

## The CAA93 image processing competition

Stephen Bullas<sup>1</sup>, Mike Canter<sup>2</sup>, Mike Fletcher<sup>3</sup>, Martijn van Leusen<sup>4</sup> & Dick Spicer<sup>3</sup>

(<sup>1</sup>*Strategic Decisions Ltd, Leatherhead, UK;* <sup>2</sup>*Brinna House, Bigton, Shetland, UK;* <sup>3</sup>*School of Computing, Staffordshire University, Stoke-on-Trent, UK;* <sup>4</sup>*Instituut voor Pre- en Protohistorische Archeologie, Amsterdam, The Netherlands*)

### 2.1. Introduction (Mike Fletcher and Dick Spicer)

In the invitation for papers for CAA93 an image processing competition was announced. Those wishing to participate were sent three PCX images to process. The purpose of the competition, in addition to providing some amusement for all concerned, was to discover the extent of image processing in archaeology, and to allow those concerned to identify each other informally. The number of requests (nearly 50) for the image disks was itself rather intriguing, indicating to us that there were a great many people interested enough in image processing to be prepared to tackle the problem. Early replies were encouraging, even though many were letters which came from frustrated individuals who had given up at the first stages. Those people who completed and entered an attempt were few: the results they submitted, however, were mainly of a very high standard.

#### 2.1.1. Creation of the test images

Although the idea of producing a set of "scrambled" images seemed very easy, it was not simple in practice. One or two rules were quickly decided upon: colour was not to be used, since this might unfairly advantage those with colour processing and printing facilities; in addition, it should be possible to extract images with fairly modest — and fast — processing. The scrambling could be either grey-level dependent, spatially performed, or both. It was decided not to use the frequency domain for scrambling, since this would probably severely limit success to those with expensive programs, and would lay us open to too many accusations of elitism. We began, therefore, by examining simple grey-level operations.

Given an image whose intensities can range from zero to 255, it might at first sight seem possible to bracket two distinct ranges, by taking a dark image and a bright one and combining their grey-levels. However, as soon as two pixel values — each from the same place but from the two separate images — are added together, there is no way of separating them into their original values. Immediately we had one of the most common image processing problems, a grey-level segmentation. All archaeologists would be familiar with the slightly different brightness of crop-mark features within a field in an aerial photo. This, then, was chosen to be a fundamental, and archaeologically useful, task. The first image of our set was to be an aerial image of a field, with a clearly (?) recognisable object superimposed.

We chose the Uffington White Horse, a figure cut into the chalk hillside on the Oxfordshire/Berkshire border, be-

cause it is so well-known, appearing as it does in almost every coffee-table book on archaeology. The image was made into a binary one, with a completely black background, and a uniformly white horse, of value 20. Subtracting this image from the one of the field left a dark horse-shape as a crop-mark, each pixel of the original image being reduced in value by 20 under the horse. Looking at this picture (Fig. 2.1), it seemed that this was already quite tough enough to give out, and we decided to produce images which were rather easier to segment, using more technically biased image processing methods to merge them.

The second image was produced by merging a vertical "view" of Clonehenge, the imaginary computer-generated site which we produced for evaluating survey strategies, with a photo of a fairly well-known real henge. This was Maes Howe, in Orkney, with its familiar passage grave inside the circular ditch. We turned this upside down, just to make life a little more difficult (Fig. 2.2). Both original images made use of the full range of possible brightness values, so adding them together would have produced saturation at the high end, as well as making extraction almost impossible. We used a chequer-board pattern, of alternating single pixels, as a mask. Thus in the first row, the odd pixels were of the first picture, and even pixels of the second. The next row was the reverse of this, and so on. It was thus only necessary to generate the same chequer board to act as a key to separate the two images, and optionally to "fill in" lost detail by a low-pass filter or some other means. An enlargement of the portion of the image around the Clonehenge barrow can be seen in Fig. 2.3.

However, it seemed to us that not everyone would be willing — or even have the facilities — to write a program to generate such a key. We decided to use at the same time an additional method of encoding the images, using their grey-levels alone, thus providing a choice for entrants. Both original images had been produced using all 256 brightness levels. But one of them was quantised into 16 intensities, such that all values from 0–15 were given the value 0, all from 16–31 given 16, and so on, up to the final values 239–255 (set to 239). The overall effect of this one operation is not very noticeable, especially on an image with a good deal of rough texture (though contouring would have been seen if we had chosen the smooth slopes of Clonehenge). As can be seen from Fig. 2.4, the histogram of the combined image immediately gives the game away: high peak values at 0, 16, 32 etc., and uniformly low levels at other values. Extraction could be accomplished by isolating only the values of modulus 16, and these would carry all the pixels from




Figure 2.1: Test image no. 1. The Uffington White Horse, subtracted from an aerial photograph.




