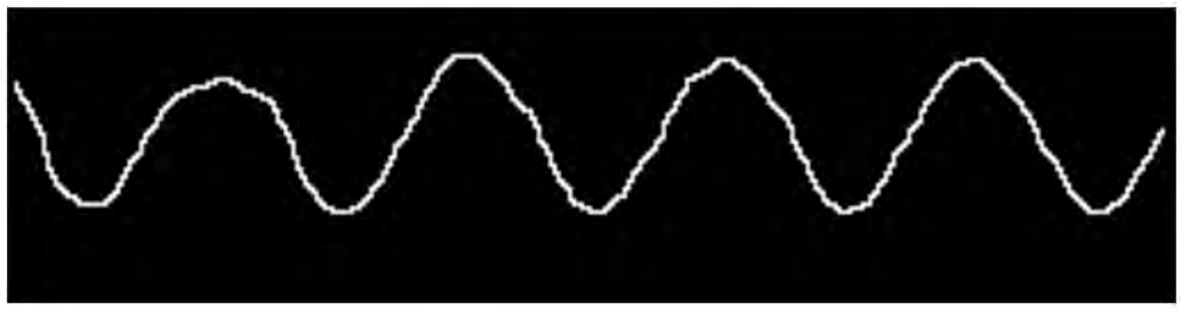

(Row 1) Edge of the thread.

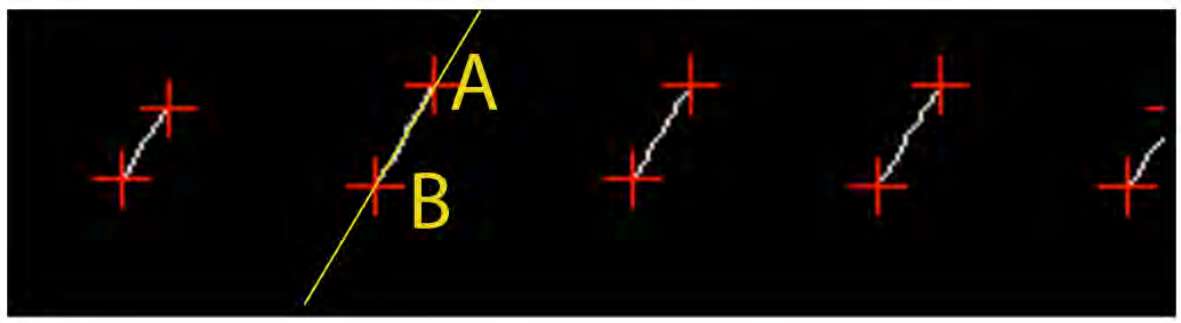

(Row 2) Flank (/). Detected using binary morphology.

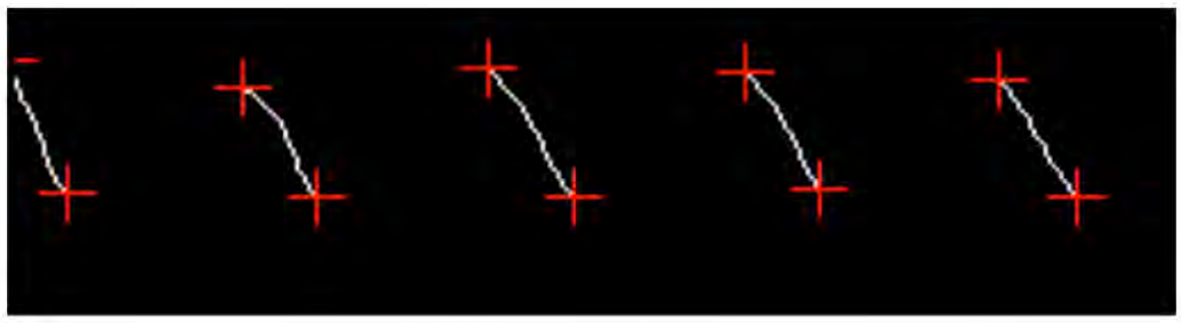

(Row 3) Flank (\). Again, detected using binary morphology.

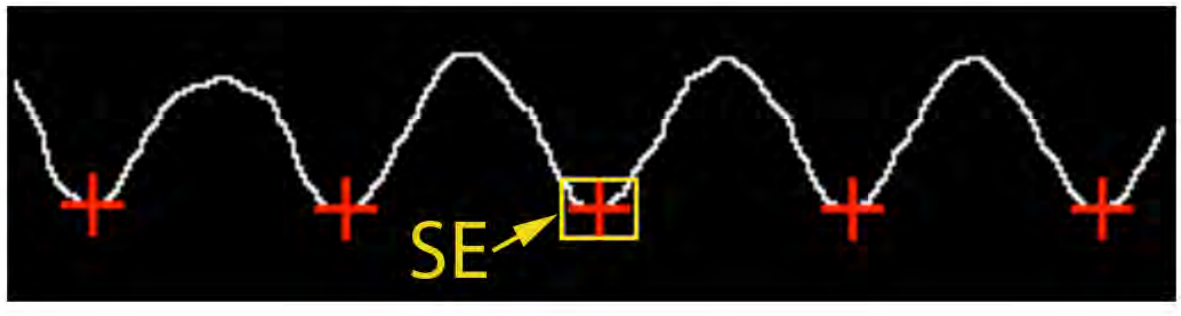

(Row 4) Troughs, found using binary morphology. Structuring element (SE) is defined by that part of the edge inside the yellow box.

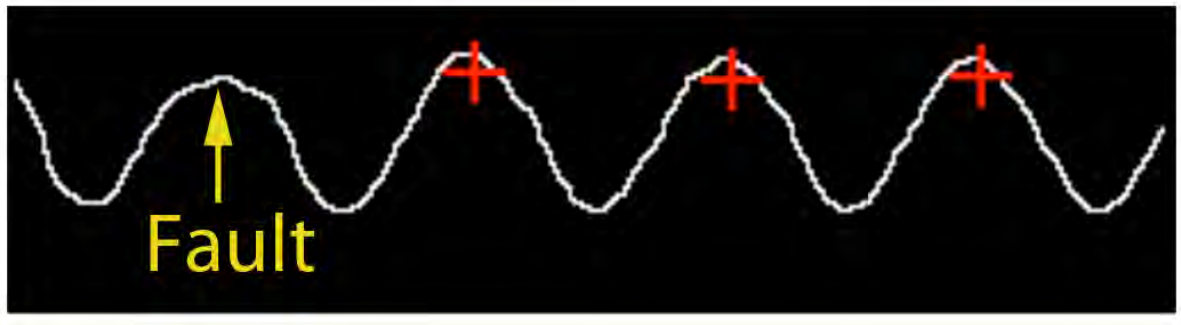

(Row 5) Peaks, found in a similar way.

INTERNAL SCREW THREAD (viewed obliquely from outside)

(Top-left) Original image.

(Top-right) Grey-scale dilation, followed by erosion.

(Centre-left) [TL] subtracted from [TR].

(Centre-right) [CL] after contrast enhancement.

(Bottom-left) [CR] after thresholding.

(Bottom-right) Small blobs in [BL] were removed and some noise reduction applied. The resulting contours have been superimposed on the original. Inspection requires that these contours are all present, intact and have the correct radius of curvature.

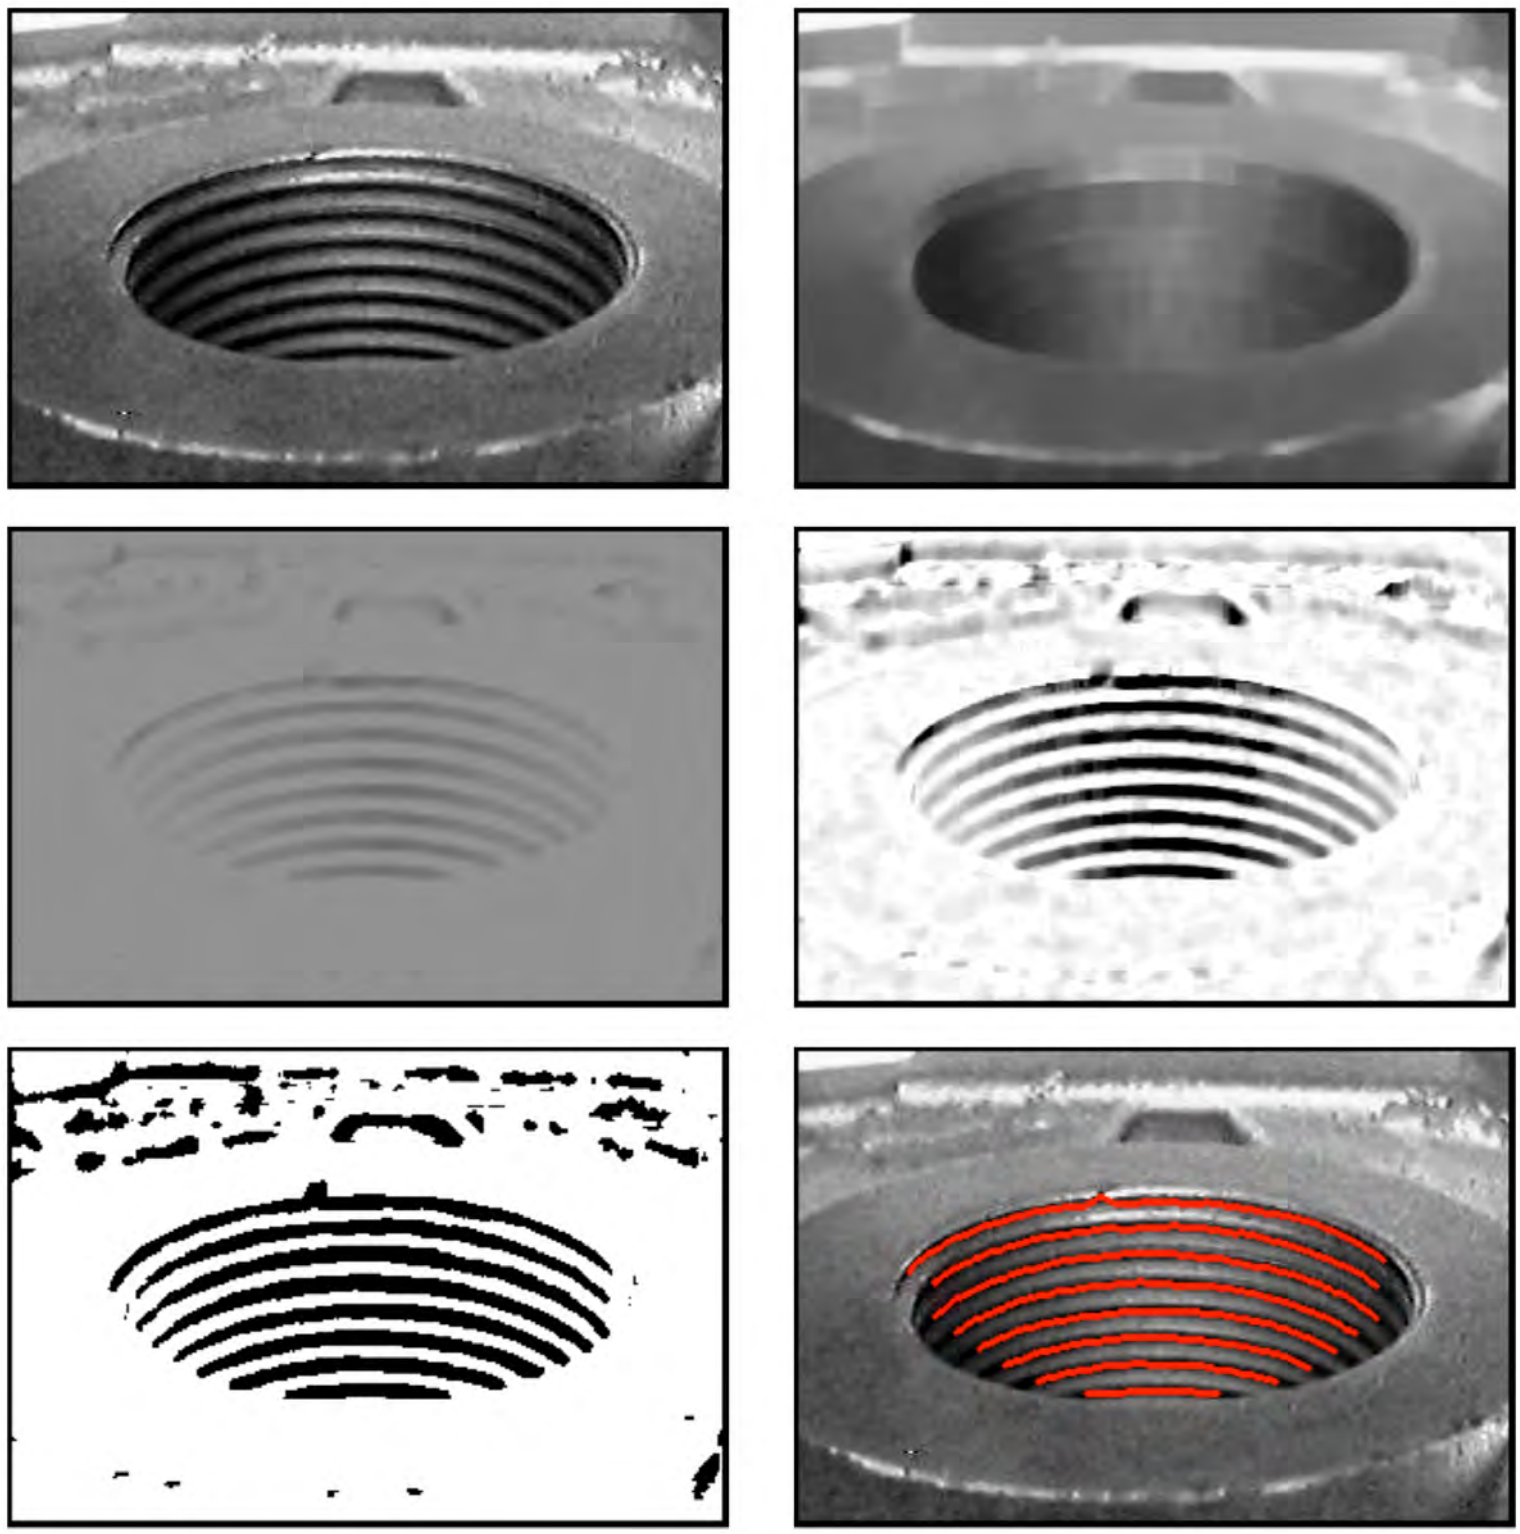

HARDNESS TESTING

To measure the hardness of a surface, we first form a tiny indentatio, by pressing a pointed tool into the surface, with a standard force. The size of the Indentation measures the hardness. An automated hardness testing procedure must tolerate some surface texture.

Vickers: Tool, diamond; URL: *https://en.wikipedia.org/wiki/Vickers_hardness_test*

Brinell: Tool, spherical URL: *https://en.wikipedia.org/wiki/Brinell_scal*

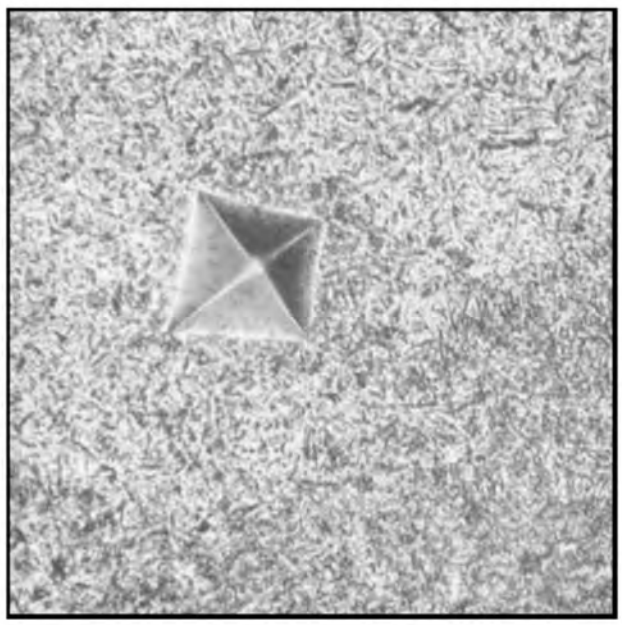

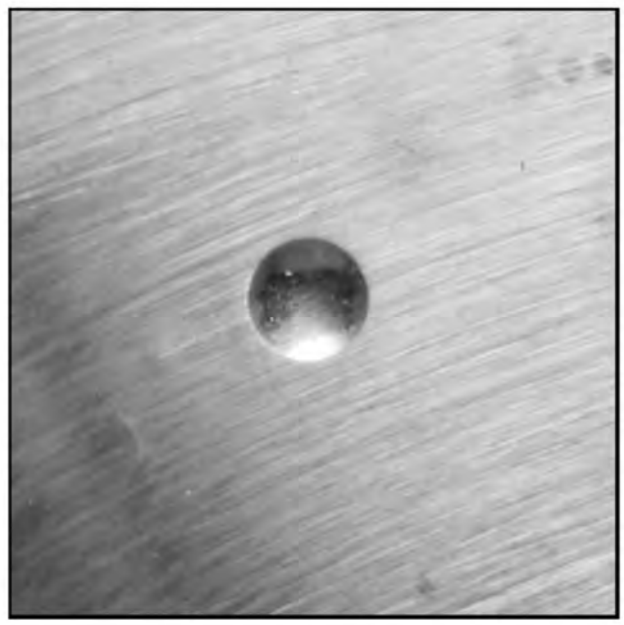

(Top-left) Vickers hardness test, Diamond indenter.

(Top-right) Brinell hardness test. Spherical indenter.

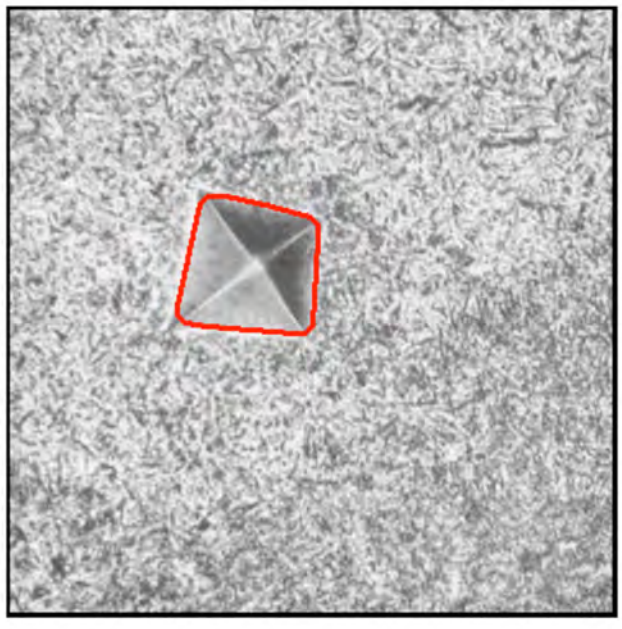

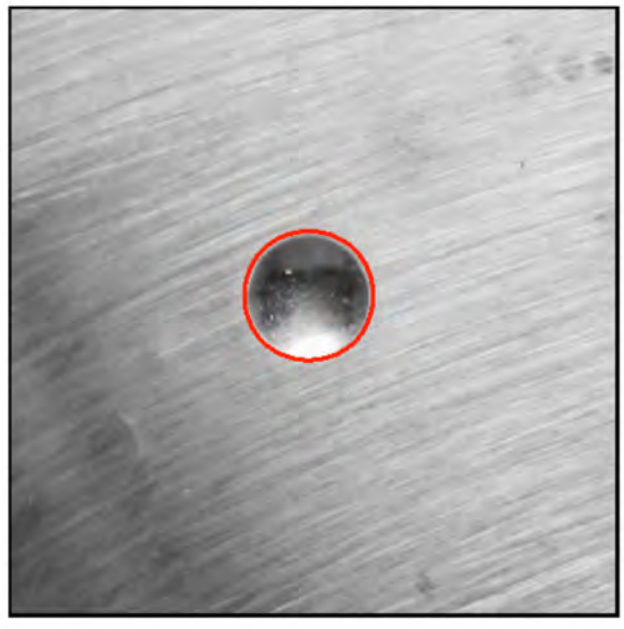

(Bottom-left) Contour around the edge of the Vickers indentation.

(**Bottom-right)** Contour around the edge of the Brinell indentation.

AEROSOL SPRAY CONE

Typical spray-cone measurements: angle; symmetry, density profiles. Inspection tasks, detect: jets, voids, large drops and pulses of spray material.

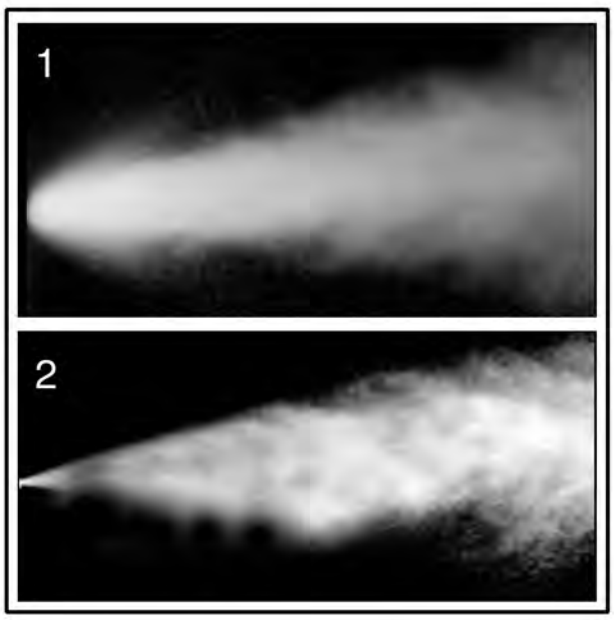

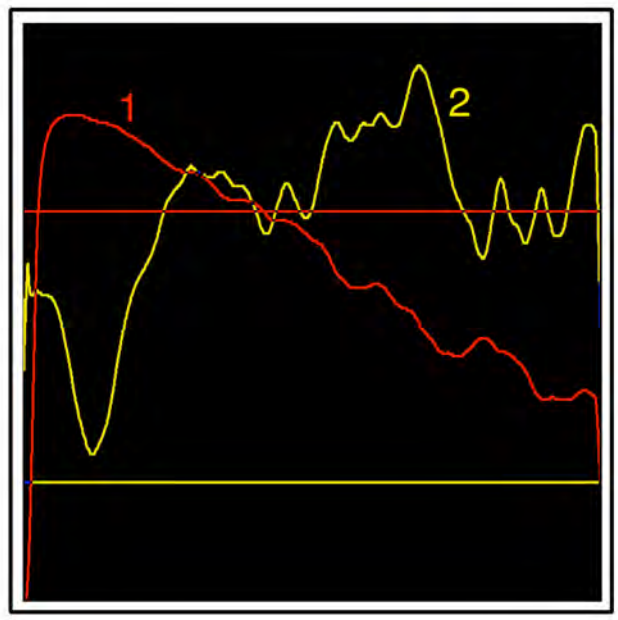

(**Top-left)** Two original images.

(**Top-right)** Row intensity profiles.

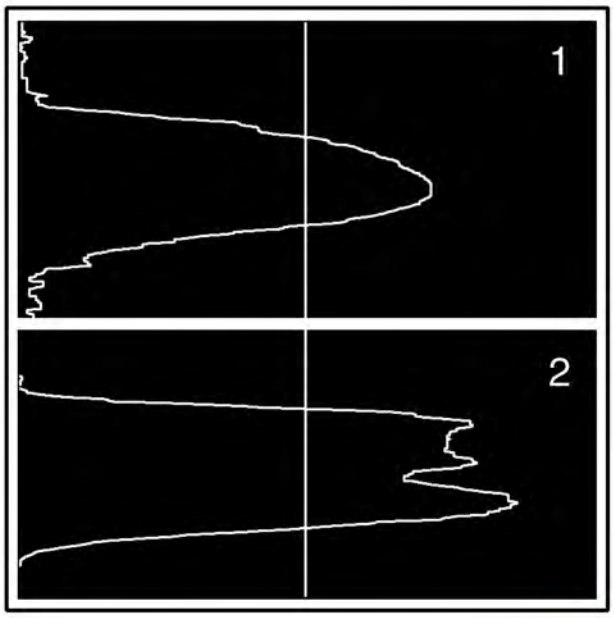

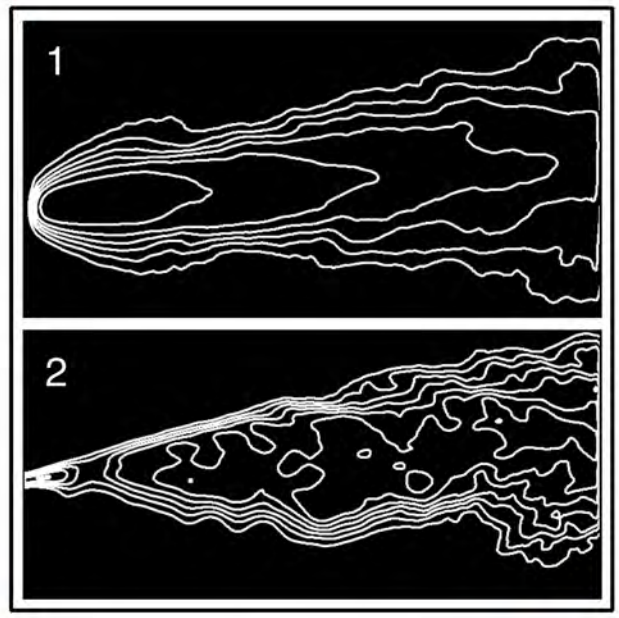

Centre-left) Column, intensity profiles.

(Centre-right) Intensity contours.

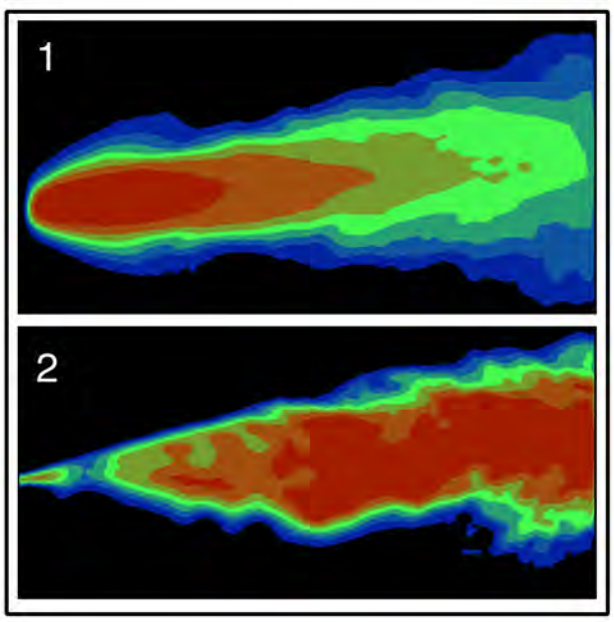

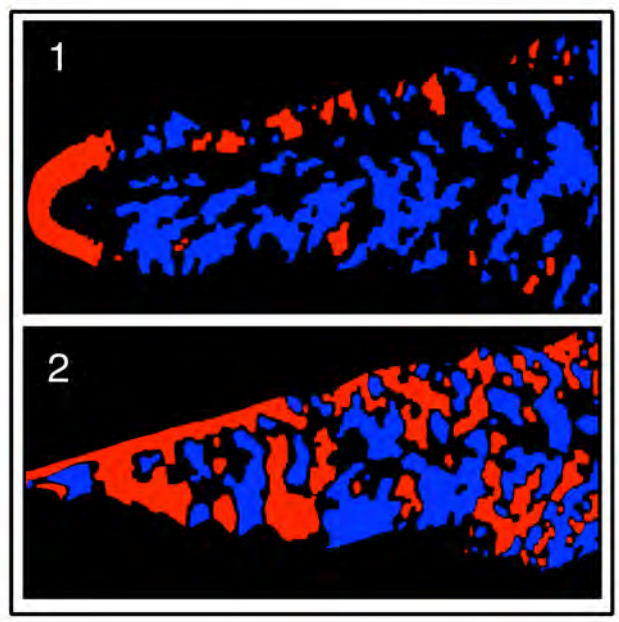

(Bottom-left) Pseudo-colour representation of the intensity

(Bottom-right) Colours indicate whether the intensity is increasing, or decreasing, as we travel from left to right.

CRACKS IN METAL (magnetic-particle visualisation)

There are several ways to enhance the visibility of cracks. See for example:

Magnetic Particle: URL *https://en.wikipedia.org/wiki/Magnetic_particle_inspection*

Dye Penetrant: URL *https://en.wikipedia.org/wiki/Dye_penetrant_inspection*

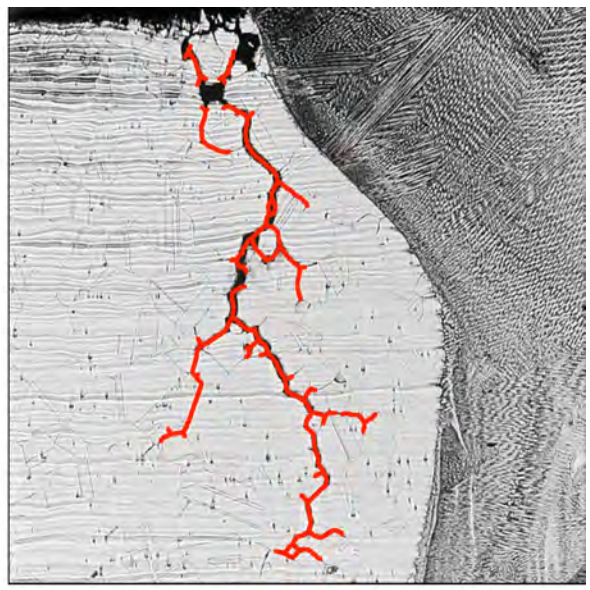

(Top-left) Original image. (No processing was used.

(Top-right) Crack detected by grey-scale morphology, plus noise removal.

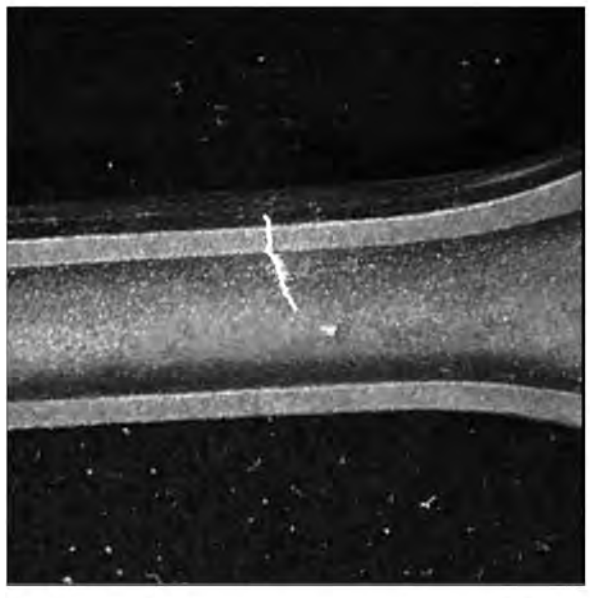

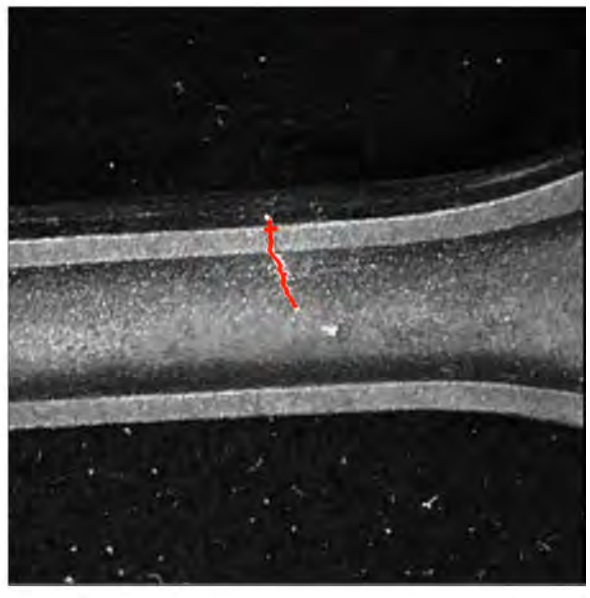

(Centre-left) Original image. (Magnetic-particle indicator)

(Centre-right) Crack detected by grey-scale morphology, plus noise removal.

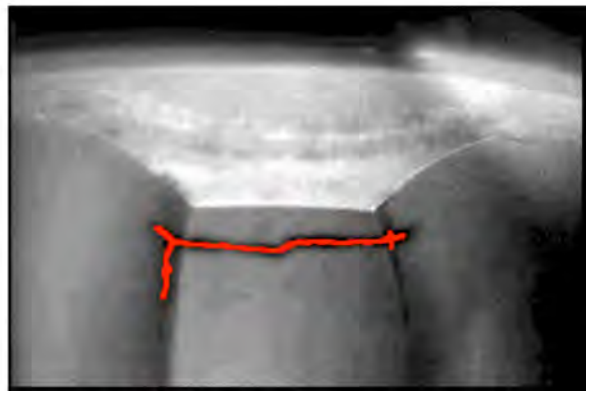

(Bottom-left) Original image. [Dye-penetrant indicator]

(Bottom-right) Crack detected by grey-scale morphology, plus noise removal.

ELECTRONICS (pattern location and inspection)

New-comers to Machine Vision envisage, expect direct image comparison (grey-scale template-matching) to be useful as a means of pattern matching. In practice, it is not used often because it is very susceptible to quantisation noise, camera noise and minor deviations compared to the reference pattern. In the application, shown here high-precision printing on a very stable substrate allows it to be used successfully. First, it is necessary to put the template in place. The two examples shown here use correlation to achieve an optimal fit.

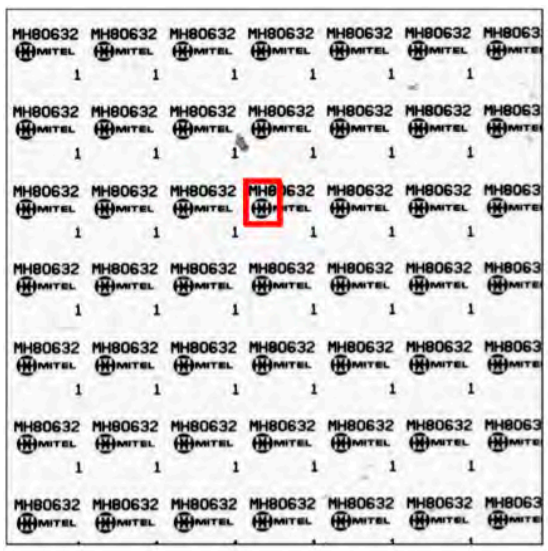

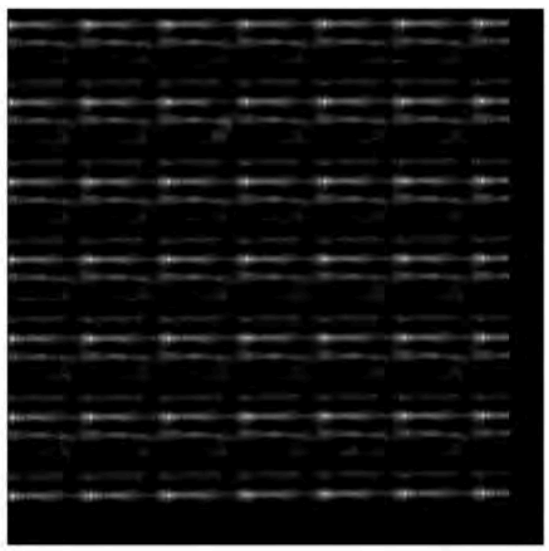

(Top-left) Original image,high-precision printing on a ceramic wafer. The sub-array within the red rectangle defines the correlation template for [TR].

(Top-right) Cross-correlation map between the template and the rest of the image in [TL].

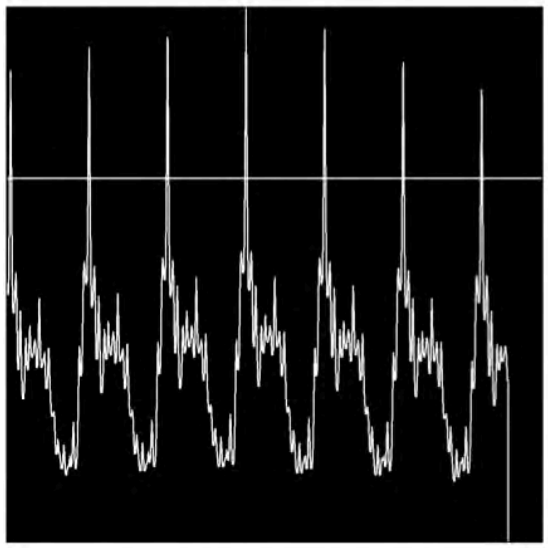

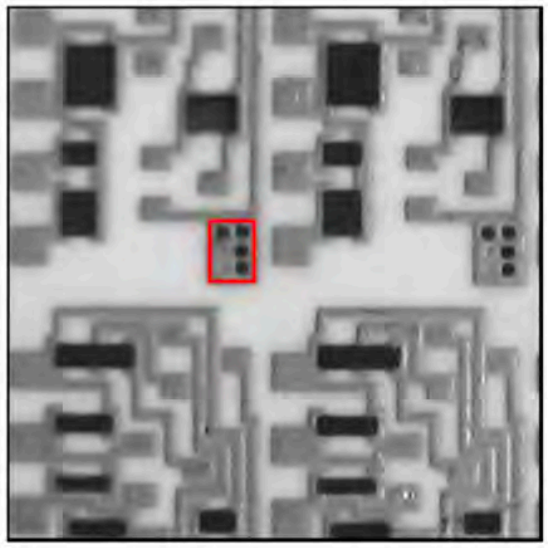

(Centre-left) Intensity profile through one row of brights spots in [TR]. The sharp, well-defined peaks show that the template fits similar patterns very precisely.

(Centre-right) Another original image, the component side of the same wafer as in [TL].

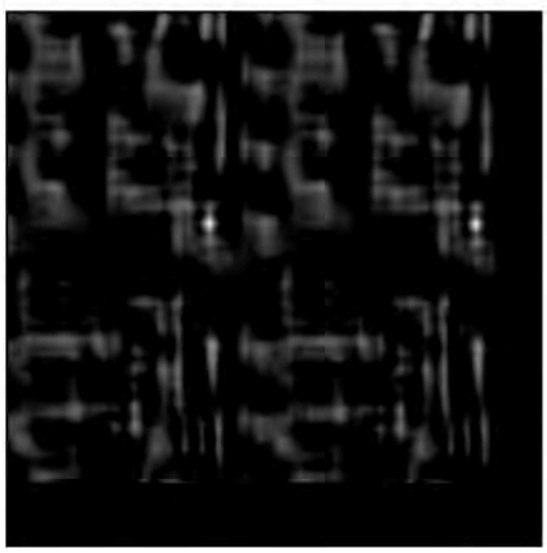

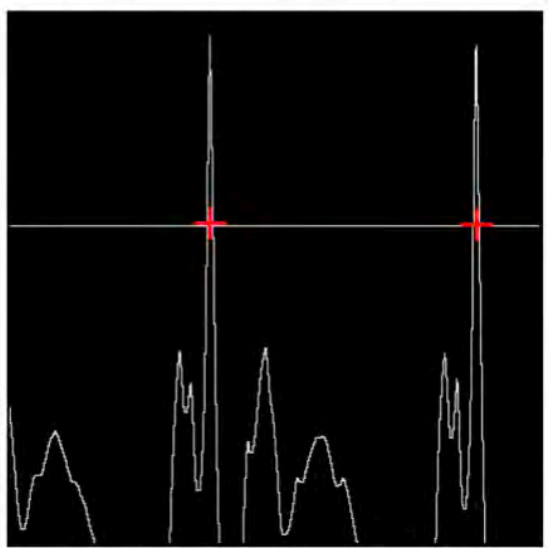

(Bottom-left) Cross-correlation map between the template and the rest of the image in [CR].

(Bottom-right) Intensity profile through the two brights spots in [BL]. Again, the sharp, well-defined peaks show that the template fits the similar pattern very precisely.

LEATHER

(Top) Animal hide. The quality and thickness of the leather varies across the hide:

- Belly, thin, soft pliable leather, suitable for gloves; shoe uppers, leather clothing.
- Back, thick tough leather, suitable for soles of shoes and boots, protective ware.

There is are several potential applications for a vision system

- Inspecting the hide for cuts, tears, thin spots, etc.
- Deciding where to place a "pastry cutter on the hide to produce componentS for gloves, shoes, etc
- To sort /check loose shoe components

(Others) Leather shapes for shoe components. At any given moment, components for several different styles and size of shoes will be present in the factory production area. For most shapes its the mirror image is also present

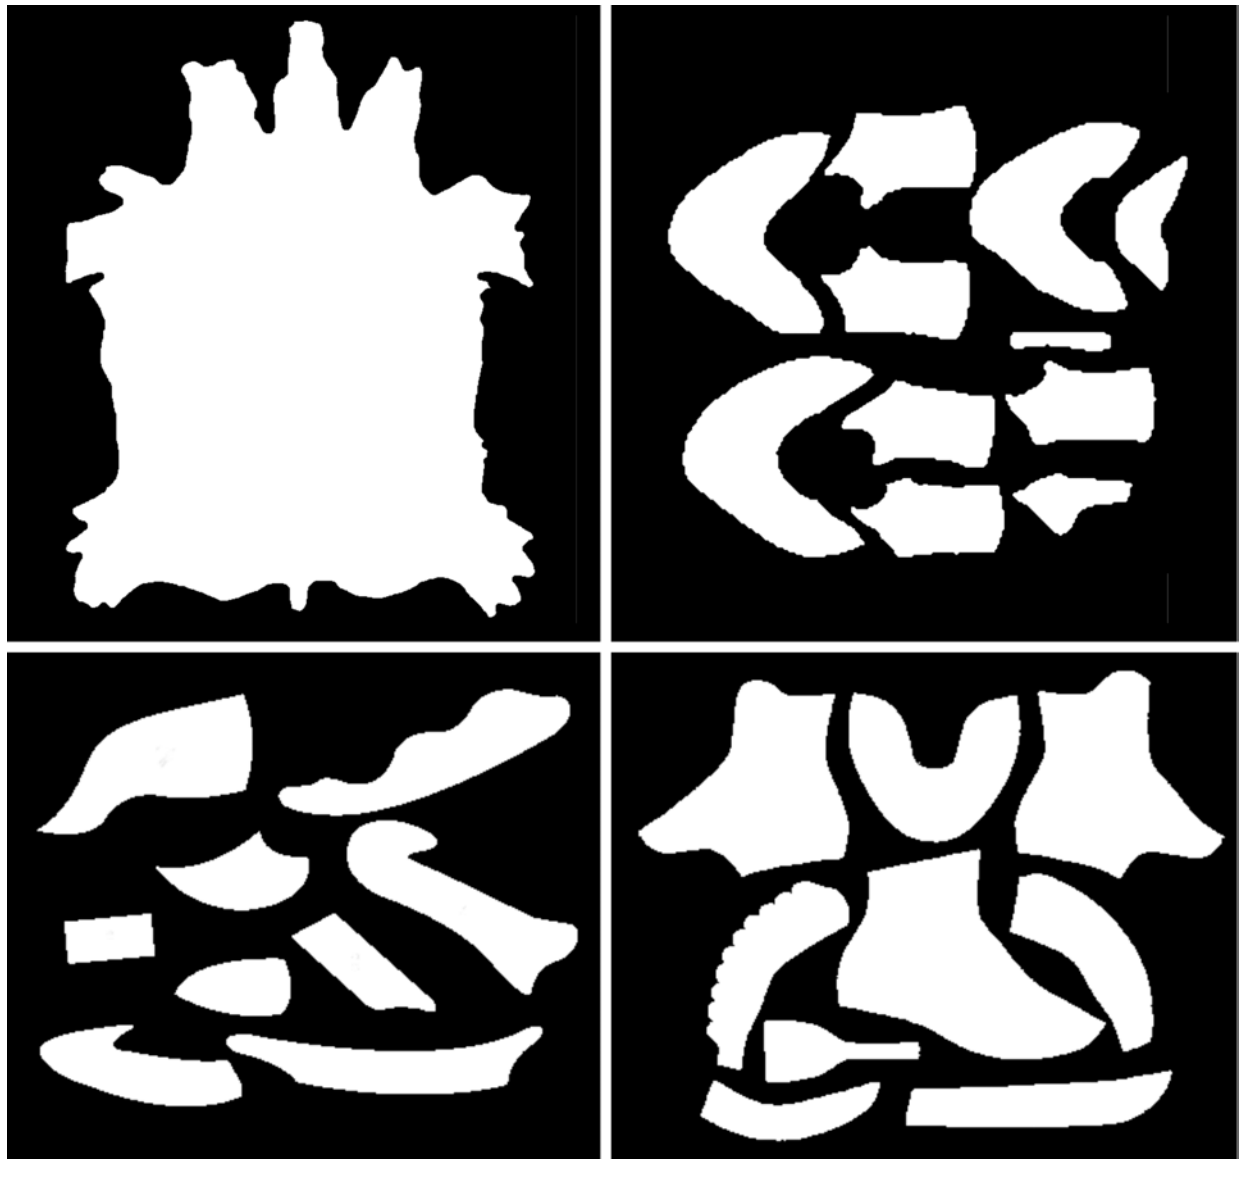

POPULATED PCB

Model-based inspection of a complex scene. While a printed circuit board presents a complex image, only a few small areas are likely to be of significant interest. Vertical lines (A and B) and horizontal lines (C and D) are identified first. These provide “anchors” that enable the inspection program to define specific regions of interest (ROIs) and then apply the appropriate image processing function to each one. The distance AB determines the horizontal scale. Lines A and B also allow us to fix the horizontal position. Similarly, lines C and D determine the vertical position and scale. These “anchors” allow the program to place the ROIs (four red rectangles.) Since each ROI covers only a small part of the board, very simple inspection checks might suffice, e.g. number of blobs of known colours (resistor colour codes), or checking that there is a large black object with three bright spots (left-most ROI, potentiometer)

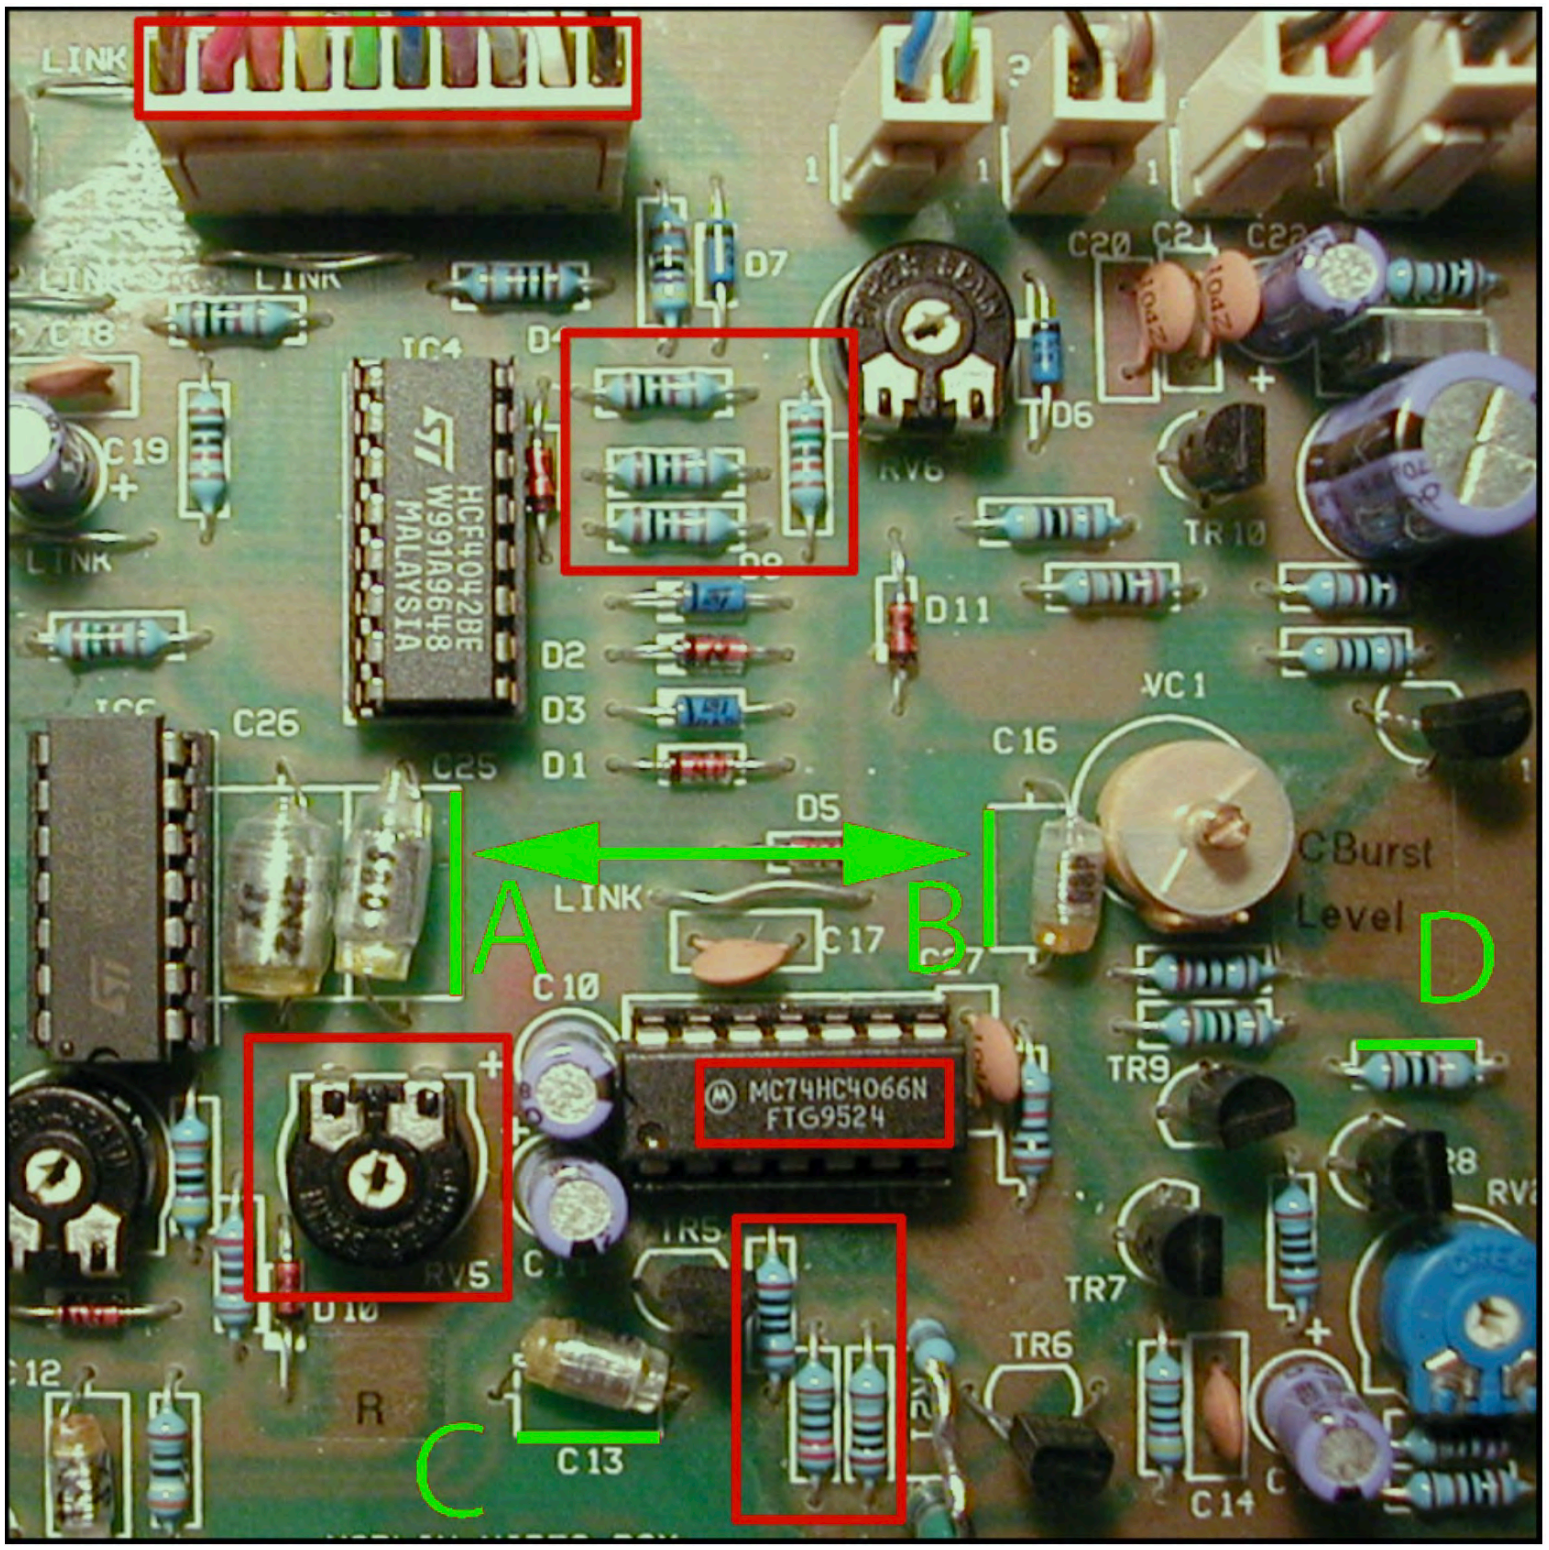

VIDEO MONITORING OF A CYCLIC MANUFACTURING PROCESS

(Top) Two frames from a video (Appendix: Movie ???)