Figure 2.3: A detail of test image no. 2, showing the chequer-board mixing pattern.




Figure 2.2: Test image no. 2. The "Clonehenge" artificial henge model, mixed with an inverted aerial photograph of the Maes Howe henge.




Figure 2.4: The histogram of test image no. 2, showing quantisation at modulus 16.




Figure 2.5: Test image no. 3. A photograph of King's Low round barrow, with a picture of Tutankhamun's mask superimposed.




Figure 2.6: The histogram of test image no. 3, showing two distinct intensity ranges.




Figure 2.7: A detail of test image no. 3, showing the random noise introduced in the Tutankhamun area.

one image, together with only a very few stray ones from the second.

Having used a uniform pattern as a mask, we examined the possibility of a random distribution. We also followed the same idea as before of using grey-levels as a secondary means of segmentation. This time we chose to have a relatively unknown image together with a very obvious one (Fig. 2.5). The base image was fairly dark: it is a picture of the top of a cross, wedged upright in the earth and leaning on the stump from which it was broken. The site is King's Low round barrow, at Stafford, and the cross is said to have been taken and erected there from a medieval site in South Wales during the 17th century. The object itself is, incidentally, missing, and this photograph is the only record we now have.

The grey-levels of this image range only from 0 to 127, but those of Tutankhamun's face were made to range upwards from 128. Observe in the histogram that there is an obvious break at the middle of the range; the imbalance is because Tutankhamun's image is only a small proportion of the picture (Fig. 2.6). Using Tutankhamun as a mask, random dots were produced so as to have roughly equal numbers of white and black dots within the central area. These, then, were to determine the allocation of pixels for the final image. Tutankhamun's right eye is shown in close-up in Fig. 2.7. Extraction of these two image was considered to be a good test of processing capability: having separated them on their intensities, some sort of filtering and enhancement would be necessary to fill the randomly spaced empty dots.

We must congratulate all those who entered the competition, and who so enthusiastically went to work on extracting the images. The discussions which arose from the competition, the large attendance during the prize-giving, and the detailed, but spontaneous, impromptu session in which the winners discussed their techniques, showed clearly how much interest archaeologists have in image processing.

Only one entry managed to deal with all three files: the research students from Stafford were, of course, able to read our minds, as well as use the same software we used to produce the originals. They managed complete segmentation, without any prompting from us, and thus deserve praise (though no prize) for this achievement. With other entrants, the ingenuity of the methods of extraction were a complete surprise to us. The following sections describe how the test images were interpreted and processed by the prize-winning participants.

## 2.2. Image no. 1: the horse (Mike Canter)

The test images were received as PCX files, and when viewed they appeared to be the result of photographic double exposure. The truth, however, was even worse, they had been "Staffordised"! (or were they "Burslemed"?). The accompanying letter included the phrase "Enjoy yourselves." This was assumed to be the main objective of the exercise, and the images were processed accordingly. The work proceeded along the lines of artistic creation rather than scientific analysis. Many alternatives were explored, and no notes were taken! What follows is therefore based upon recollections of the steps retained.

Two suites of programs were available:

1. Graphic Workshop
2. IDRISI GIS

PCX files are accepted by almost all raster based graphics programs. Unfortunately the IDRISI GIS is one of the exceptions to this rule. The images were therefore converted into TIFF format using Graphic Workshop software. Examination of the file description, using TIFIDRIS, indicated an image of 480 rows by 640 columns and a header file of 1568 bytes. The image was imported into IDRISI GIS using its



Figure 2.8: Image no. 1 pre-processed by windowing and "hill shading" in IDRISI.




Figure 2.9: Removal of background interference by grey-scale reclassification.

PARE routine to remove the header file and set out the image coordinates and reference system.

There are however several dozen variations on the TIFF format. The result is that TIFF files output from one program are not necessarily the same as TIFF files that can be




Figure 2.10: The Uffington White Horse cloned to produce a racecourse effect.

read by another specific program. And so it came to pass that IDRISI accepted part of the data that was in the original PCX files. (The results of this process do not seem to be exactly repeatable. They probably depend upon the phase of the moon and the incantations used by the computer operator!)

The resulting IDRISI .IMG file was treated as a 256 level grey-scale image. By looking at 16 levels at a time, the levels of grey-scale containing the information to be selected were identified. In the case of the "horse" image, these were classes 70 to 120. The horse was extracted as a window to reduce the time taken by some of the following techniques.