(Centre-left) Histogram of the difference of Intensities of two frames, at the same phase in different cycles.

(Centre-right) As in [CL] but for different phases. (i.e. Histogram of difference of intensities in [TL] and [TR].)

(Bottom) Multi-image histogram. The intensity histogram of each frame in the video is represented by the intensities along one row in this image. It is easy to detect and analyse the the cyclic pattern, so that same-phase frames can be directly compared. Notice that the video is played twice.

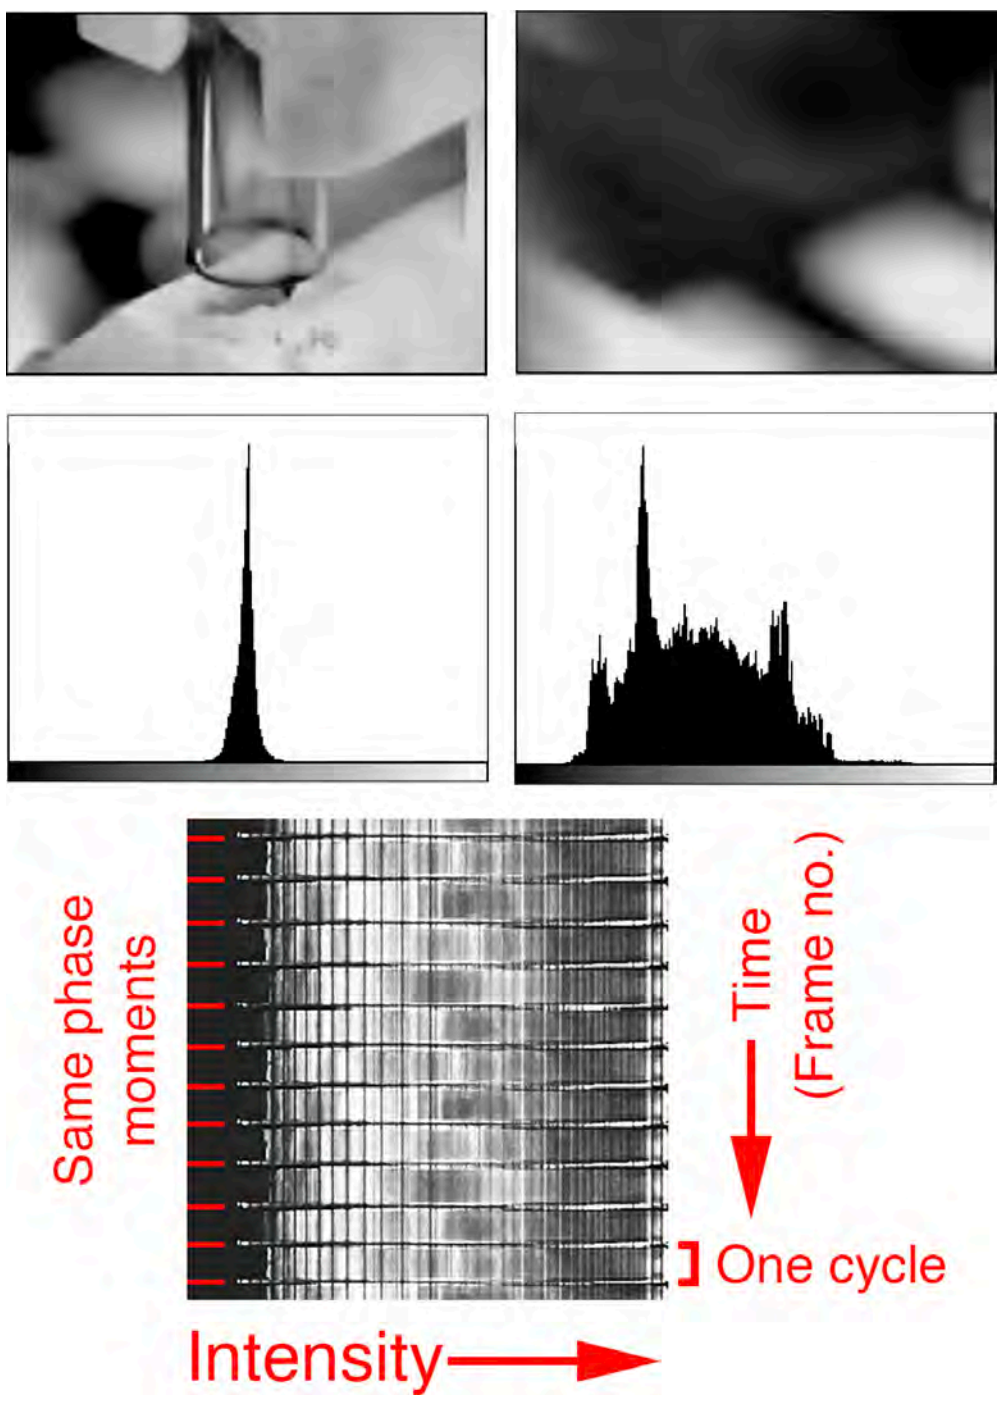

Section 7

SPORT & LEISURE

SAILING (monitoring sail trim on a racing yacht)

SailSpy is a real-time vision system, developed for automatically measuring sail shapes and masthead rotation on ocean-going racing yachts. It was used by the New Zealand team in two America's Cup challenges in 1988 and 1992. SailSpy uses four video cameras mounted at the top of the mast to provide views of the headsail and mainsail on either tack. The cameras are connected to a computer below deck using lightweight cables mounted inside the mast. Images received from the cameras are automatically analysed by computer which calculates sail-shape and mast-rotation parameters. The sail-shape parameters are calculated by recognising sail markers (dark ellipses) on the sails, and the mast-rotation parameters by recognising markers painted on the deck. Uncontrolled twisting of the mast can lead to sudden, catastrophic breaking

Using a vision system to analyse sail shape on high-performance racing yachts is common practice now.

References

https://www.researchgate.net/publication/252548588_SailSpy_a_vision_system_for_yacht_sail_shape_measurement

SailSpy: a vision system for yacht sail shape measurement, O J Olsson, et al Proc. SPIE 1823, *Machine Vision Applications, Architectures, and Systems Integration*, (1 November 1992);

SPORT (Soccer, Cricket, Tennis)

Many sports require action replay but this does not necessarily require a vision system to measure or interpret images. However, vision systems are employed for a wide range of sports, including cricket, tennis, Gaelic football, badminton, hurling, rugby union association football (soccer), volleyball, etc. Hawk-Eye is one such system. It is able to track the flight of a cricket ball and is even able to predict its future trajectory (after bouncing), to make critical decisions about *leg-before-wicket* appeals. (Predicting the bounce of a spinning cricket ball is beyond its capability - and many batsmen!) In tennis, the system is used to locate the exact point of bounce and thereby make critical decision about who wins a rally.

Monitoring and tracking the movements of players on a football field is useful for after-the-event analysis and can assist coaches and television commentators

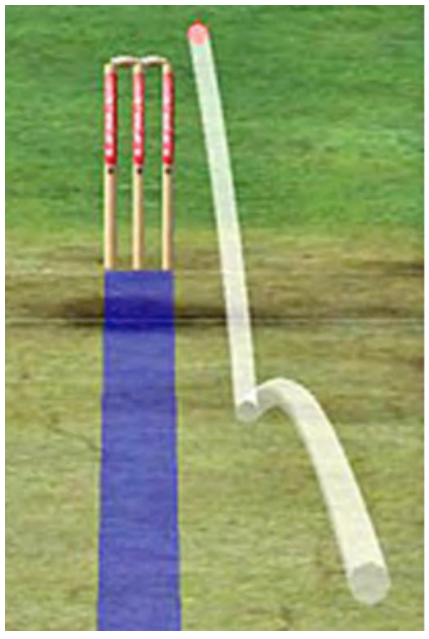

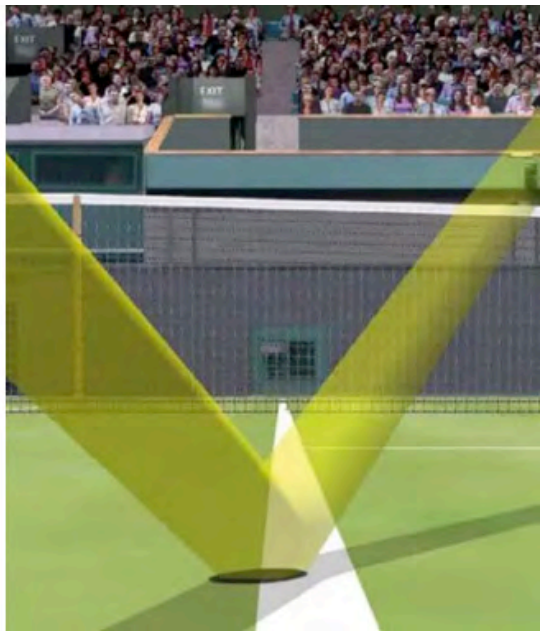

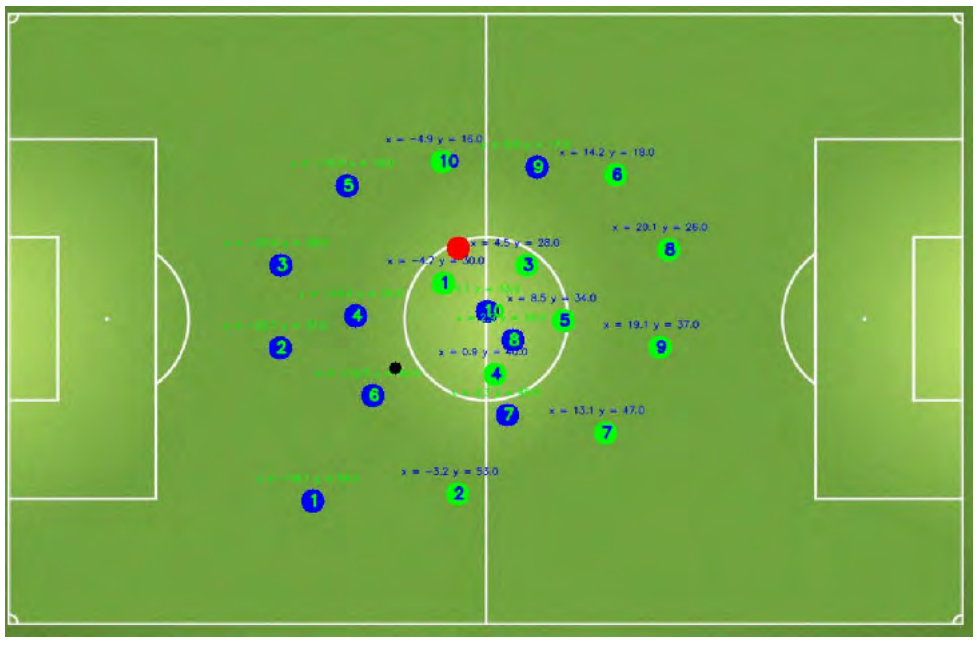

(Top-left) Cricket-ball trajectory analysis and flight prediction. This is important for judging *leg-before-wicket* appeals, which hinge on very narrow margins.

(Top-right) Tennis-ball bounce analysis. The outcome of a match may depend on whether the ball bounces in/out of court.

(Bottom) Player-position monitoring during a soccer game. The simulated plan view of the pitch indicates player positions and is calculated by a pitch-side vision system

SWIMMING POOL (detecting potential danger of drowning)

Multi-camera vision systems have been devised to detect possible dangerous (drowning) situations in swimming pools. [URL *http://www.poseidon-tech.com/us/system.html*]

A Poisedon system, installed at a pool in Bangor, Wales, UK saved the life of a 13-year old girl. [URL http://www.youtube.com/watch?v=qfKEBfPobmE]

(Top-left) Still from a video of a swimmer, taken using an underwater camera.

(Top-right) Two frames from this video sequence have been subtracted.

(Bottom-left) Binary image extracted from [TR].

(Bottom-right) Outline of the swimmer's body was extracted fro (BL] and superimposed



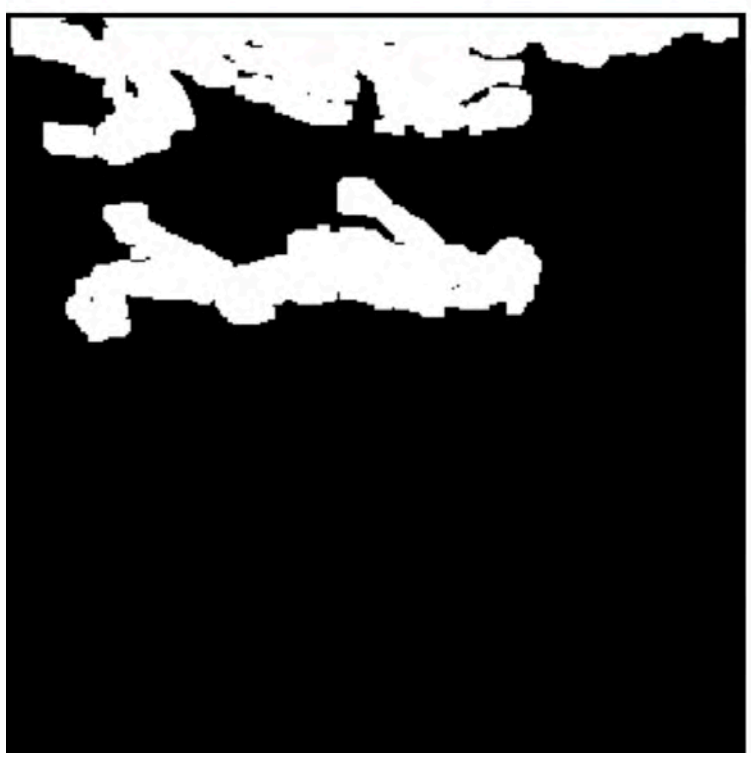


PRIVACY & SECURITY

FINGERPRINTS

(Top-left) Binary fingerprint, obtained by filtering, thresholding and reducing noise the original grey-scale image. This is a good-quality reference print. A scene-of-crime finger print will often be from a smaller part of the finger-tip and of inferior quality

(Top-right) "Focus" of the ridge-swirl pattern, identified by combining multiple binary morphology. The angled red line shows an estimate of the most prominent ridge direction and provides a crude estimate of the "orientation" of the fingerprint pattern (not the finger).

(Bottom-left) Features of interest, obtained by applying the skeleton operator to [TL]: The skeleton limb ends (ends of ridge arcs) are marked with red spots) . The skeleton joints (ridge bifurcations) are indicated by blue spots blue.

(Bottom-right) Two fingerprints can be matched by comparing their corresponding "star maps" like this. Inter-star distances are nearly constant, so there is no need to adjust the size of the prints during the matching process. The scene-of-crime star-map is rotated progressively (around the red cross, focus) to find the best fit to that derived from the reference print. The angled red line in [TR] suggests a good starting point for this search.

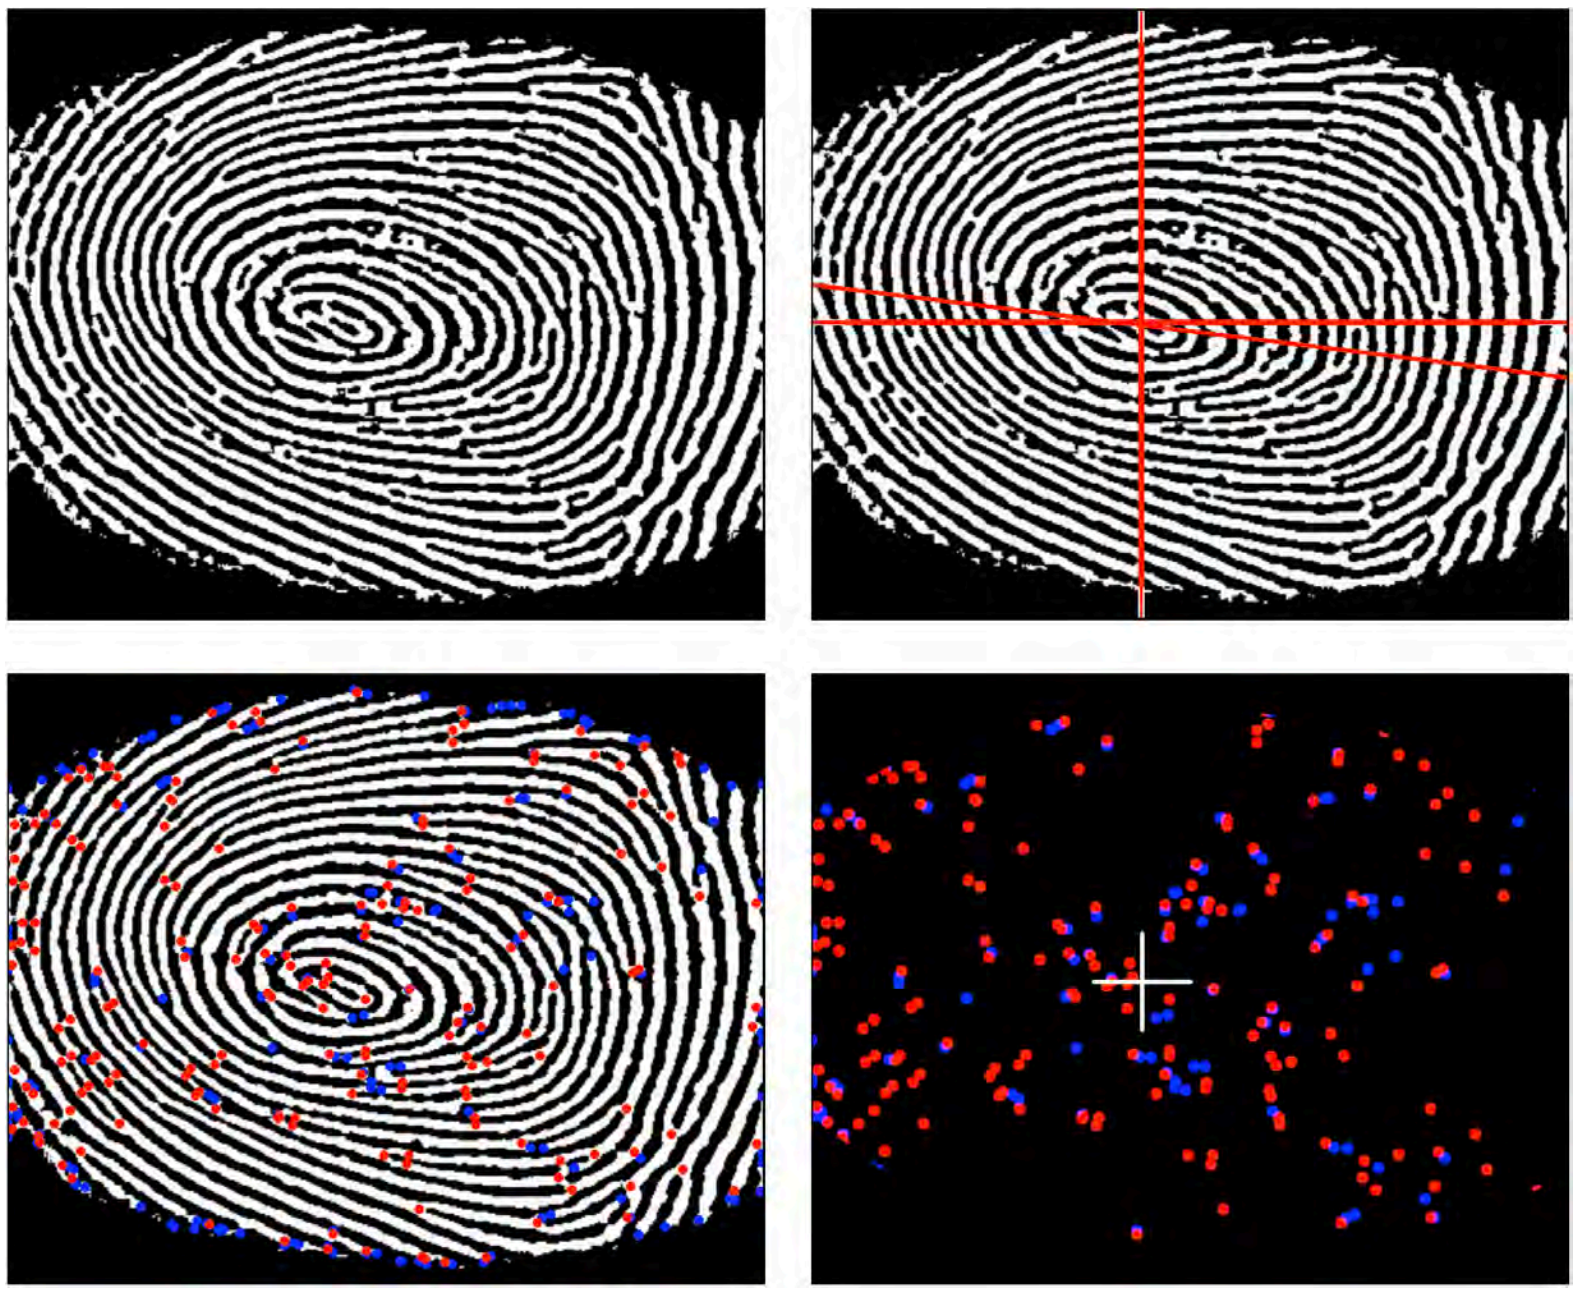

IRIS RECOGNITION

(Top-left) Original image.

(Top-right) Iris and pupil isolated.

(Bottom-left) Circles fitted to the edges of the pupil and iris.

(Bottom-right) After filtering and thresholding the grey-scale. The filtering relied on subtracting a blurred image from a very blurred image. Matching this pattern to a set of stored reference patterns, requires additional computation, akin to reading a (circular) bar code.

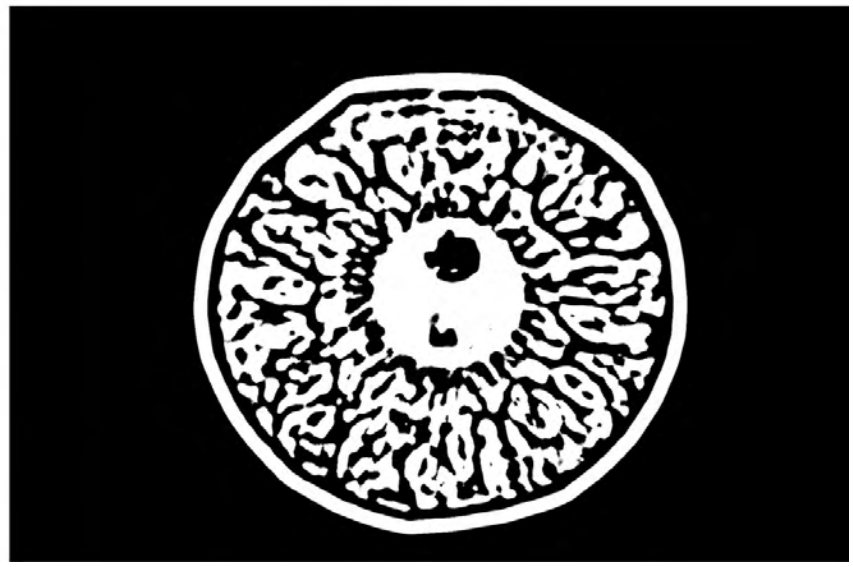

BANK NOTE (security detail on old-issue UK £5 note)

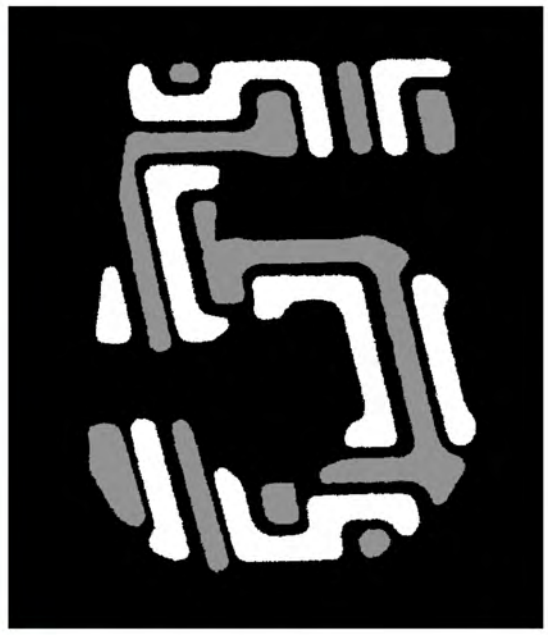

(Top-left) Fluorescent ink is made visible using ultraviolet light. This image was captured viewing light within the visible waveband.

(Top-right) Image segmented by intensity thresholding.

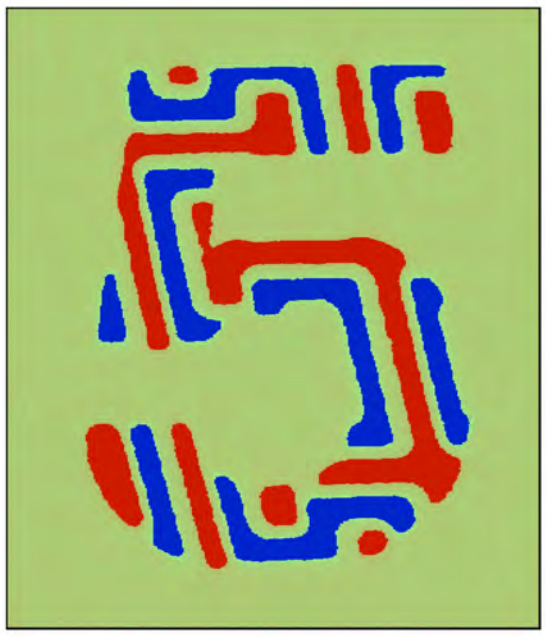

(Centre-left) Image segmented using colour separation.

(Centre-right) Visual image in white light.

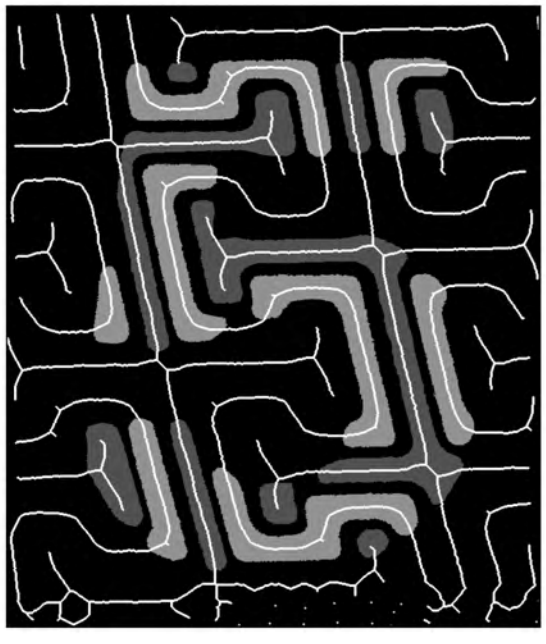

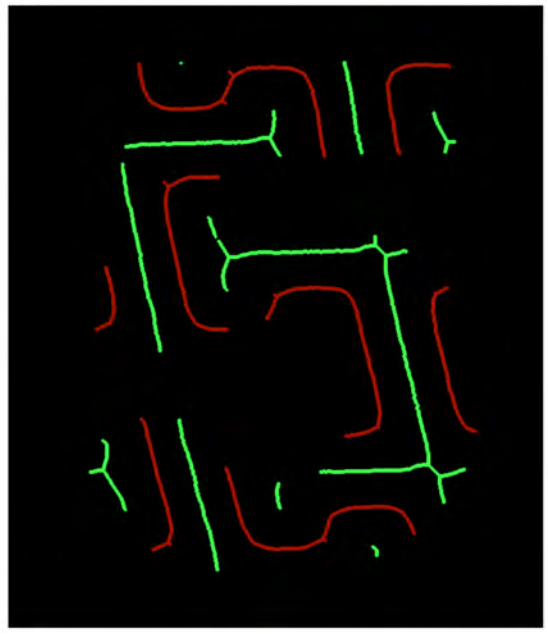

(Bottom-left) The meandering curve was obtained by processing [CR].