The horse, oriented SW-NE, was obscured by plough-lines running NW-SE. The IDRISI SURFACE routine was used to "hill-shade" the image, using a sun direction parallel to the plough lines. This reduced the contrast in the plough marks and increased that of the outline of the Horse (Fig. 2.8). The IDRISI CONVERT and SCALAR functions brought the grey-scale values back to positive integers so that they could be viewed in detail. The Horse pixels were identified as grey-scale levels zero to 19 (dark) and 30 upwards (bright). The intermediary levels 20 to 29 formed a strong lumpy background that interfered with the "horse". The histogram for the image was then RECLASSed. Levels 0-19 were reclassified into class 10 and levels 30 upwards into class 14. All other pixels were classified as 5. This removed much of the background interference and produced a Horse outline in two strong tones (Fig. 2.9). It had the added advantage that the image could be viewed on a 16-level screen and avoided the distortions arising from autoscaling.

For any particular level of grey the RGB values are equal. Using the COLOR/K mode of IDRISI the balance of Red, Green, and Blue was adjusted to show a coloured horse on a green background. A degree of edge enhancement was then achieved by plotting the grey-scale levels in ORTHO mode

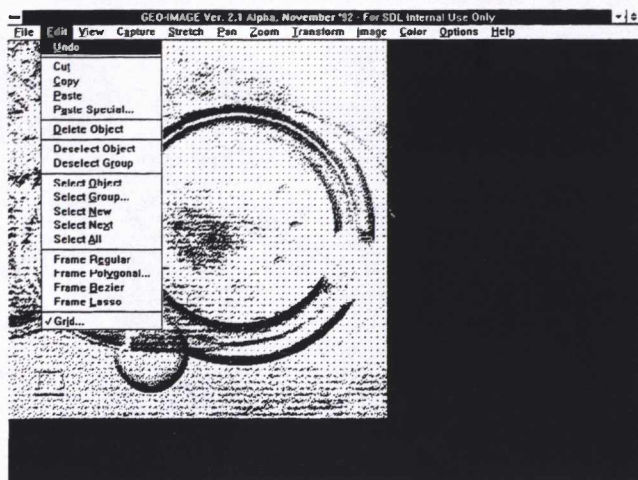


Figure 2.11: The subset of "interesting" features in image no. 2, created in GEO-IMAGE.

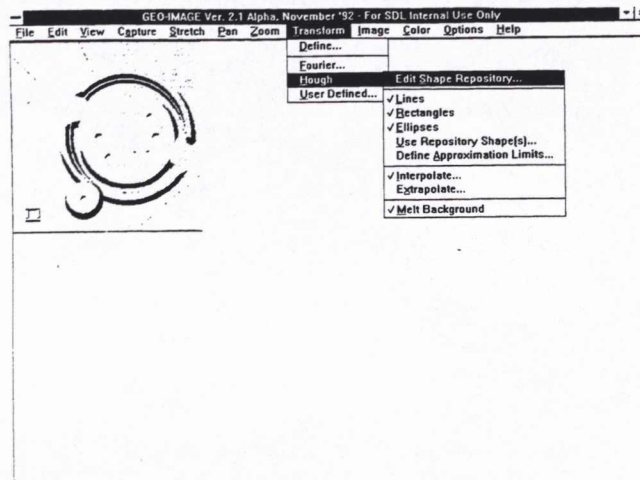


Figure 2.12: The process of two- and three-dimensional morphological object selection in GEO-IMAGE.

and draping the coloured image over it. The result was a slightly three-dimensional image that was coloured. The "racecourse" effect was produced by using the IDRISI CONCAT routine to paste several of the above images together in a staggered form (Fig. 2.10, reproduced here as a grey-scale image).

It was at about this point that the machine suffered from hard-disk indigestion. A combination of file deletion and run-length encoding produced enough space for completion of the exercises. A subsequent sweeping elimination of all intermediate files made it impossible to retrace the steps taken. Those wishing to undertake further research are warned that the White Horse at Uffington is to be found in the National Monuments Record and not in the Real Ale Guide!

### 2.3. Image no. 2: mixed henges (Stephen Bullas)

In common with all three images, image no. 2 was supplied in PCX format on a 3½" floppy disk. Deliberately, there was no supporting documentation, no clues as to what might be contained in the file, no count of the number of objects sought and no indication of how they were created or might be related (see Fig. 2.2). From an initial inspection of the file using Microsoft® Paintbrush, it seemed that the image contained a number of 2- and 3-dimensional objects ("artefacts") overlaying a variety of landscape "features". At this stage, it was not clear precisely how many artefacts were present and which of these were "solid" and which were "drawn". Similarly, although the outline of a large elliptical feature, possibly a prehistoric hillfort, together with a variety of circular, rhomboidal and linear features could be detected by the naked eye, it was a simple matter to start seeing pseudo-objects in the on-screen picture which, in reality, might not exist. Finally, it was not easy to ascertain whether

the background landscape was an aerial photograph, a computer-enhanced geophysical image or a hand-drawn picture.