(Bottom-right) Relating the UV and visible images by ANDing the meandering curve in [BL] with the blobs derived from the UV image [TR].

HOLOGRAM (and other surfaces producing diffraction patterns)

(Left) Original images. White light was projected onto the same credit-card hologram, from six different directions. Only the angle of illumination was varied.

(Right) Colour features derived from [L].

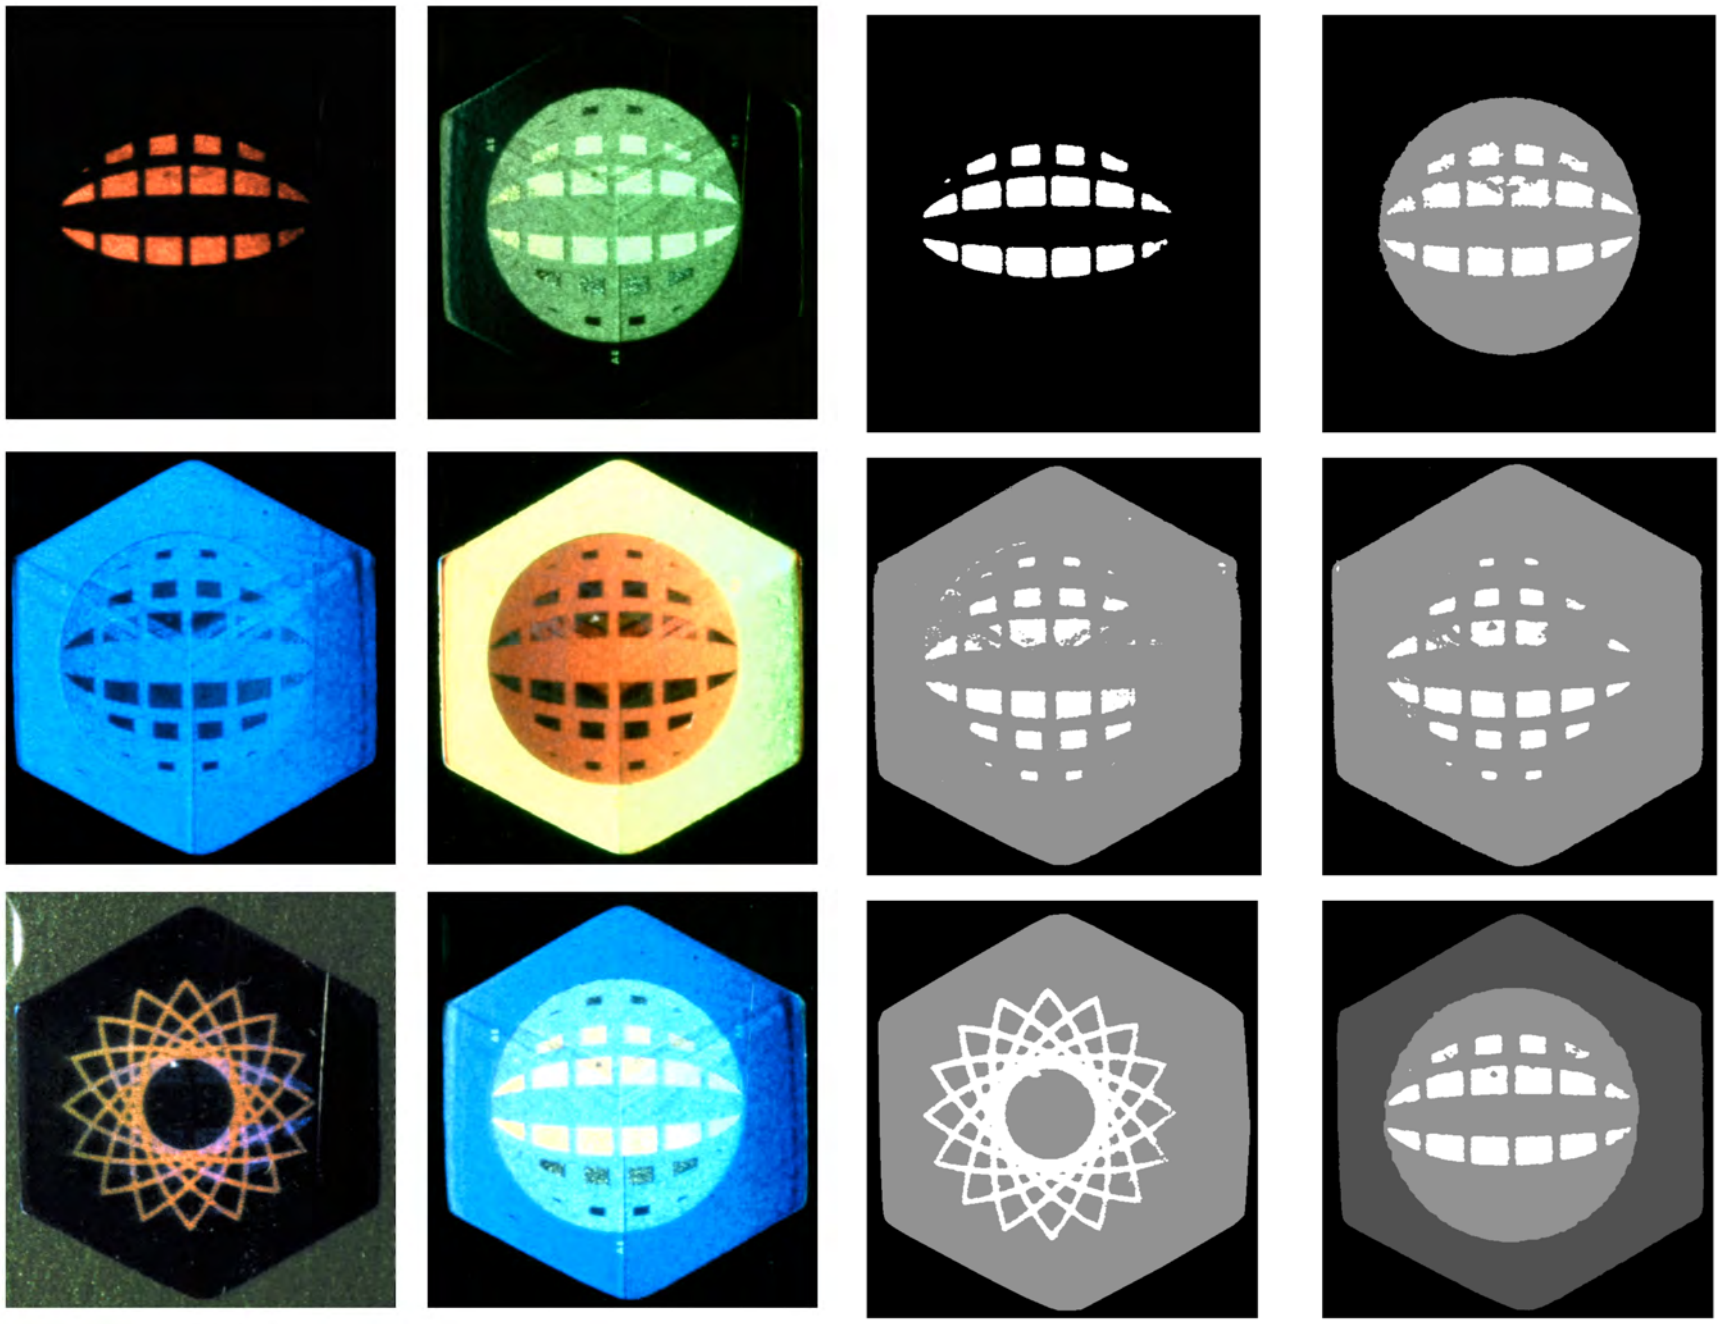

SECURE CODING OF A BINARY IMAGE (1-TIME PAD)

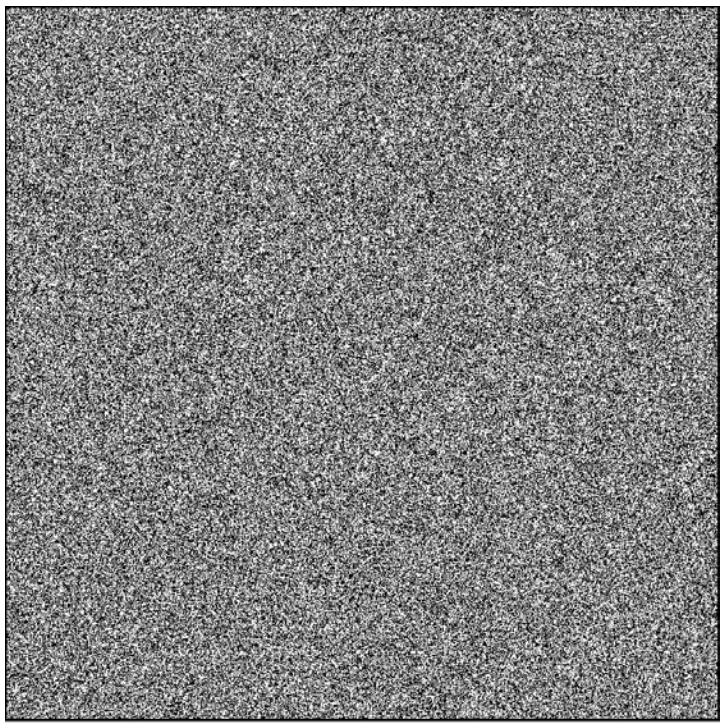

(Top) Pseudo-random binary image, generated in MatLab using the Mersenne Twister algorithm. Important points to ensure high security:

1. 50% of the pixels in this image must be white.
2. The random image must be used only once. (This is a “1-time pad” coding method.)
3. In a real security application this image would be held by both sender and receiver of the coded message.

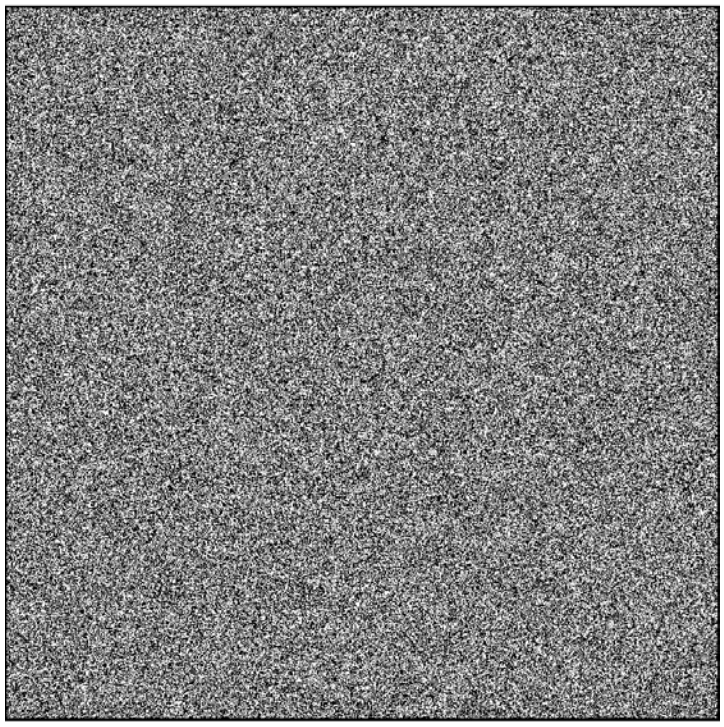

(Centre) Message after coding. This was generated by combining [T] with the binary image to be hidden (QR code) using exclusive OR].

Y

(Bottom) Result of combining [T] and [C] using exclusive OR. This is identical to the original QR code

WATERMARKING (EMBEDDING A BINARY IMAGE)

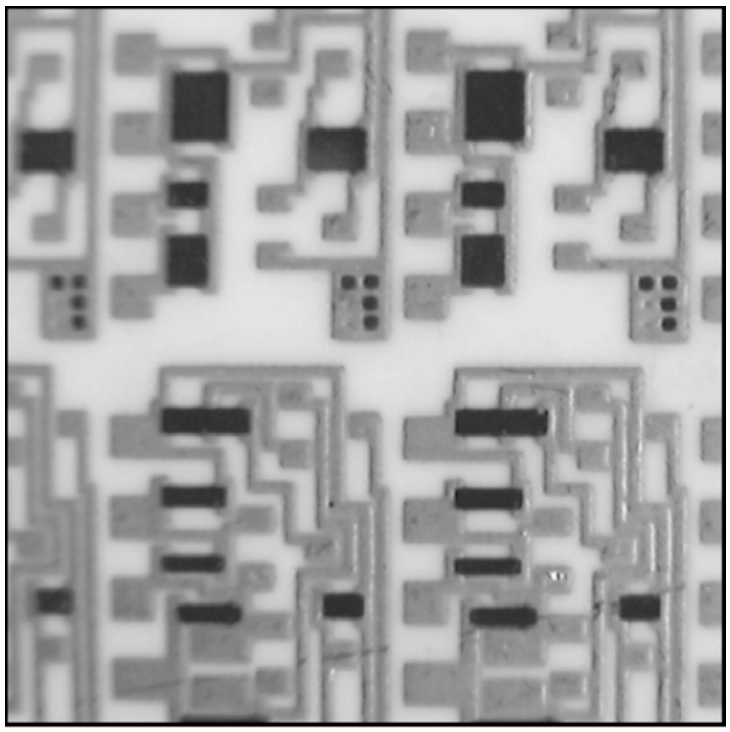

(Top-left) Original grey-scale image. This will be used as decoy, to obscure the presence of the message in [C].

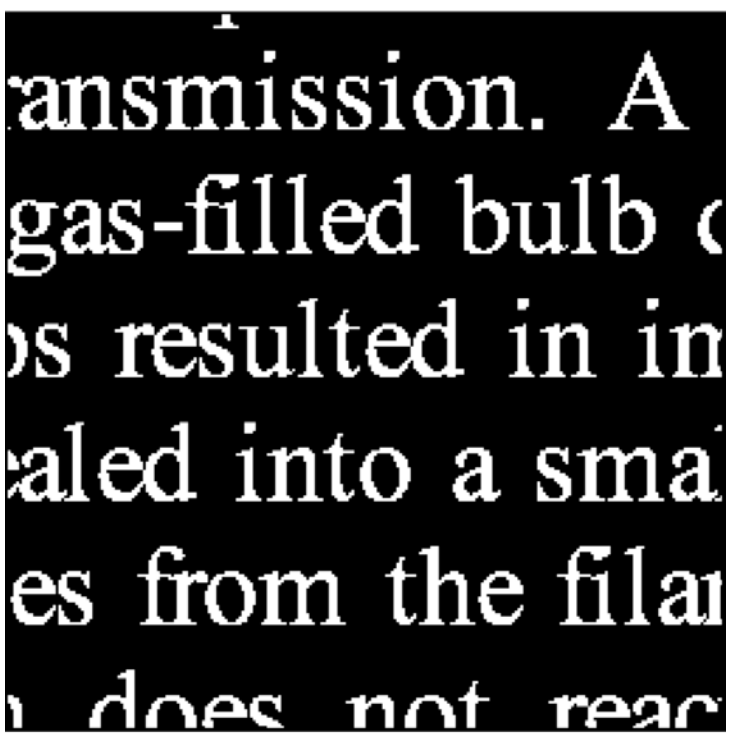

(Centre) A binary image, representing a "secret message".

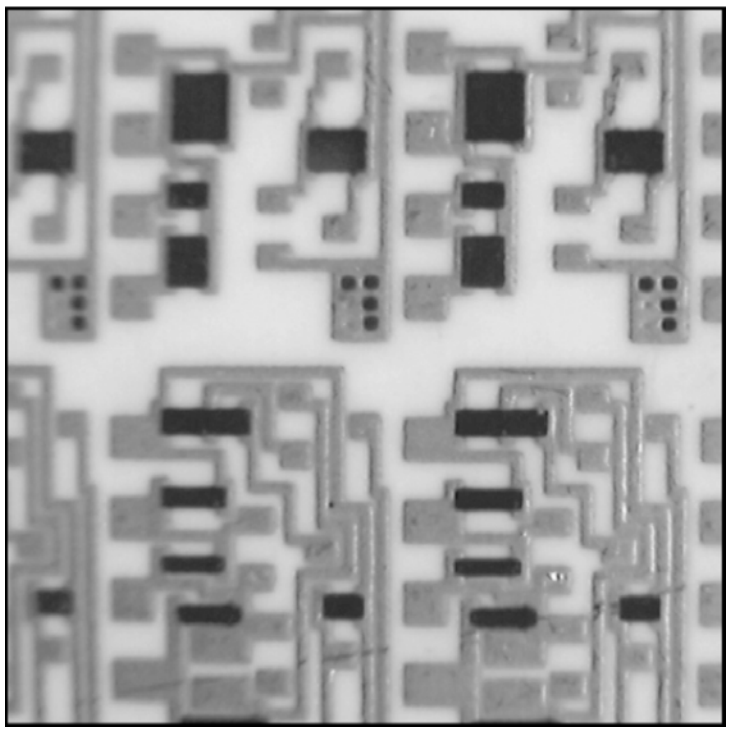

(Bottom) [C] embedded in [T] by inverting the least-significant bit at each location corresponding to a white pixel in [C]. Although this and [T] are visually indistinguishable), [C] can be reconstructed exactly, by simple image subtraction.

WATERMARKING A COLOUR IMAGE / HIDING 3 QR CODES

(Top-left) Original colour image.

(Top-right) Watermarks added. Notice that this and [TL] are visually indistinguishable

(Centre-left, Centre-right, Bottom-left) QR-codes representing the six verses of the hymn *"Amazing Grace"* by John Newton. These were superimposed on the R, G and B channels of [TL] respectively.

(Bottom-right) Bit reversal of [TR]. The three QR-codes can be extracted from this image by selecting the least-significant bit in each of the RGB channels in turn and then thresholding.

SENSING THE COLOUR STATE OF PHOTOCHROMIC INK (explanation)

The photochromic material (PCM) used in this application undergoes a reversible change of colour (clear to pink to clear) in response to electromagnetic radiation (UV and visible light). Both colour states are stable and the switching is fast. This PC differs from that used in photochromic spectacle lenses, which has one stable and one quasi-stable state

PCs have considerable potential for adding extra security features to bank notes, credit cards, identity cards, etc. Combined with Machine Vision and the coding methods outlined later,

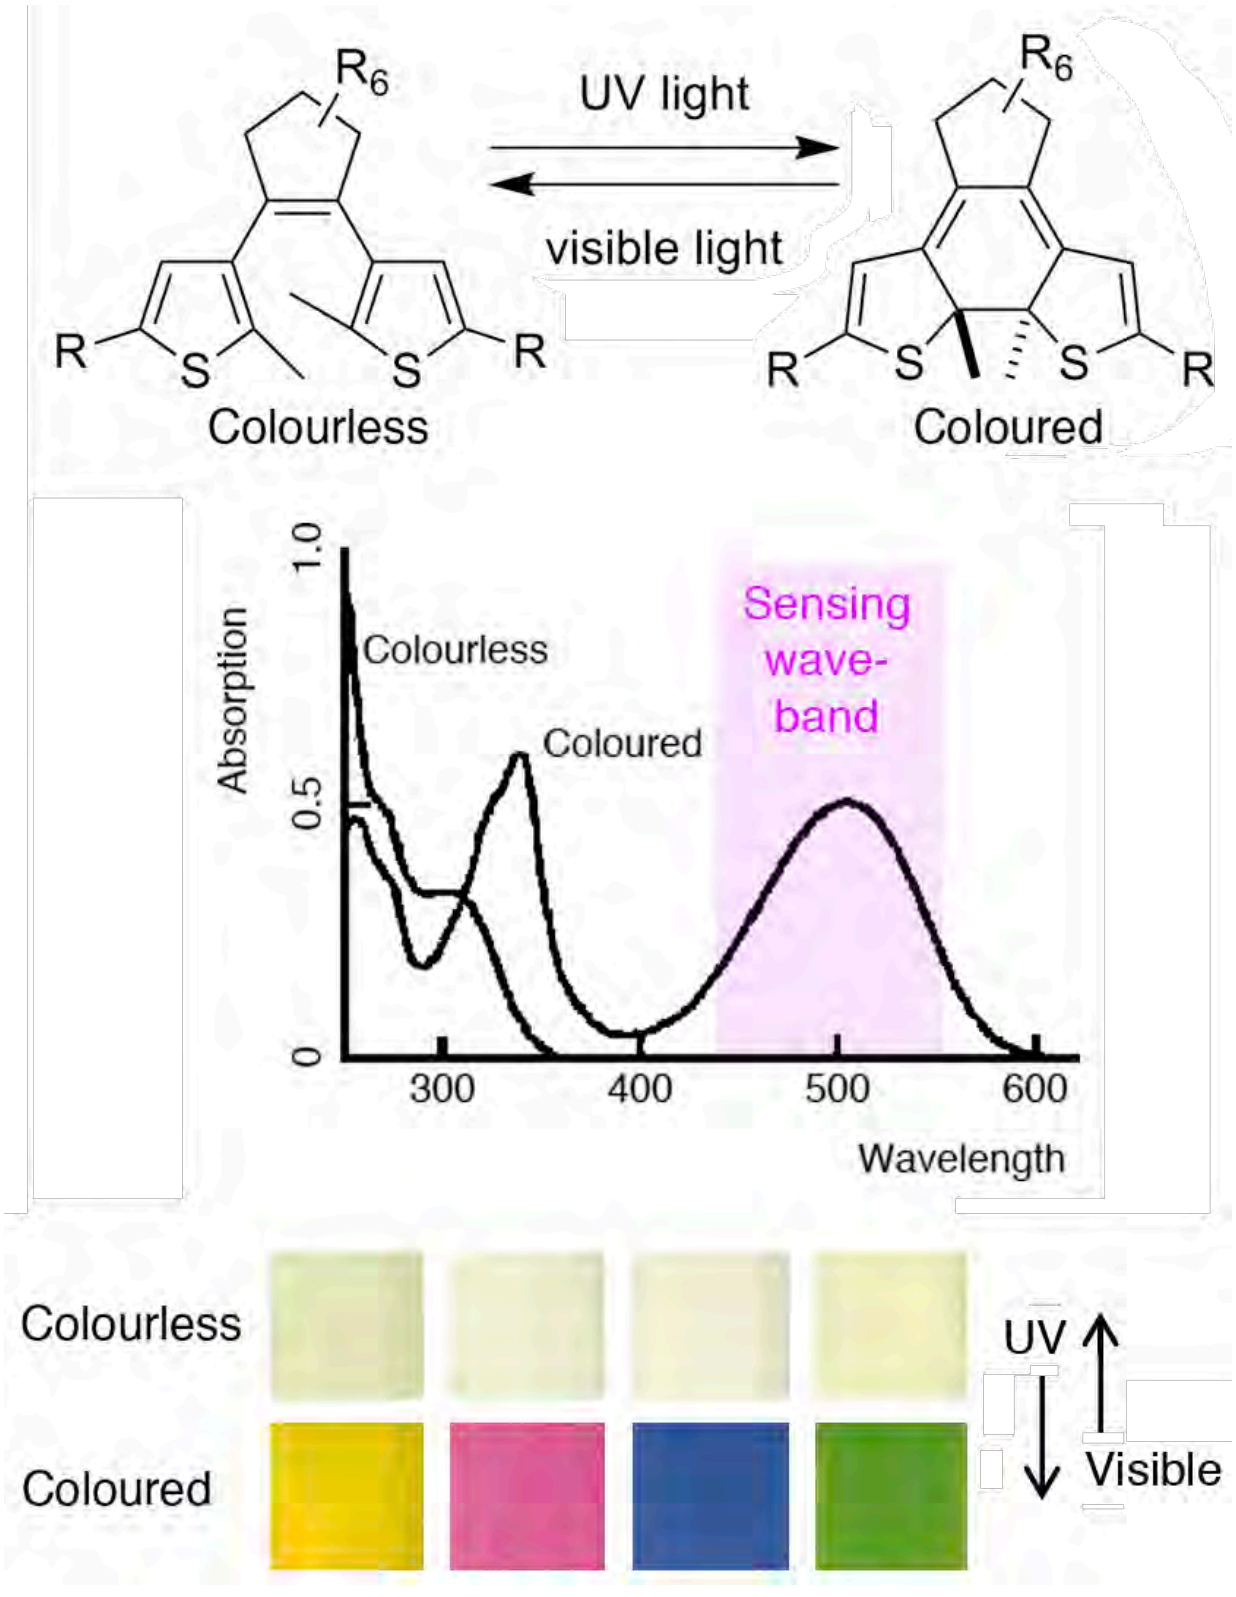

(Top) The shape of the PCM molecule can be switched to the "colour" state by illuminating it with UV. This can be reversed, using visible light.

(Centre) Reflectance spectra of PCM. Observing the energy within a narrow spectral band (mauve) allows the colour switch to be detected reliably.

(Bottom) Colours of four types of diarylethene ink.

[Source: Proc Jpn Acad Ser B Phys Biol Sci. 2010 May 11; 86(5): 472–483.]

SENSING THE COLOUR STATE OF PHOTOCHROMIC INK (results)

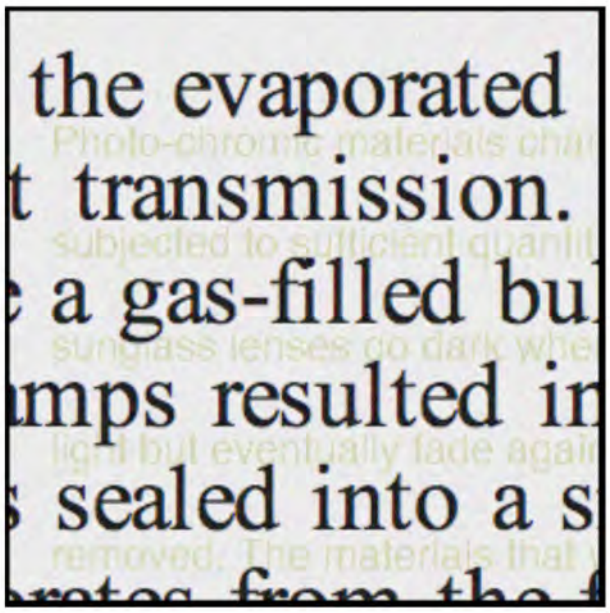

(Top-left) Photochromic ink in its colourless state. The PC and black inks are deliberately superimposed.

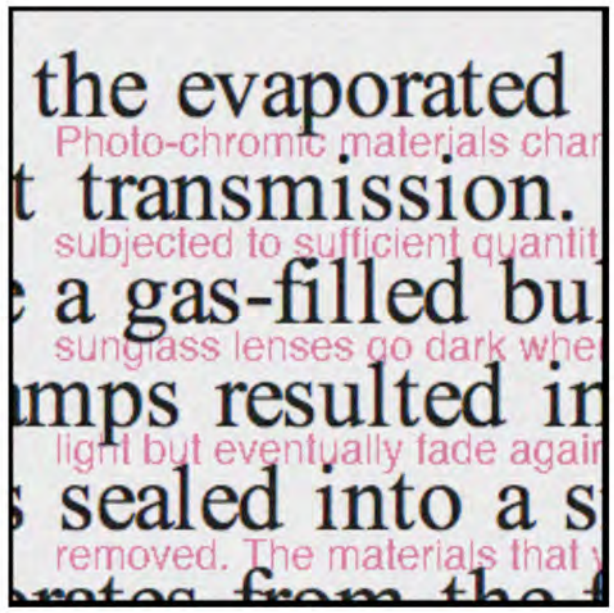

(Top-right) Photochromic ink in its coloured state.

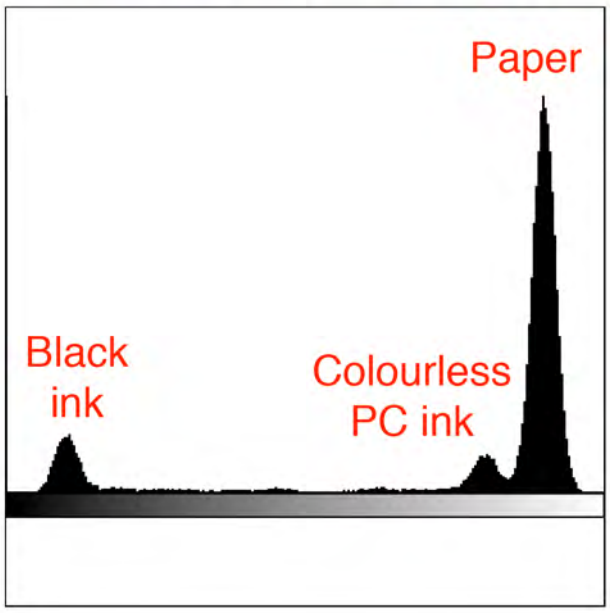

(Centre-left) Intensity histogram of [TL].

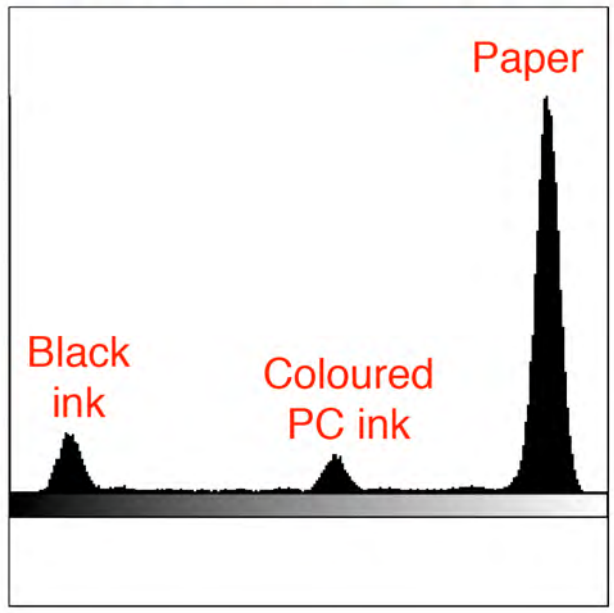

(Centre-right) Intensity histogram of [TR].

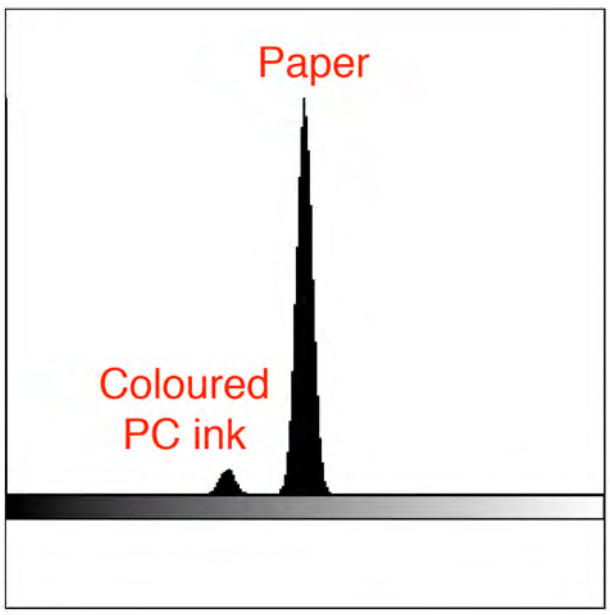

(Bottom-left) Histogram of the difference in the intensities of [TL] and [TR]. Notice that there is a well defined valley, providing a clear guide for thresholding.

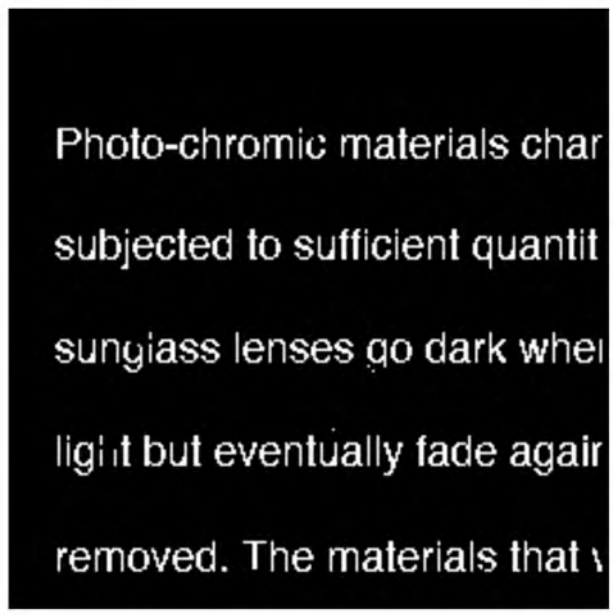

(Bottom-right) Threshold of the difference image (i.e.subtracting the intensities of [TL] and [TR]).

Section 9

RECOGNISING PRINTED CHARACTERS, LOGOS & MOTIF PATTERNS

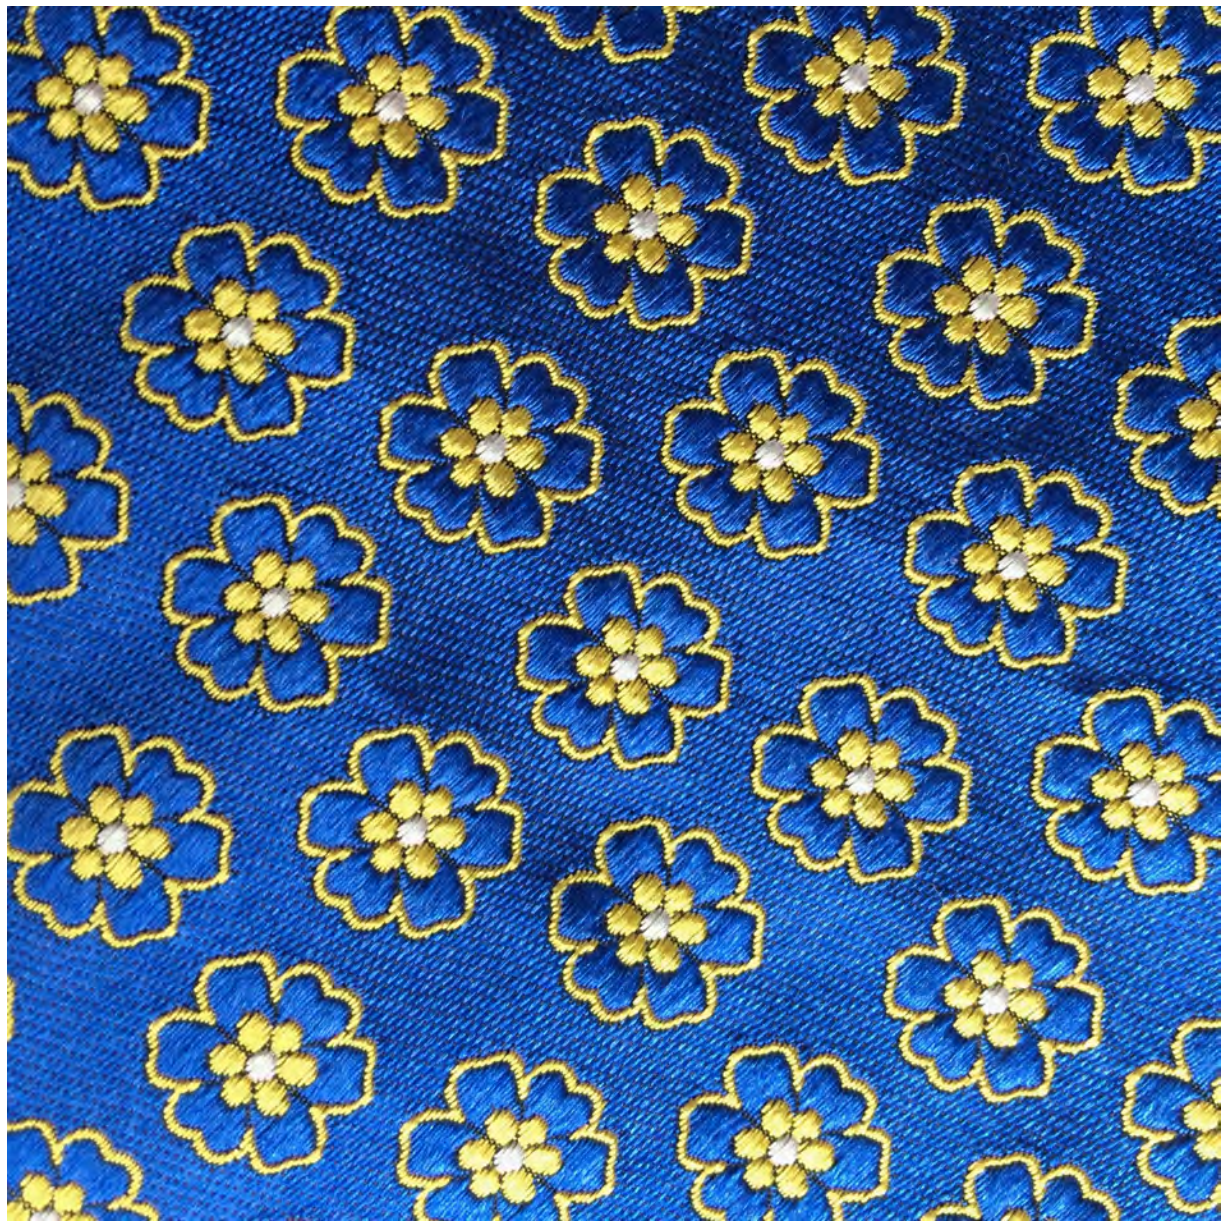

OPTICAL CHARACTER RECOGNITION (OCR)

Illustrating basic principles of character recognition on well-formed printed alpha-numeric characters, the value symbols of playing cards: 2,3, ...,10, J, Q, K, A The processing was designed to ignore serifs, since they are very small and may produce erratic results

(1) Original grey-scale images.

(2) Binary images, derived from column 1..

(3) Graphs showing numbers of white pixels along each row in column 2.

(4) Graphs showing numbers of horizontal black-white transitions along each row

(5) Spots indicate the skeleton limb ends.

(6) Spots indicate the skeleton joints.

(7) Major bays.

(8) Lakes (Euler number is faster)

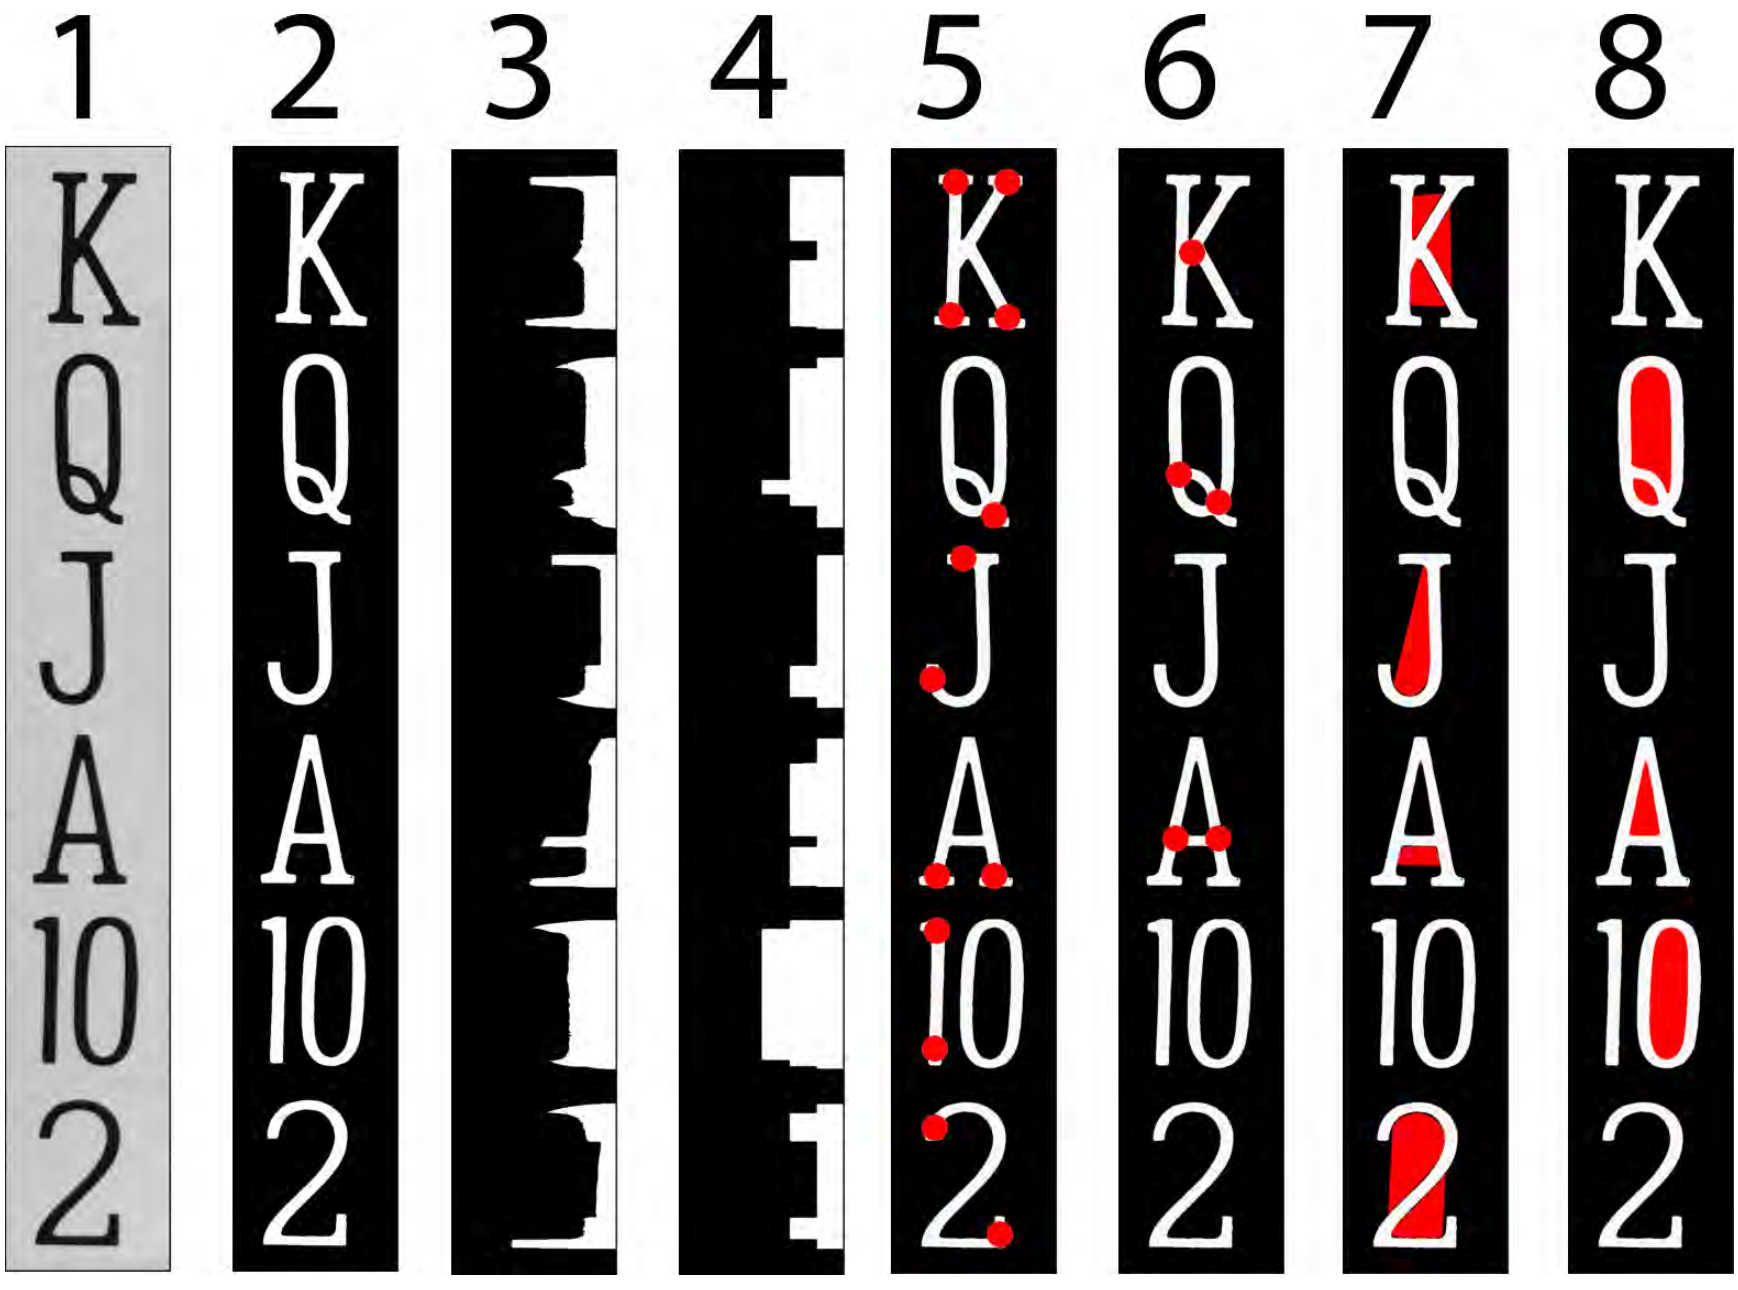

OPTICAL PATTERN RECOGNITION (playing cards)

A commercial system has been developed for checking the integrity of new packs of playing cards. This application does not represent a large area of application for Machine Vision. However, playing cards do possess many features in common with printed cartons. As a result, they provide a useful model for demonstrating how Machine Vision can contribute in this very much larger market.

In this illustration, morphology is used to recognise **Q** and small-♥ and the same patterns inverted.

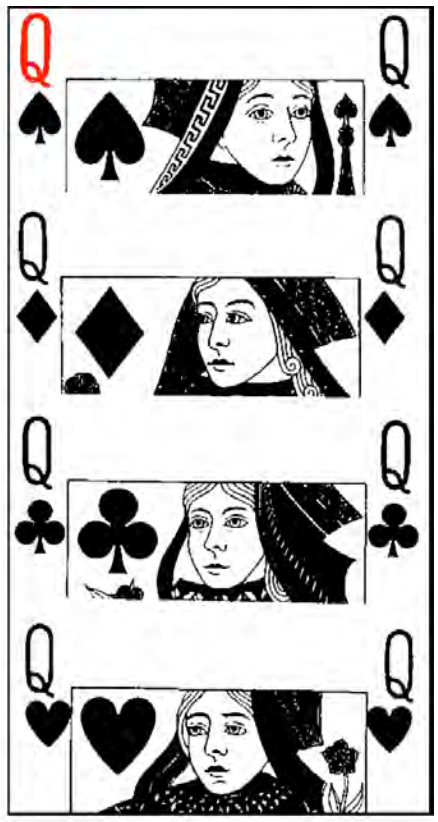

(Top-left) Original image: collage of four cards.

(Top-right) B-image. This gives the greatest contrast (out of RGB) for the red features in [TL]. The red **Q** here was used to define the structuring element for morphology.

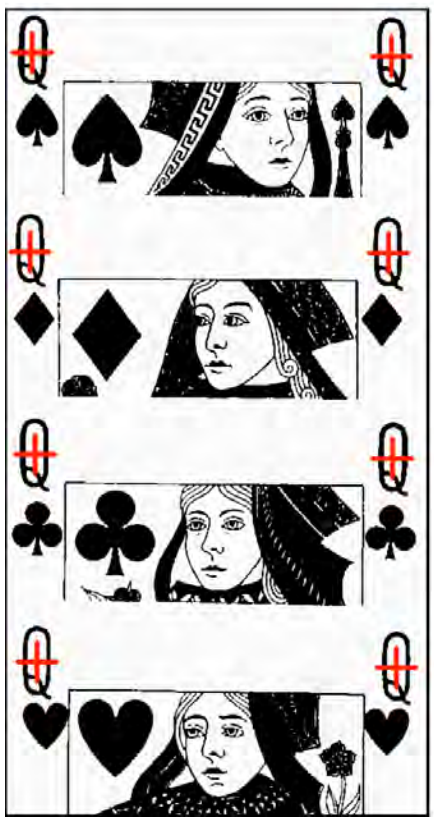

(Bottom-left) Qs detected

(Bottom-right) All **Q**s and small ♥s detected. This is a composite of four different morphology operations, detecting **Q**, ♥ and these patterns inverted.

OPTICAL CHARACTER RECOGNITION (OCR)

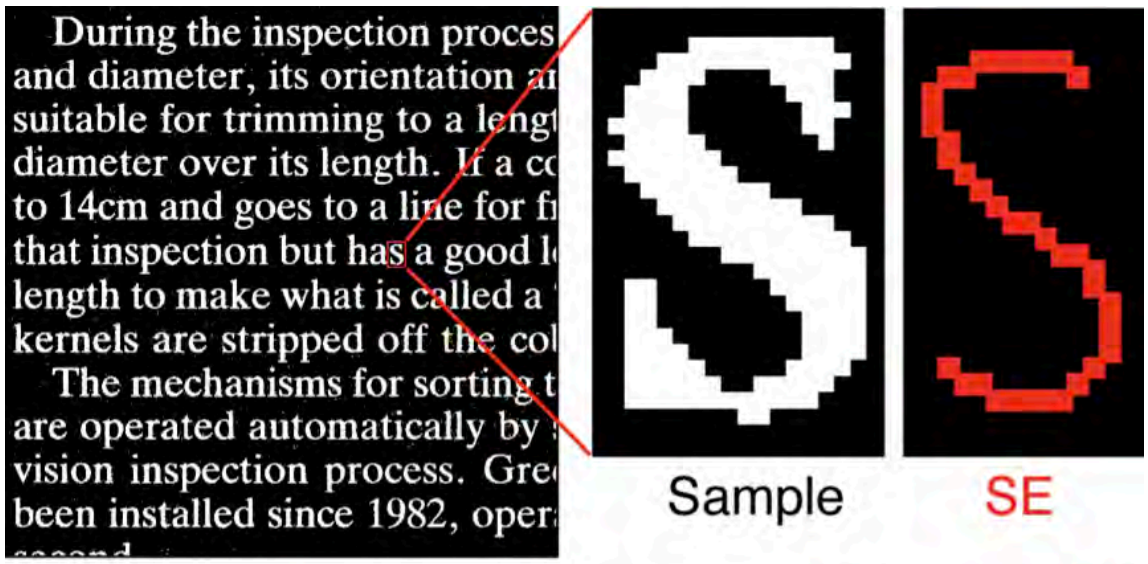

(Top-left) Original image. **(Top-right)** Create a Structuring Element (SE) by thinning a sample of lower-case "**s**".

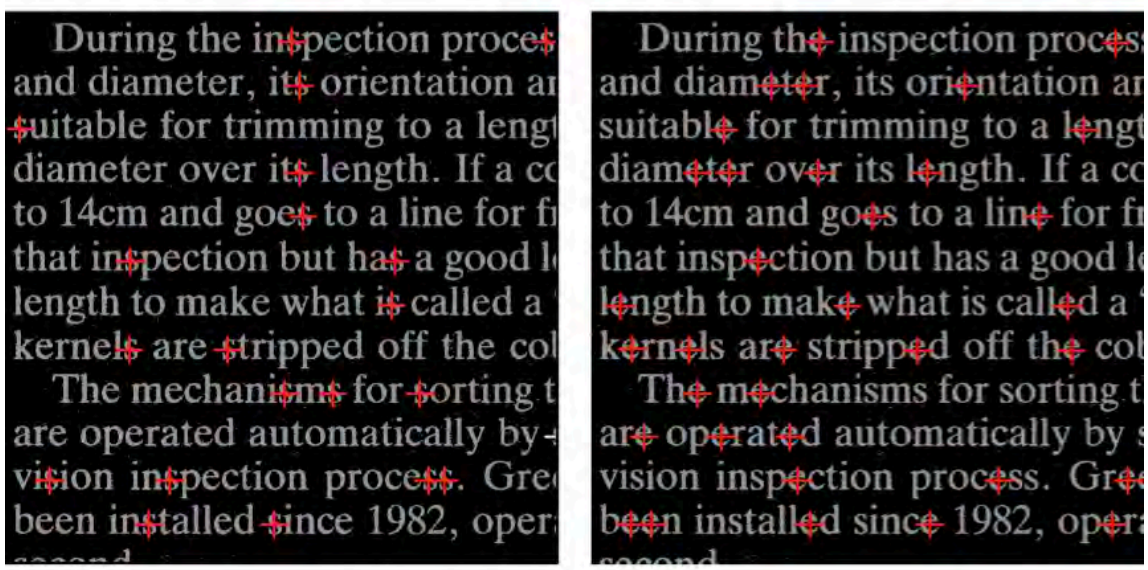

(Centre-left) Red crosses indicate where the SE fits the text. **(Centre-right)** Letter "**e**" detected. SE designed for that purpose.

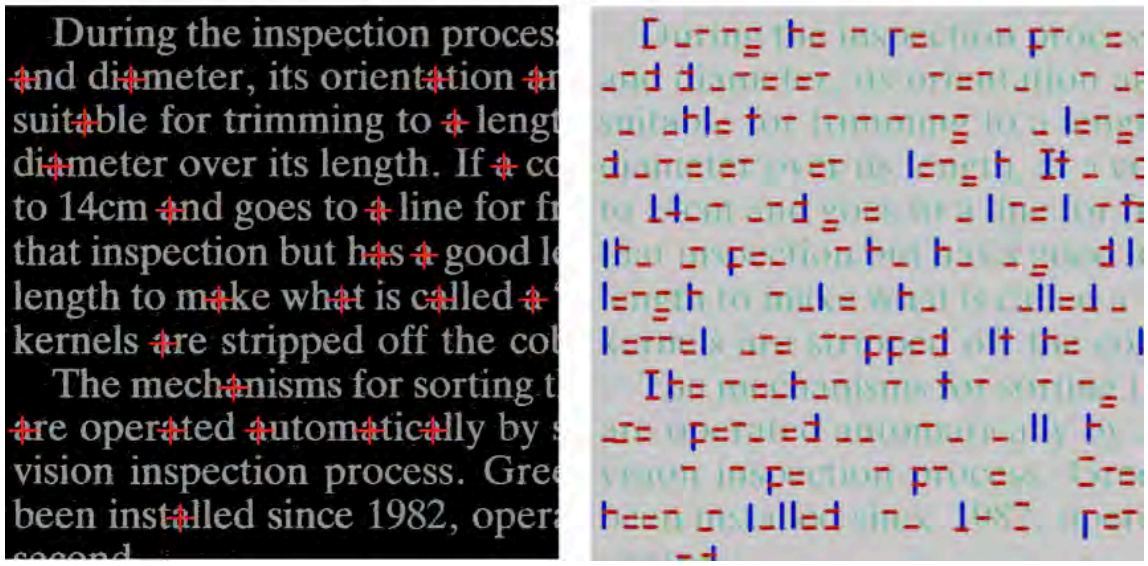

(Bottom-left) Letter "**a**" detected. SE designed for that purpose. **(Bottom-right**) Parts of letters detected by erosion with vertical bars

COMPLEX REPEATING PATTERNS (woven silk neck-tie)

New-comers to Machine Vision almost invariably expect more of template-matching than is achieved in practice. Slight variations in the product, that are barely discernible by eye, often lead to a poor match between the repeating pattern and template. For this illustration, a template was derived from one of the repeated patterns (elephant in the red rectangle). Cross-correlation (a grey-scale function) finds a good match between the template and each of the other elephants. Once the (local) best-fit positions have been found, it is a trivial matter to complete the inspection process, using simple image subtraction.

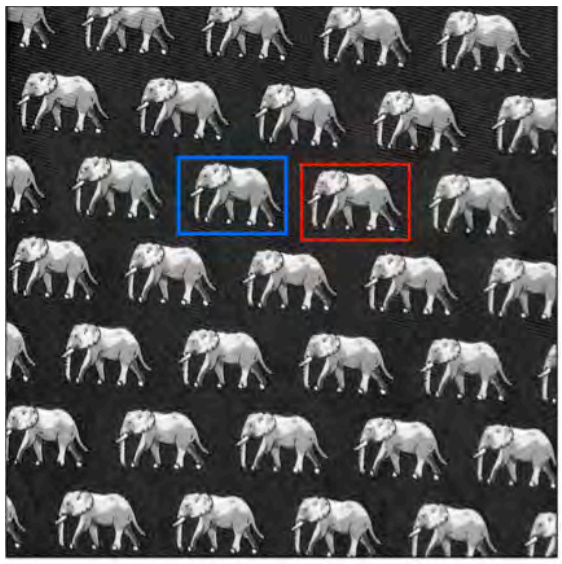

(Top-left) Original image.

(Top-right) Template enlarged.

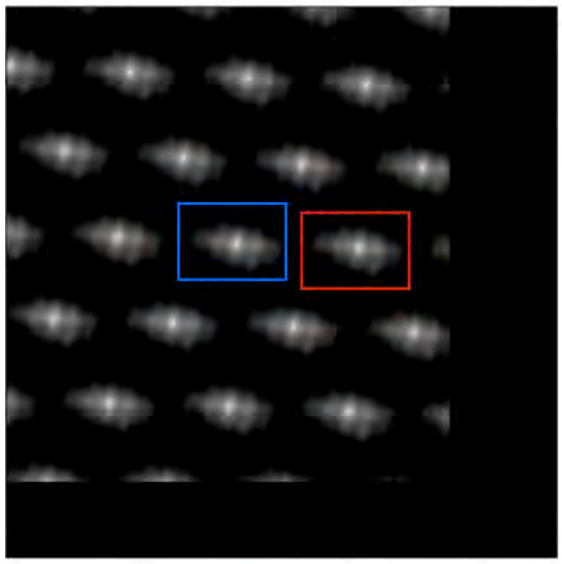

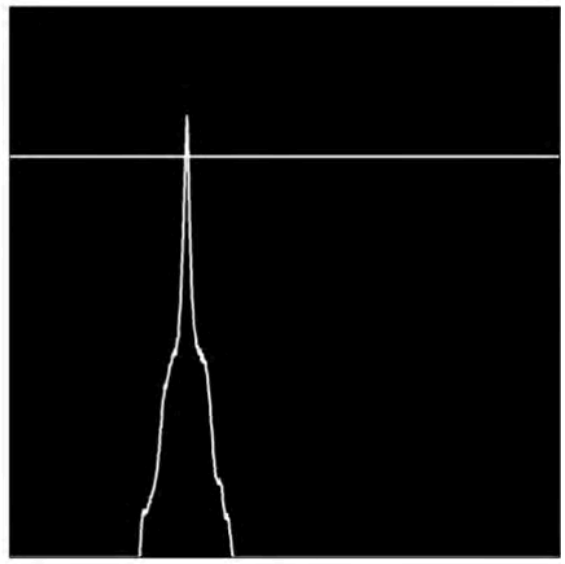

(Centre-left) Cross-correlation map.

(Centre-right) Intensity profile through the spot in the blue rectangle in [CL].

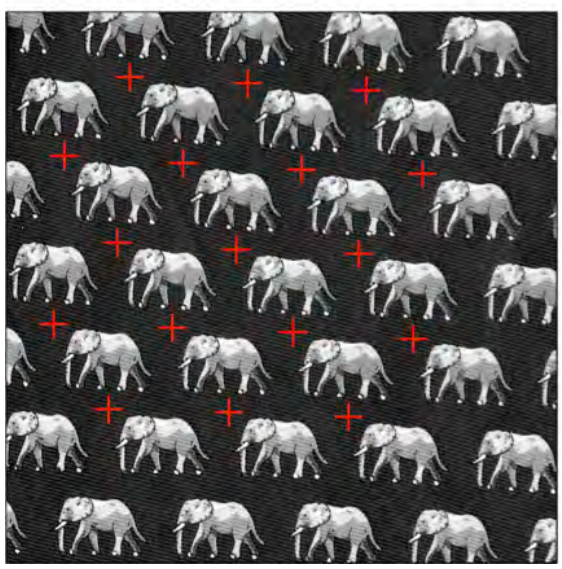

(Bottom-left) Centres of the spots in [CL]. Each cross indicates where the template matches an elephant.

(Bottom-right) Template compared to each best-fit sub-image. For a perfect match, the corresponding rectangle is uniformly grey.

QR CODE (locating “anchor” points, prior to reading)

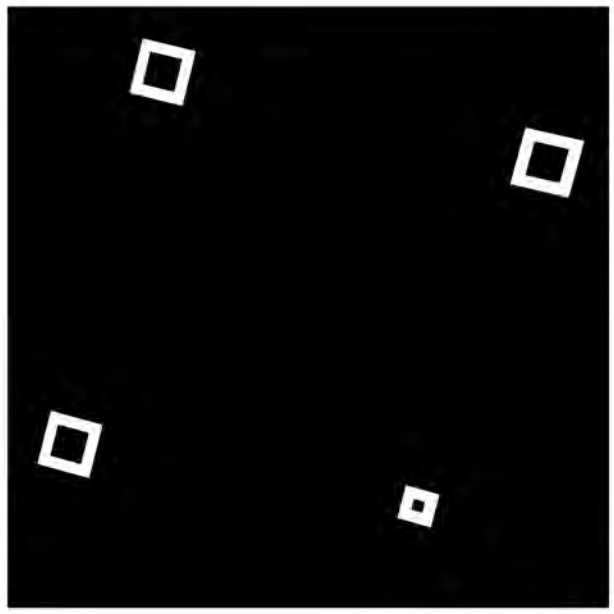

(Top-left) Original image

(Top-right) “Anchor" points, identified as large (black) islands in islands.

(Centre-left) Lines joining the anchor points provide reference axes for normalising orientation

(Centre-right) Normalising orientation (rotating [TL] so that the upper red line in [CL] is horizontal.

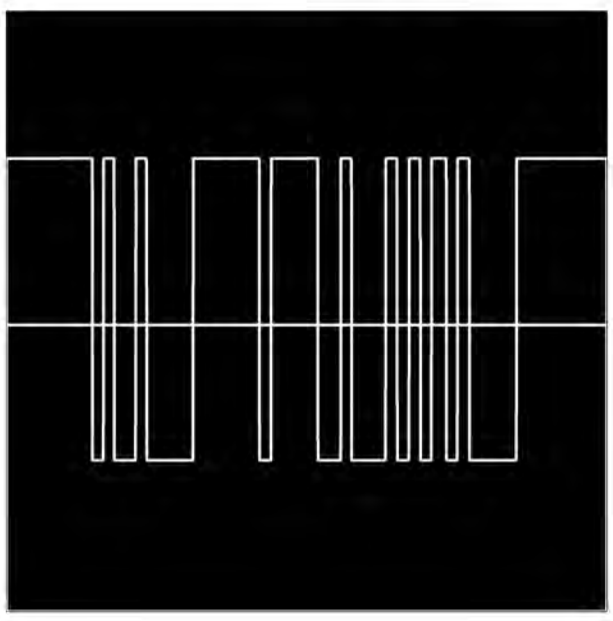

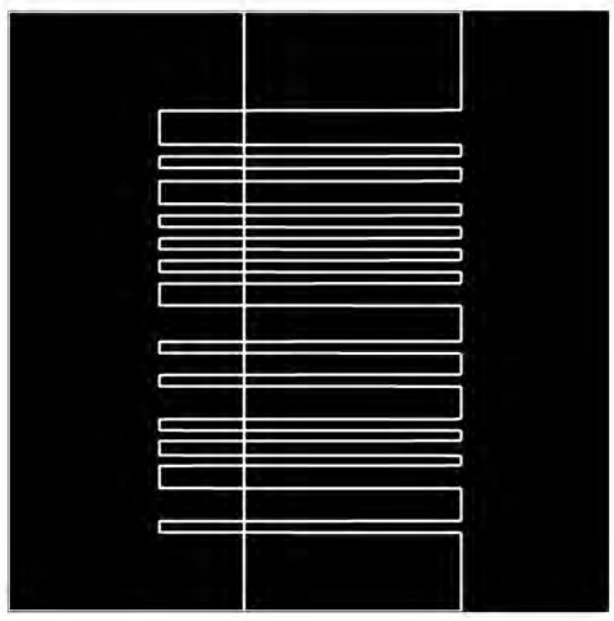

(Bottom-left) intensity profile along the horizontal line is like a bar-code, representing a series of 0s and 1s.

(Bottom-right) intensity profile along the vertical line also produces a "bar-code" of 0s and 1s..

Section 10

ALLSORTS

ASBESTOS FIBRES

(Top-left) Original image.

(Top-right) Binary image.

(Bottom-left) Hough Transform of [TL].

(Bottom-right) Straight lines detected in [TR] by finding the intensity peaks in [BL].

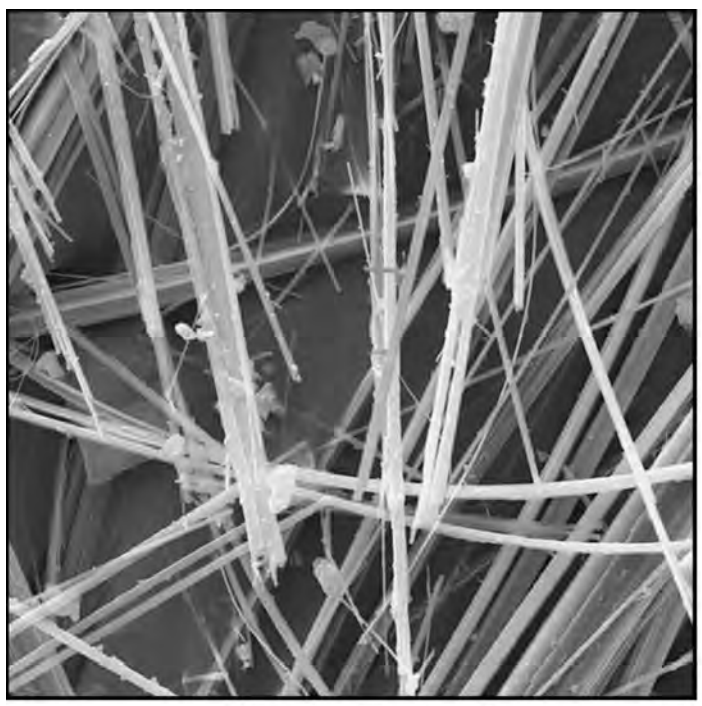

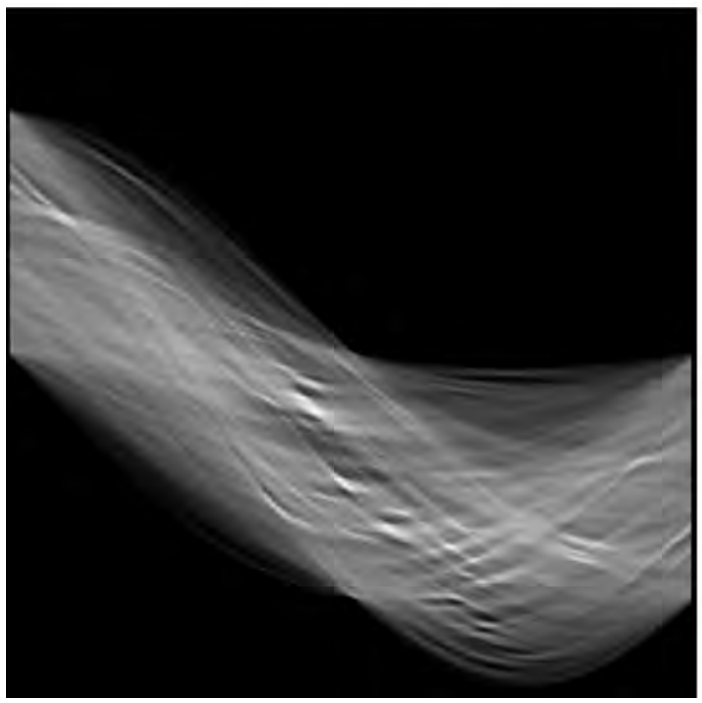

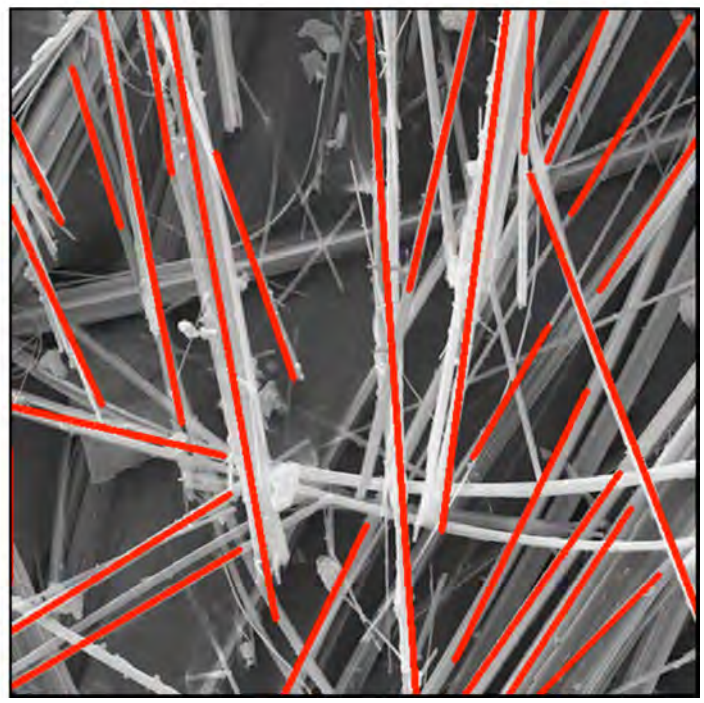

DUST (compact particles)

(Top-left) Original (binary) image.

(Top-right) Colour coding blob area. Big blobs (dust particles) are red; mid-size, orange and green and small ones blur.

(Bottom-left) Showing the number of blobs in the indicated size ranges.

(Bottom-right) Blobs were removed progressively, by erosion, using a 3x3-pixel square structuring element. The right-most column indicates how many blobs remain after each stage. E.g. 64 blobs “survived” repeating erosion 6 times.

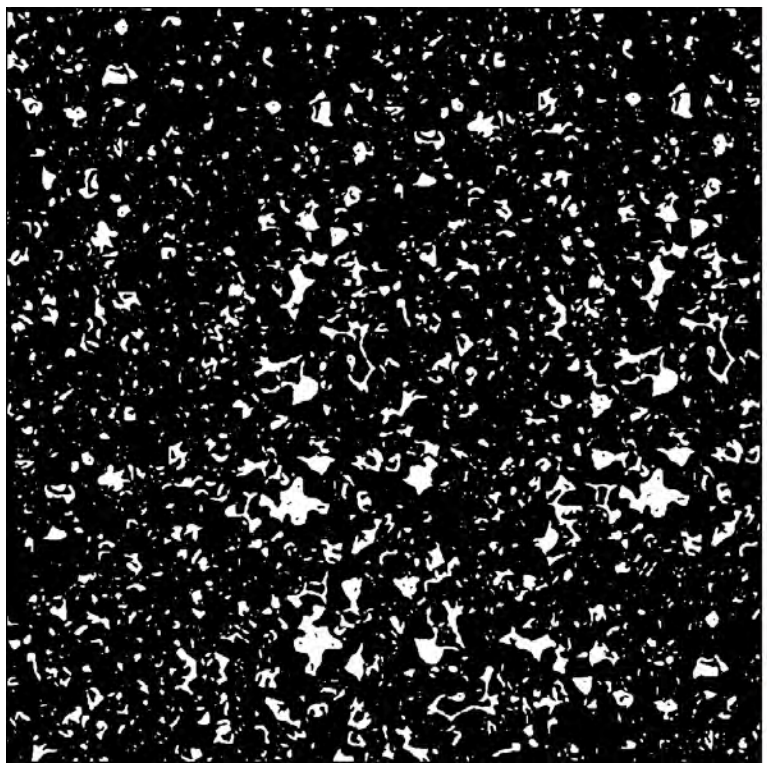

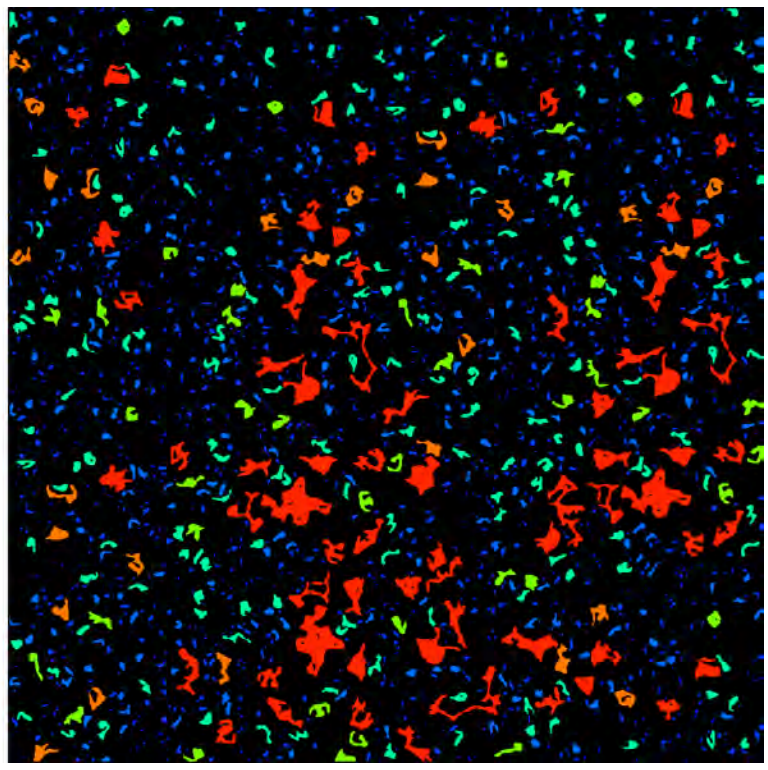

Table 1: Size (area) distribution	
Blob area (pixels)	**No. of blobs**
< 25	1732
25 - 50	306
50 - 100	235
100 - 200	165
200 - 300	41
300 - 400	27
400 - 500	20
> 500	48

Table 2: No. of blobs “surviving” erosion	
Level of erosion	**No. of blobs**
No erosion	1603
1	1087
2	655
3	404
4	224
5	133
6	64
7	36

DUST (short, usually isolated fibres)

(Left) Original image.

(Right) Ends of the fibres, detected by finding the limb-ends of the skeleton.

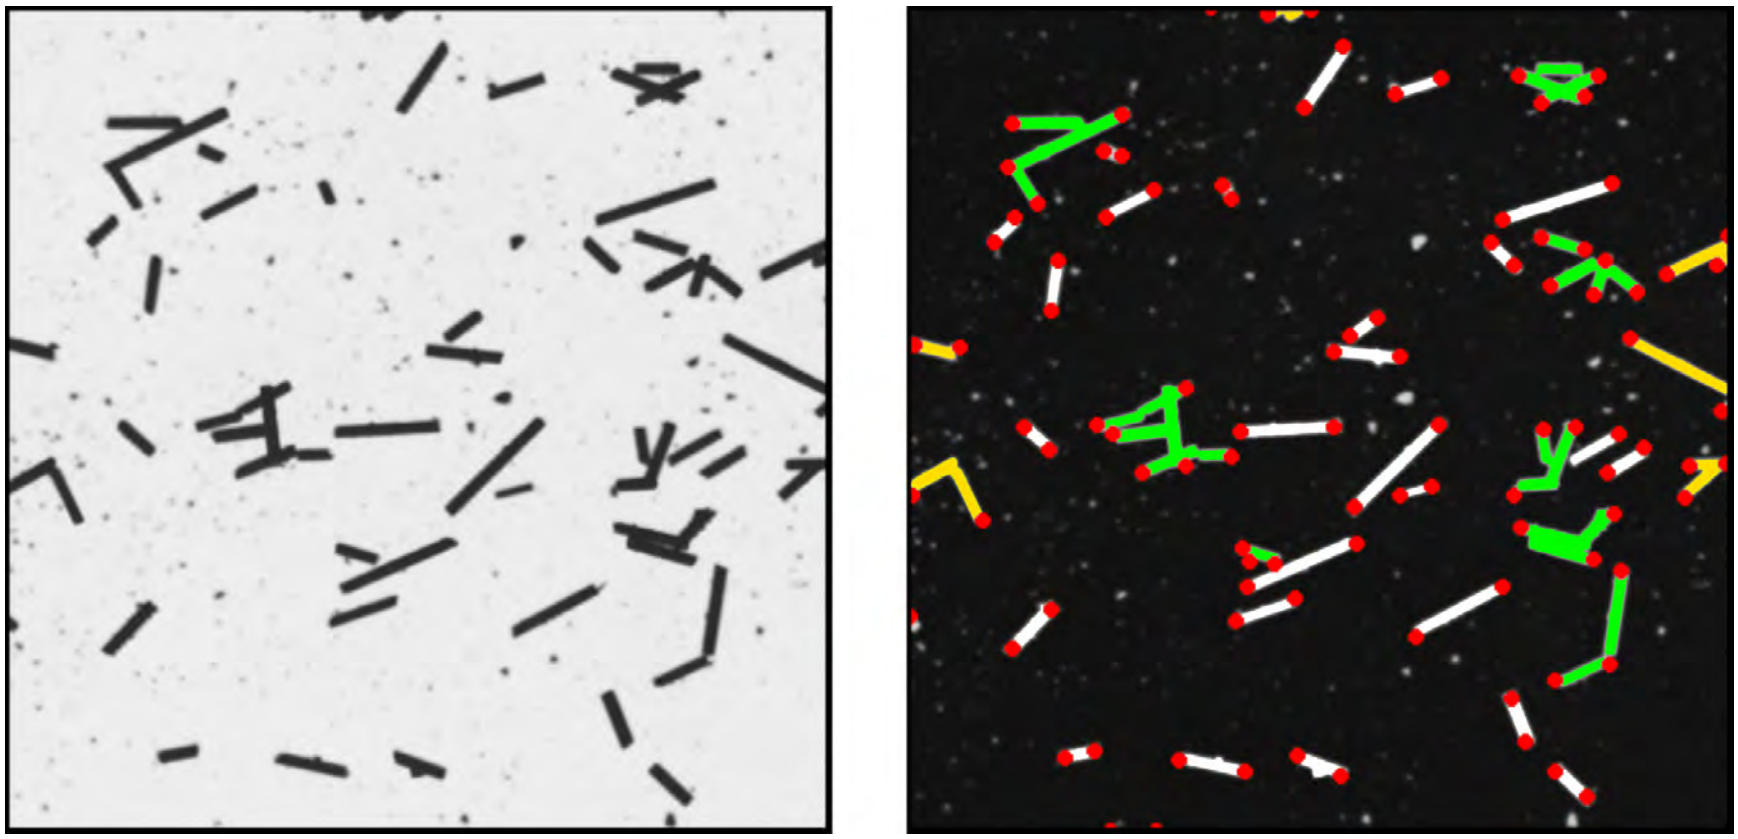

COUNTING BACTERIA ON A CULTURE PLATE (low density culture)

In this application, separating the RGB colour components is particularly useful. The yellow grid is superimposed on the culture-growth material (impregnated paper sheeting) and provides an aid for human visual analysis. A pre-calibrated machine vision system does not need the grid: the G-image simply ignores it, while its contrast is greatly enhanced in the B-image.

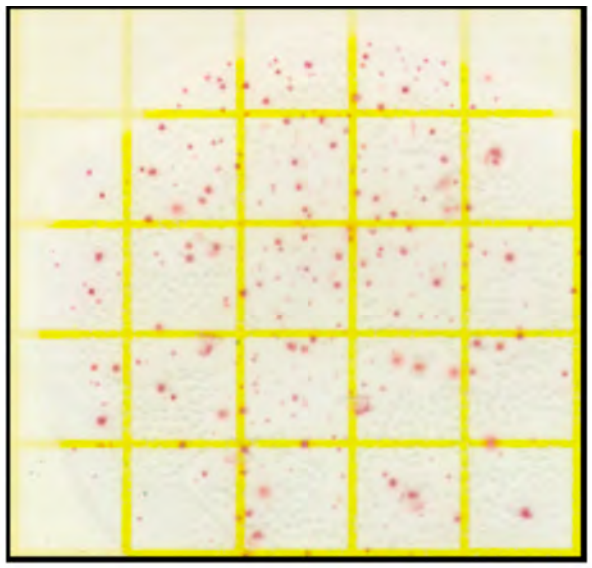

(Top-left) Original image.

(Top-right) R-image.

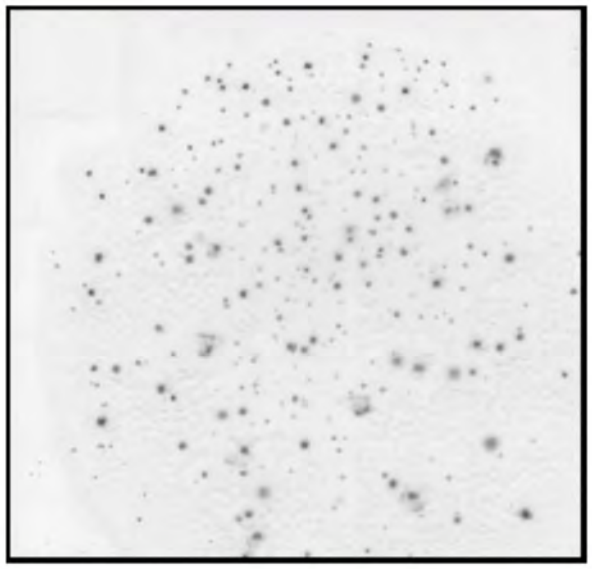

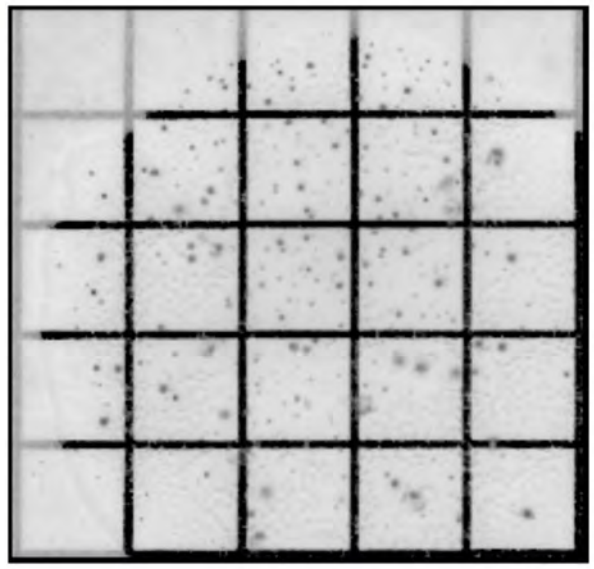

(Centre-left) G-image.

(Centre-right) B-image.

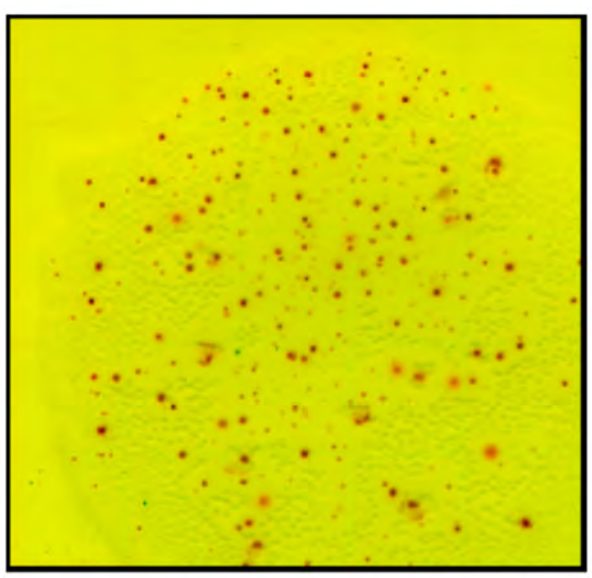

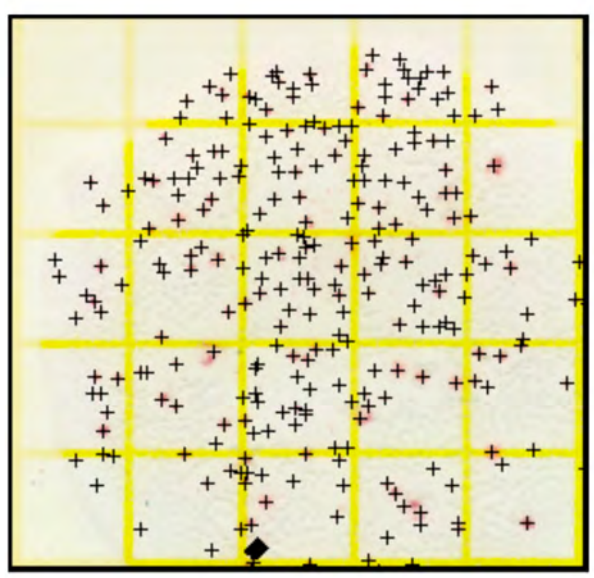

(Bottom-left) A colour image was reconstituted after each of the RGB channels was enhanced separately.

(Bottom-right) Crosses mark the centroids of the red spots found by a colour-recognition filter designed to detect "red" while ignoring "yellow".

COUNTING BACTERIA ON A CULTURE PLATE (high density culture)

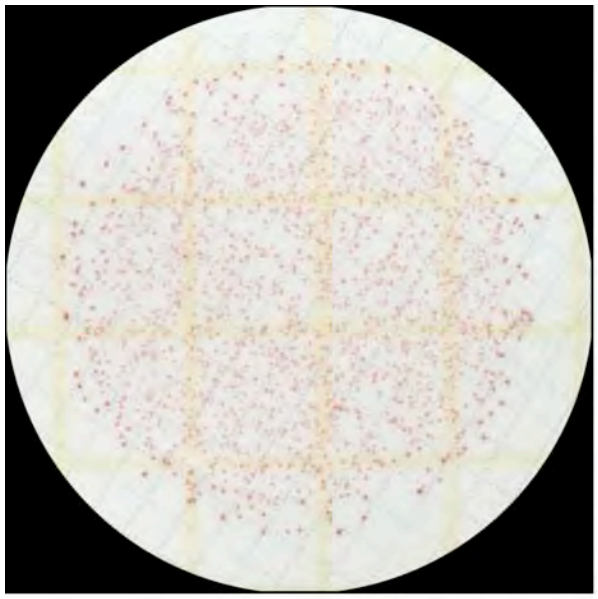

(Top-left) Original image.

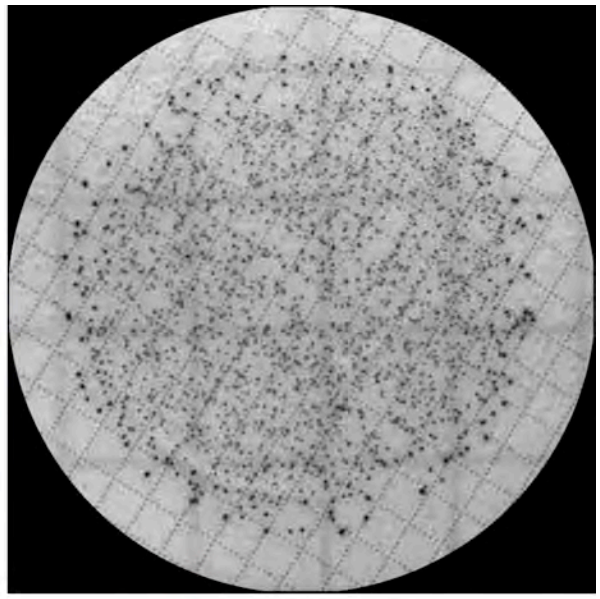

(Top-right) Intensity component of [TL]. (I-Image).

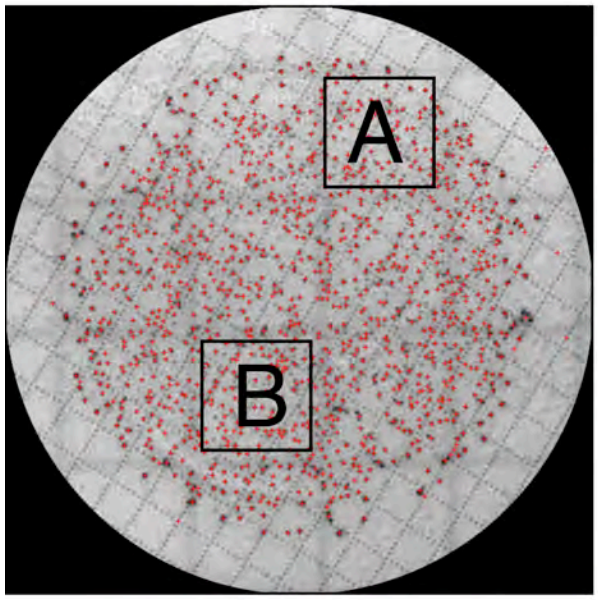

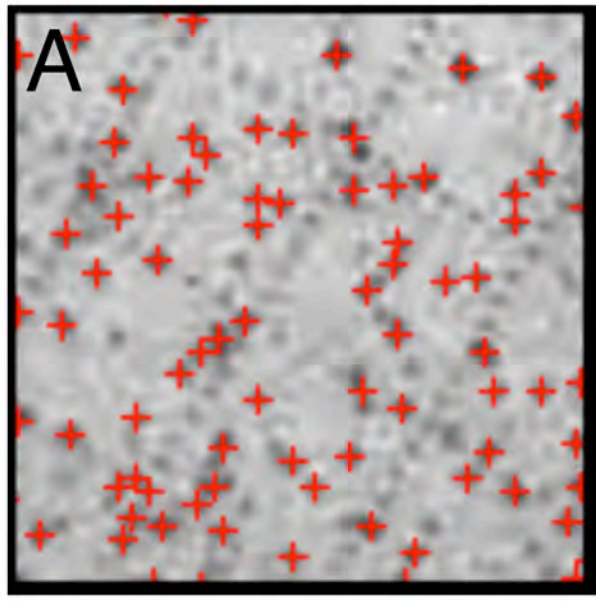

(Centre-left) Placing a cross on every dark spot.

(Centre-right) Magnified view of region A in [CL].

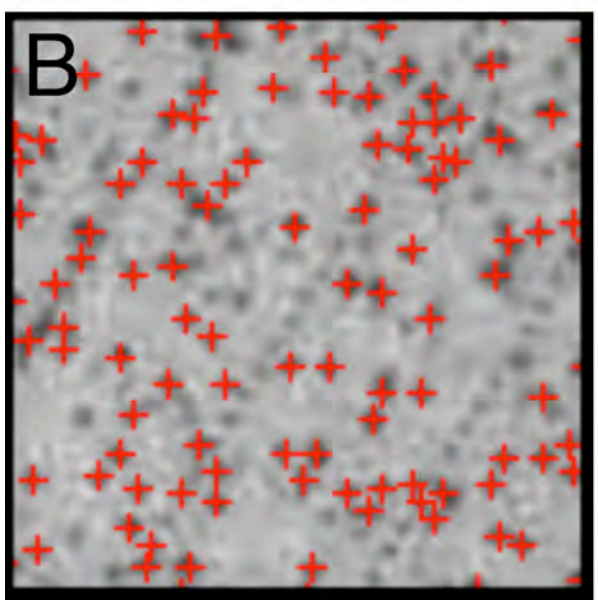

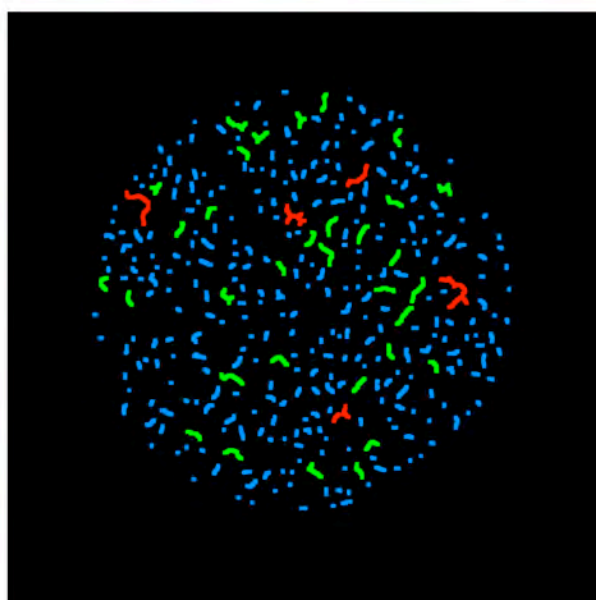

(Bottom-left) Magnified view of region B in [CL].

(Bottom-right) Green: Three or more merged spots; Red: Two merged spots; Blue: Isolated spots. Blob-size is used as the criterion for deciding whether merging has occurred.

CREASED FABRIC

A large high-street clothing retailer was interested in quantifying the creasing properties of fabrics, such as those used in shirts, blouses and dresses.

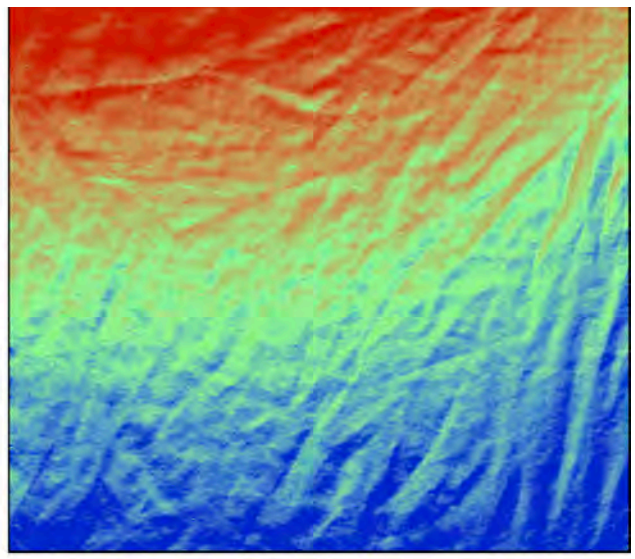

(Top-left) Original image. Low-angle illumination. No attempt was made to crease the fabric in a controlled manner..

(Top-right) Pseudo-colour.

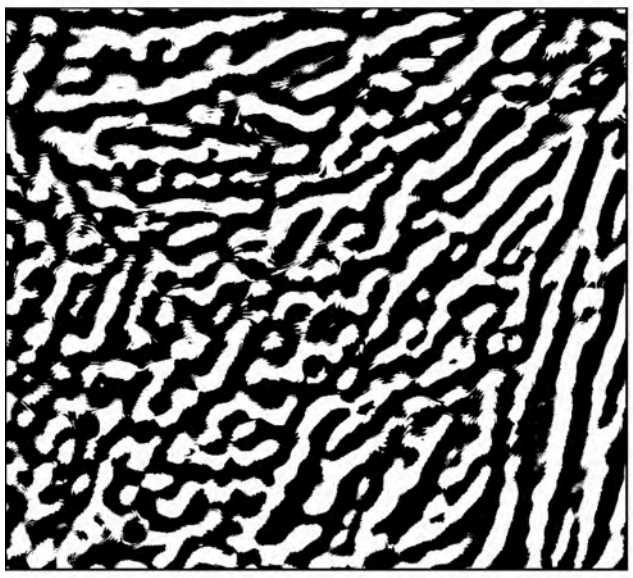

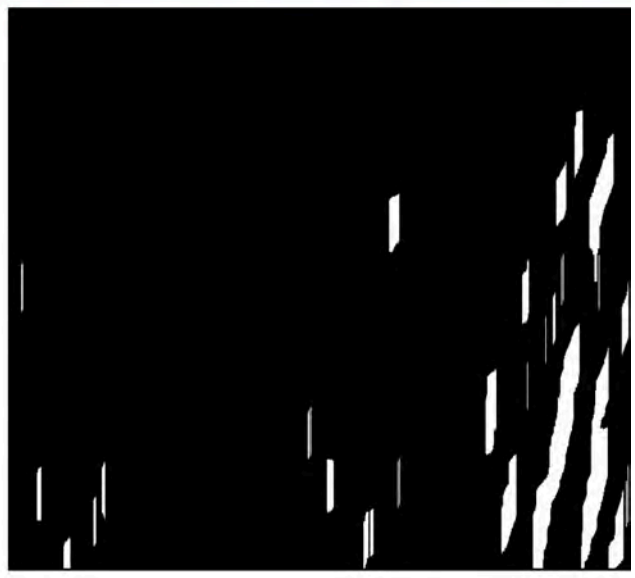

(Centre-left) Binary image, after filtering and thresholding.

(Centre-right) Strong vertical features detected by binary morphology.

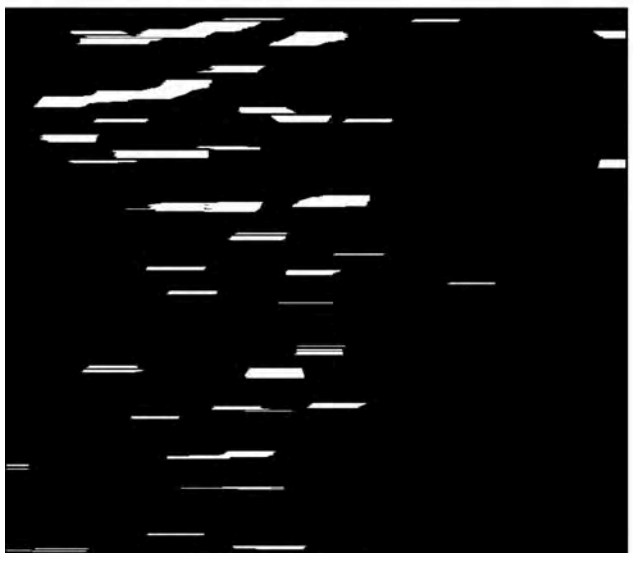

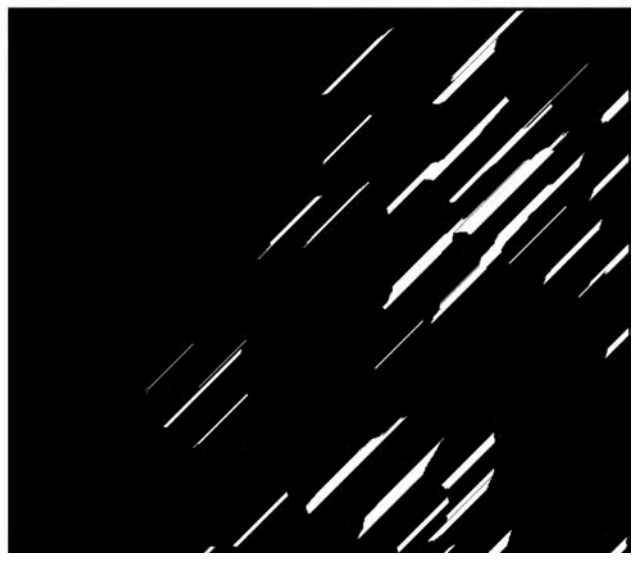

(Bottom-left) Strong horizontal features-detected.

(Bottom-right) Strong diagonal (/)

KNITTED FABRIC

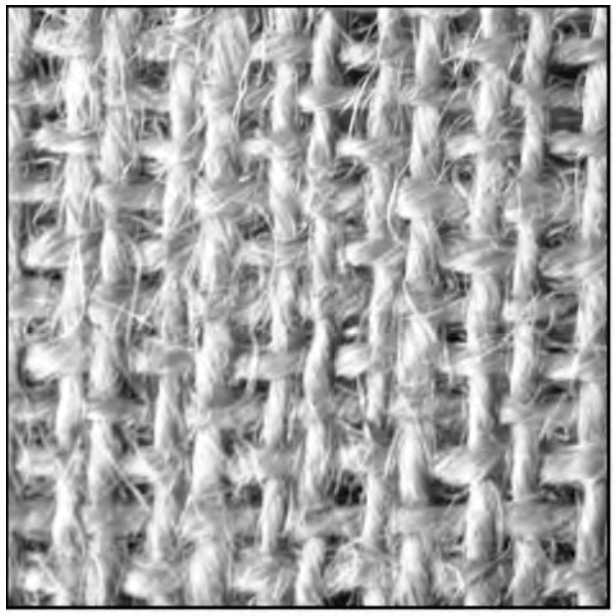

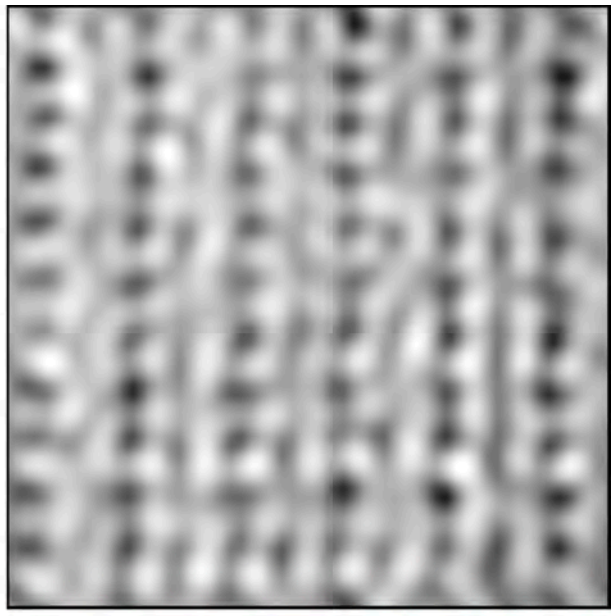

(Top-left) Original image.

(Top-right) Low-pass (blurring) filter.

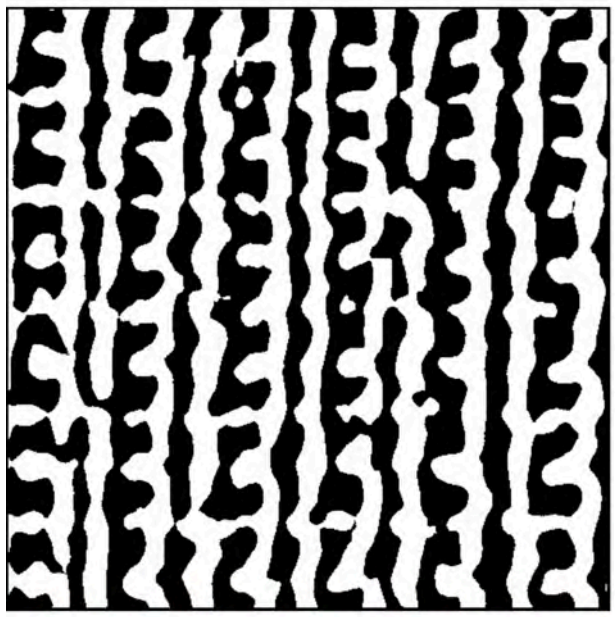

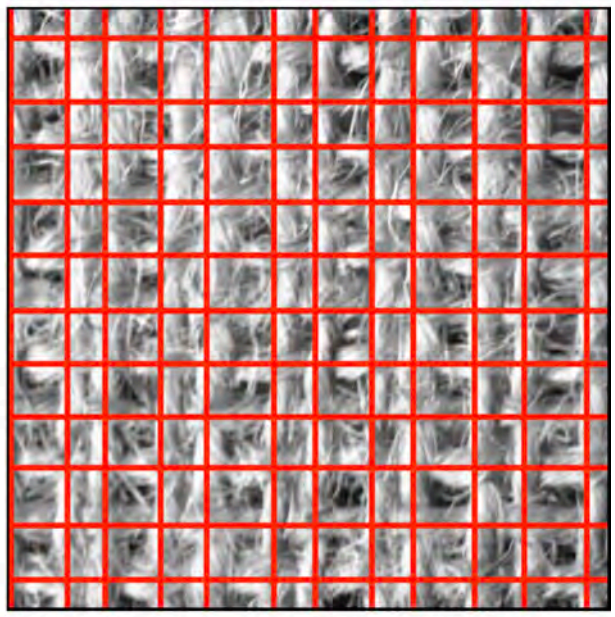

(Centre-left) Directional low-pass filtering, followed by thresholding.

(Centre-right) Vertical strands were detected by processing with morphology operators [CL]. Horizontal strands were similarly identified.

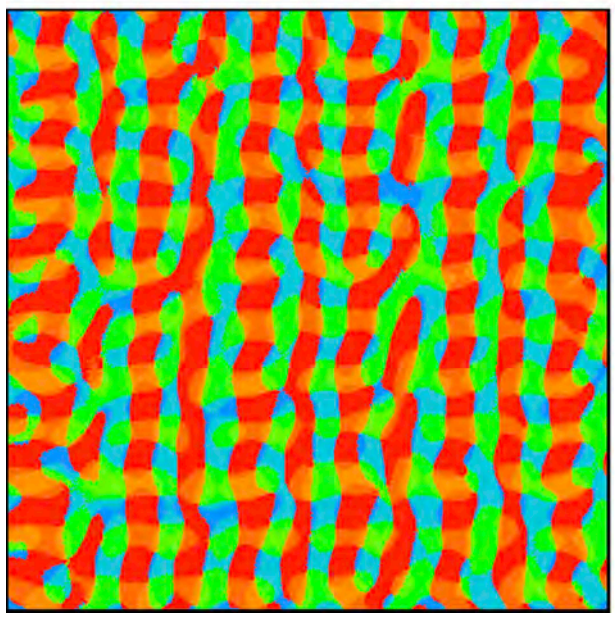

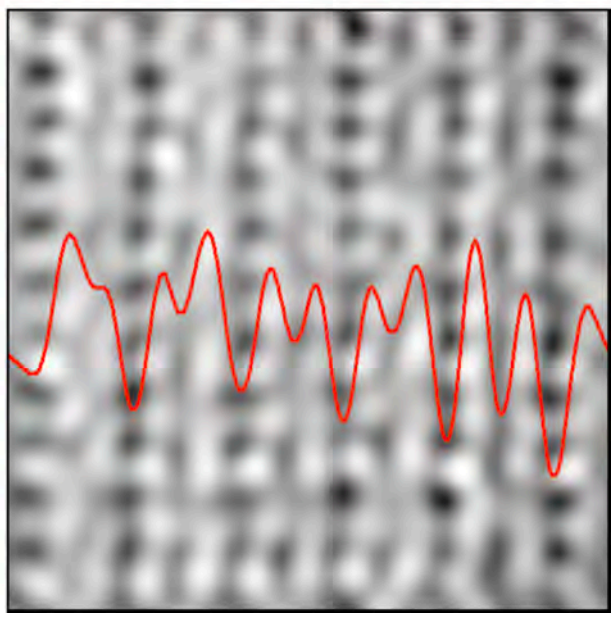

(Bottom-left) Direction of intensity gradient in [TR], colour coded.

(Bottom-right) Profile of the column-average intensity.

TEXTURE (travertine floor tiles)

Texture inspection is needed in the food industry (e.g. baked goods), leather working, car body paintwork, wood-working, decorative crafts It is a complex process combining elements of Machine Vision and Machine Learning; a full explanation is far beyond beyond the scope of this short chapter. Let it suffice to say that each texture sample is processed in order to generate a set (strictly, a vector) of numerical measurements. A texture sample is then classified by comparing its corresponding measurement vector to the vectors derived from known "good" textures. There is no difficulty in generating a multitude of measurements from a texture image. Machine learning is needed to decide which measurements are most effective and how they should be combined. It is usually impossible to write a formula for recognising "good" texture, in which case, teaching by showing is needed.

(Top-left) Four samples of travertine bathroom tiles.

(Top-right) Intensity. (I-image)

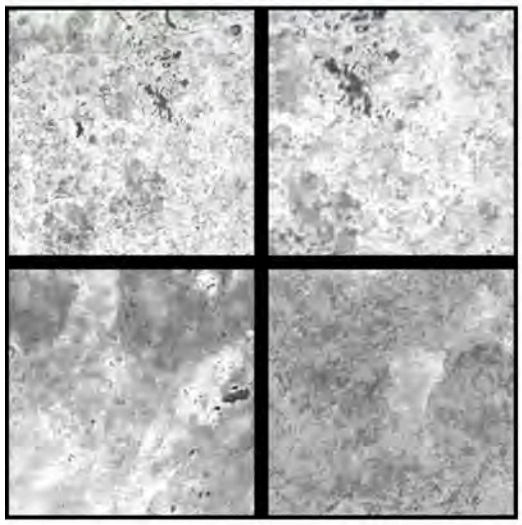

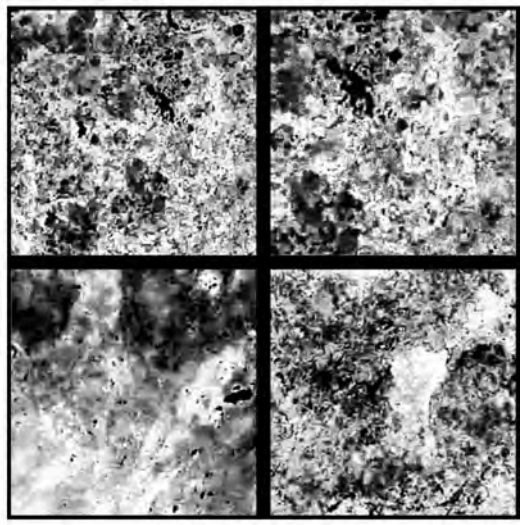

(Centre-left) Saturation gives good contrast. (S-image)

(Centre-right) Histogram equalisation of [TR], applied to each quadrant separately.

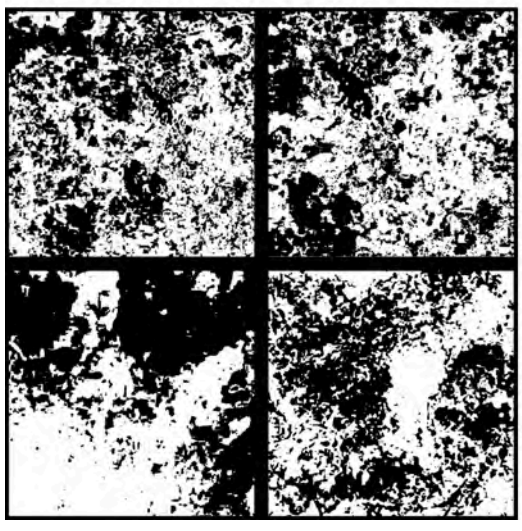

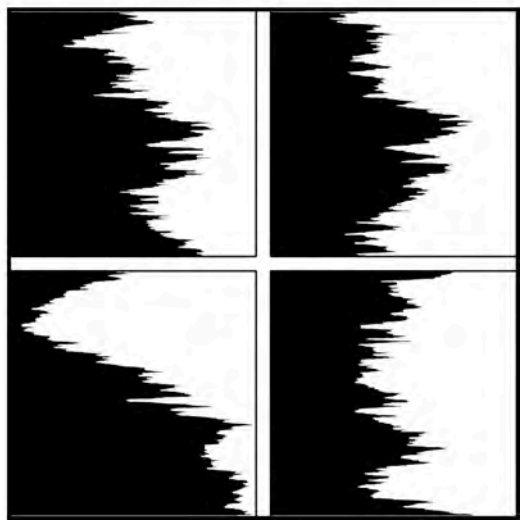

(Bottom-left) [CR] thresholded at mid-grey. Exactly 50% of each image is white

(Bottom-right) *Row-integration of [BL] showing number of white pixels in each row*

TEXTURE (wood floor tiles)

Texture analysis is a complex process combining elements of Machine Vision and Machine Learning; a full explanation is far beyond beyond the scope of this short chapter. Let it suffice to say that each texture sample is processed in order to generate a set (strictly, a vector) of numerical measurements. A texture sample is then classified by comparing its corresponding measurement vector to the vectors derived from known "good" textures. There is no difficulty in generating a multitude of measurements from a texture image. Machine learning is needed to decide which measurements are most effective and how they should be combined. It is usually impossible to write a formula for recognising "good" texture, in which case, teaching by showing is needed.

(Top-left) Four samples of wood floor tiles.

(Top-right) Intensity. (I-image)

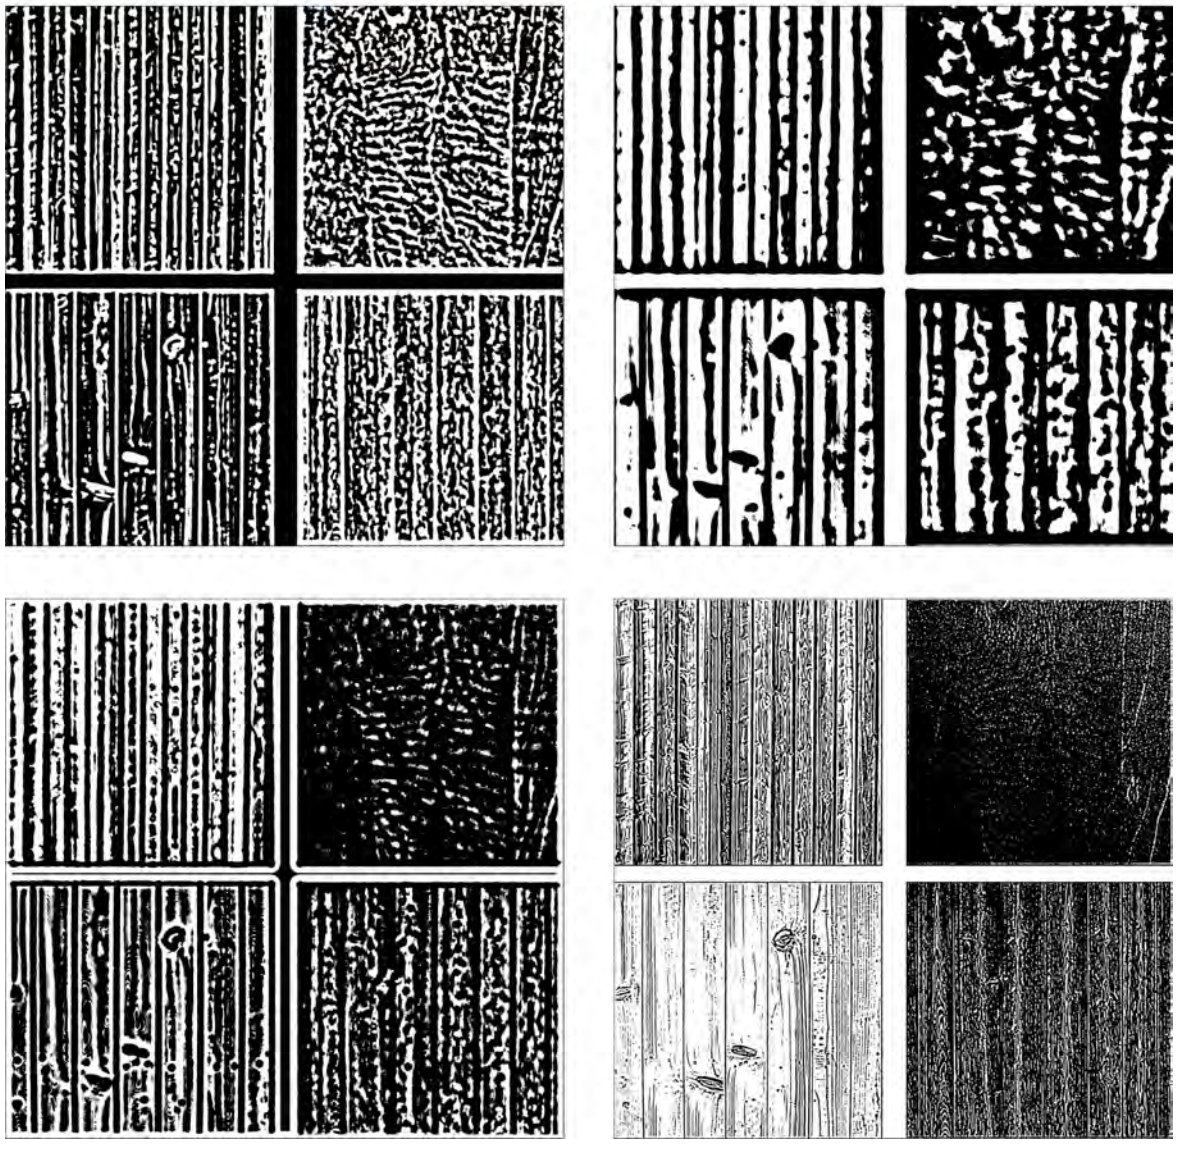

(Others) Saturation gives good contrast. Various type of filter. Notice how well the filter distinguishes different textures.

WATER DROPS ON GLASS (size analysis)

(Top-right) Original image. Back lighting

(Top-right) Binary image derive from[TL].

(Bottom-left) Blobs shaded in [TR] according to their areas: biggest blobs are brightest.

(Bottom-right) Pseudo colouring of [BL].

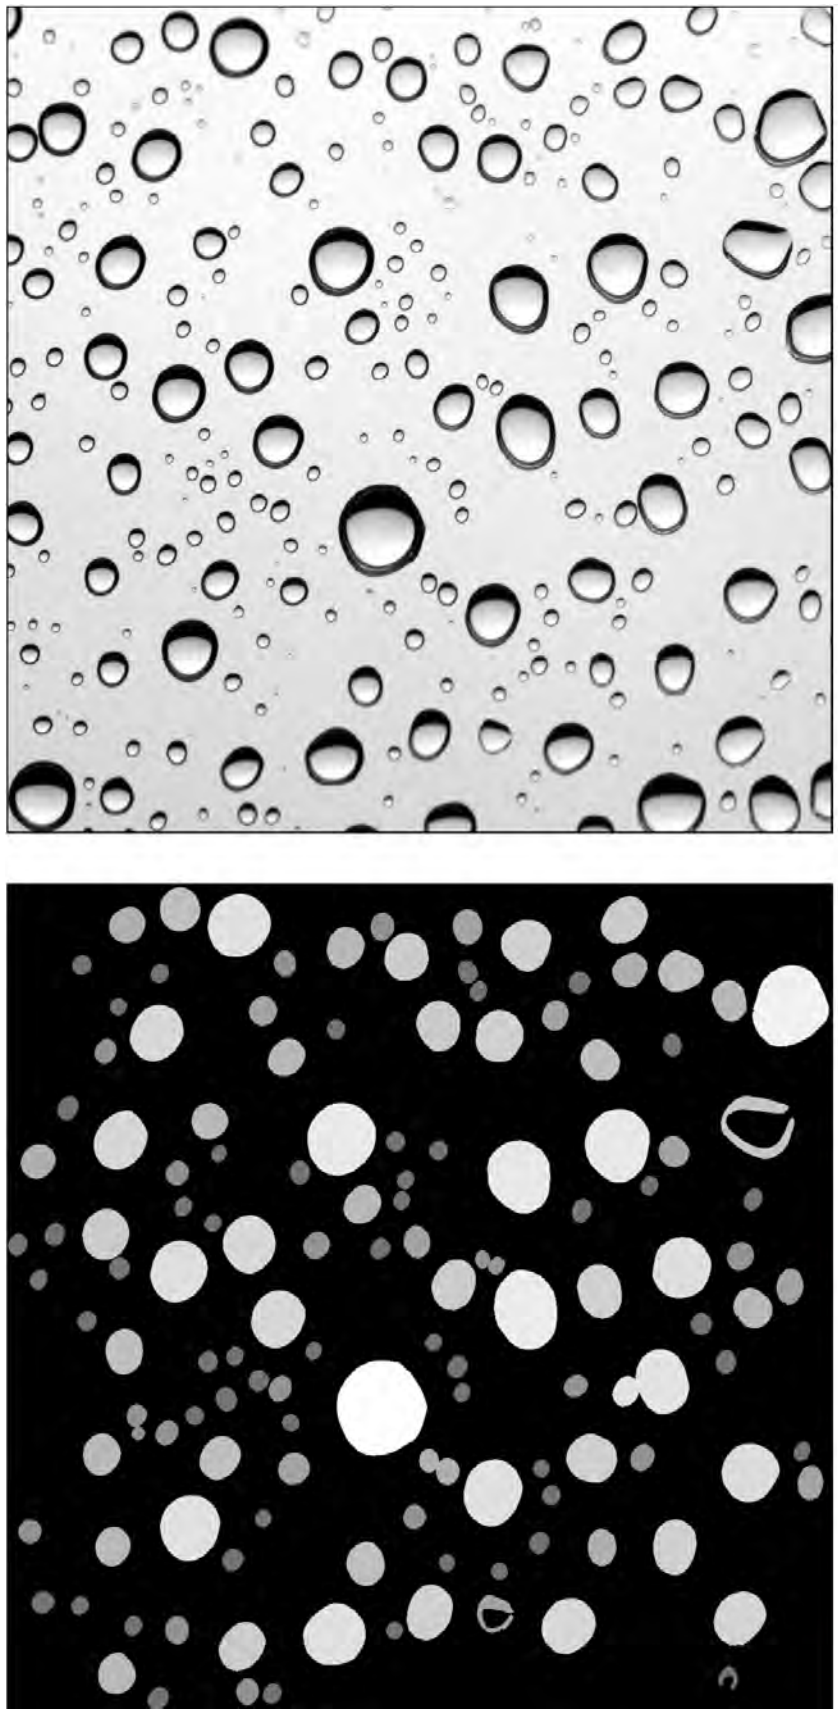

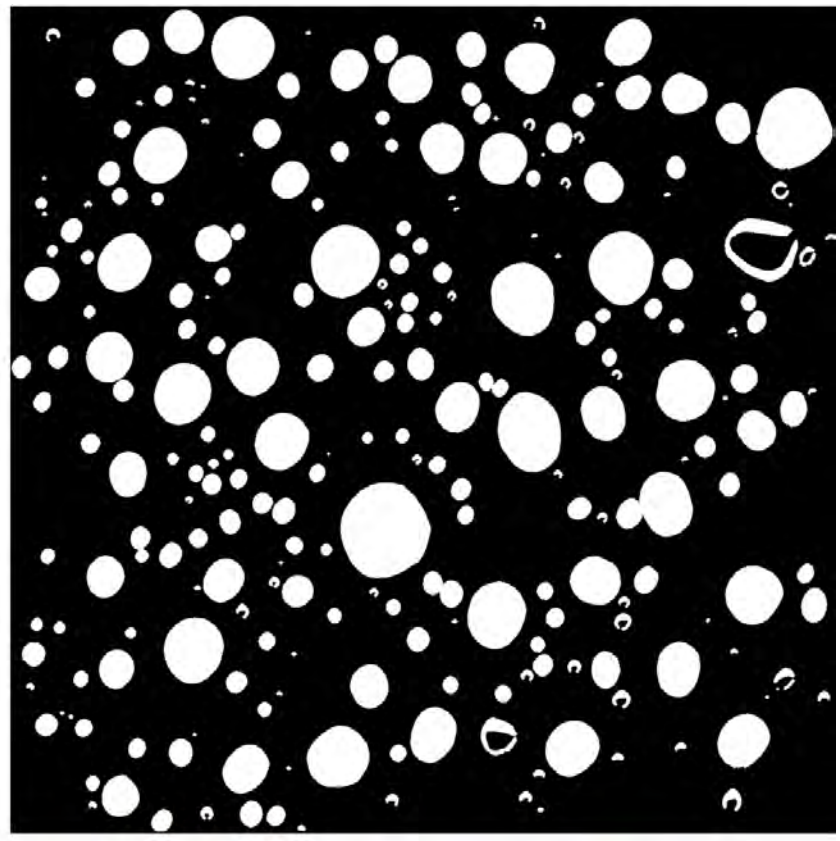

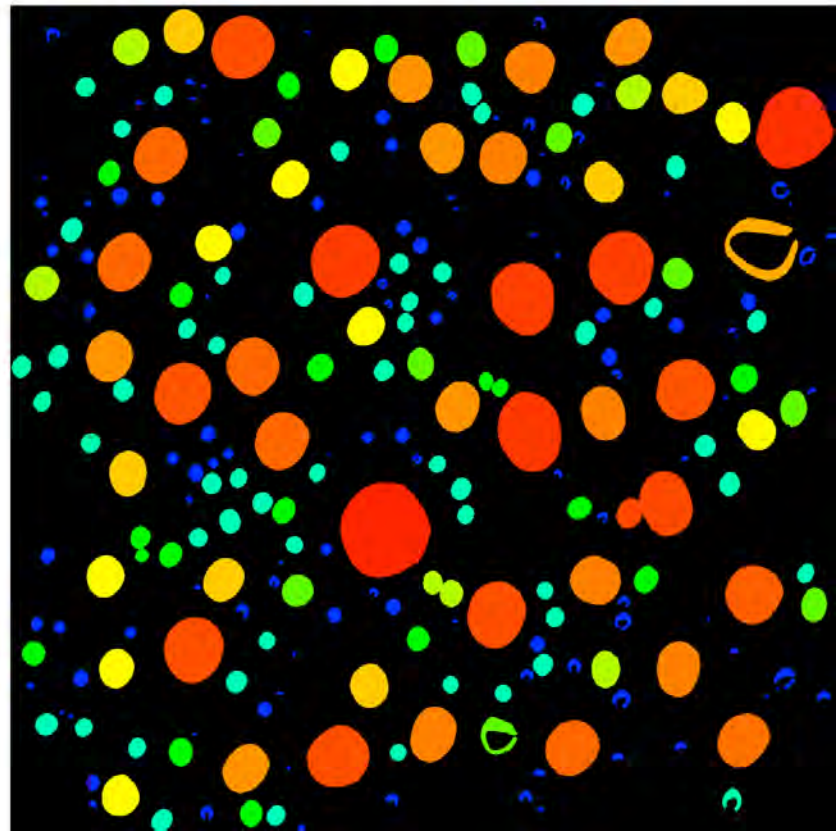

SKIN - MOLES

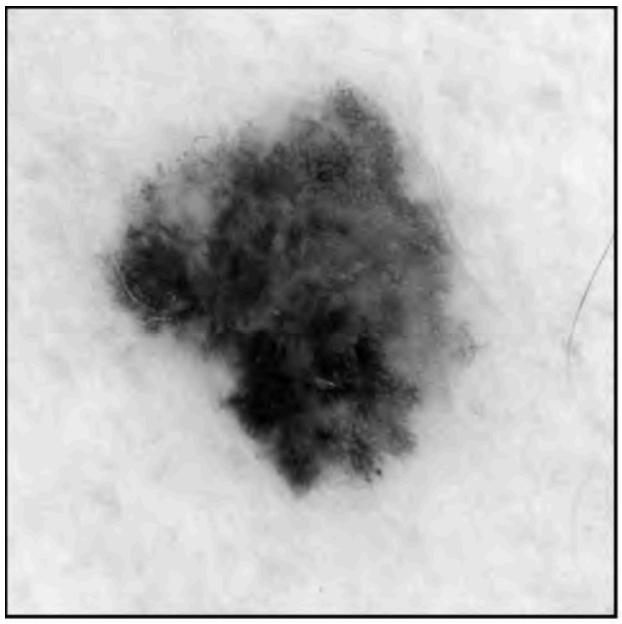

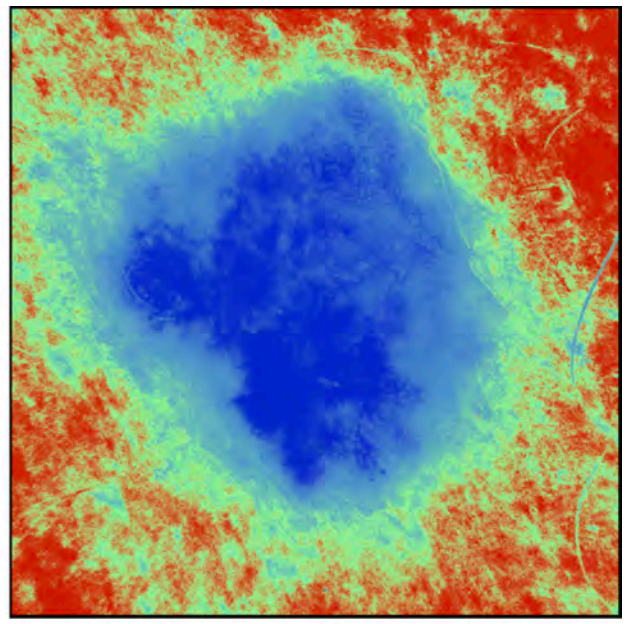

(Top-left) Original image.

(Top-right) Pseudo-colour.

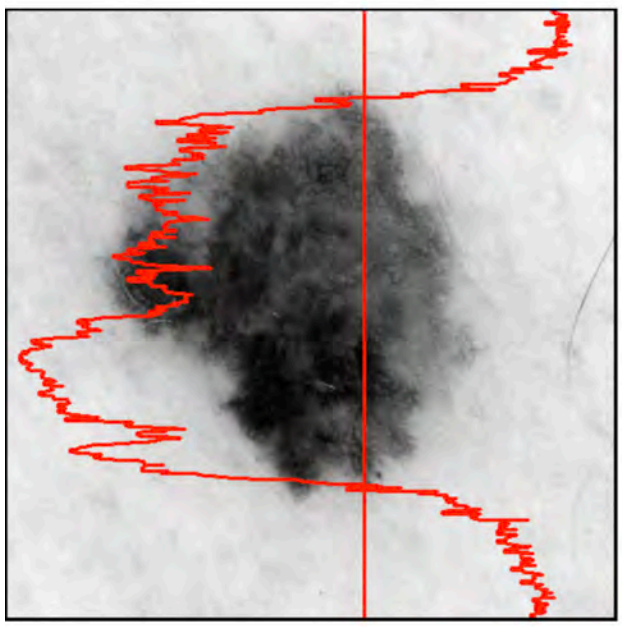

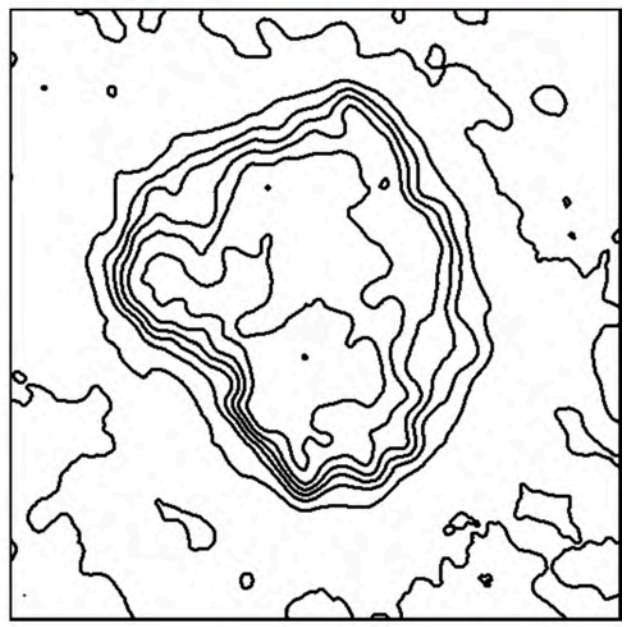

(Centre-left) Intensity profile.

(Centre-right) Intensity contours.

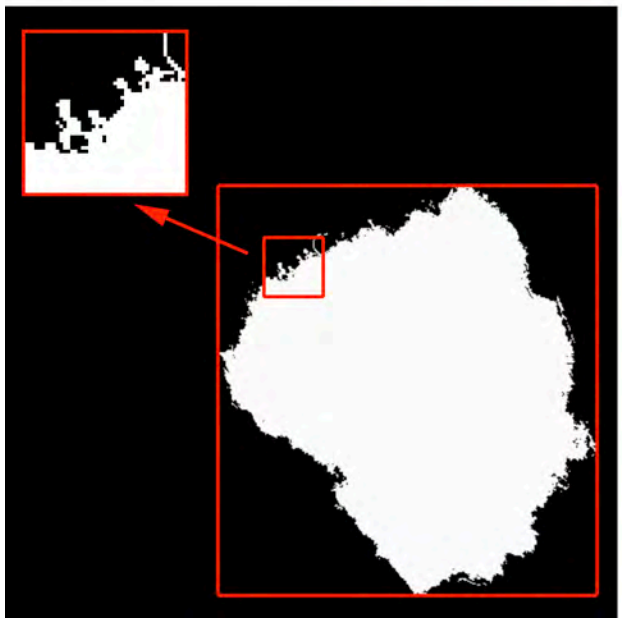

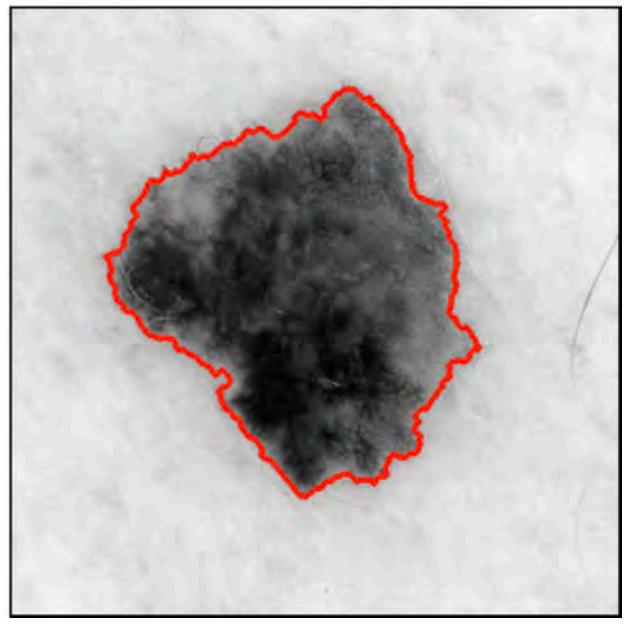

(Bottom-left) Binary image, Inset, showing the “ragged” edge.

(Bottom-right) Edge contour.