It was decided, therefore, that the first task would be to separate out the "artefacts" from the "features" and to place the results into two new picture files ready for the second task of visual analysis and mathematical enhancement. The software chosen to accomplish both these goals was GEO-IMAGE, an SDL proprietary product used internally by the company as an image-recognition and enhancement tool in the provision of consultancy services to government and industry. The hardware used was a 25MHz Intel® 486 SX-based personal computer with maths co-processor, 8Mb of RAM and a high-speed, large-capacity disk, a 1280 x 1024 colour monitor with graphics accelerator and a mouse. The operating system environment was MS-DOS® Version 5.0 running Windows™ Version 3.1 in Enhanced Mode. Microsoft® Paintbrush was used both for the initial inspection of the image and for the saving of the output as PCX files. Pictures were imported and exported between the two products using the standard Windows™ Clipboard.

#### 2.3.1. Isolating the artefacts

The Bezier Curve and Lasso Framing functions of GEO-IMAGE were used to define and encompass irregular-shaped areas of potential interest as subsets of the total image. By using subsets rather than the whole of the image, specific objects may be more clearly identified and the processing time required during subsequent analysis can be significantly reduced. The subsets were then "associated" with one another using the GEO-IMAGE Edit facility, the resultant Group including both artefacts and features (Fig. 2.11).

A series of Hough Transforms were then applied to the image. In this type of transform, an image is subjected to comparison with the dimensions of a known object. In its simplest form, only two-dimensional shapes are considered. Although this can demand considerable processing power in its own right, GEO-IMAGE also possesses the capability of

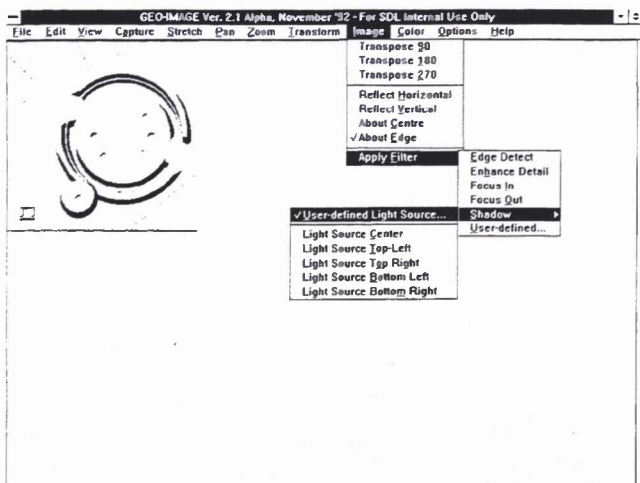


Figure 2.13: Shadowing of artefacts by GEO-IMAGE.




Figure 2.14: The results of artefact processing.

attempting to identify three-dimensional objects by morphological comparison. However, instructing it to do so can result in an exponential increase in the processing time taken. Fig. 2.12 shows the selection process required to achieve this result. All objects having linear, rectangular or elliptical shapes were sought. An approximation limit was specified such that the definition of these shapes had to be good, but not perfect. In order to help identify potential objects, the program was allowed to interpolate "missing" sections within small, predefined limits. Although not realised at the time, this feature of GEO-IMAGE automatically overcame the fact that the original image had been "chequer-boarded"; the small limit specified helped ensure that the large artefact did not become "joined" into a complete circle at its four ends. Once a shape was identifiable to the program, it was instructed to attempt to mould it into its 3-dimensional counterpart using surrounding image areas as its components.

To finalise the artefact processing, all parts of it which were unrecognisable as regular shapes (i.e. by using the default repository of lines, rectangles and ellipses only) were melded into a neutral background; a shadowing effect was then introduced using GEO-IMAGE's variable Light-Source feature (see Fig. 2.13). Processing of this image took about two hours to identify four drawn (i.e. 2-dimensional) and eleven solid (i.e. 3-dimensional) objects, including seven of the eight hemispherical indentations within the main circles. Much of the time taken could be attributed to the iterative nature of the calculations rather than to the referencing of the look-up tables contained within the shape repository. Since none of the artefacts was individually selectable by the user (no documentation accompanied the original image), the software was forced to examine the decomposed data, which form the image, many millions of times. Increasing the machine speed to 66 MHz, reducing the size of the subset area and the number of different shape-types being sought, limiting the comparison to 2-dimensional objects only and excluding the Interpolation process could all contribute to the probable reduction of processing-time to a few minutes.

### 2.3.2. Examining the landscape features

Largely due to these timing considerations, it was decided that the original image as supplied, including the landscape, its features and the overlaid artefacts, being a much larger area and including many varied shapes (for example, at least one immediately-visible rhomboidal feature rather than plain rectangular ones) in addition to incorporating the whole subset discussed in the previous section, would be processed omitting most of the advanced functionality. Accordingly, the Hough Transform recognition (and, by default therefore, the interpolation facility) was excluded. However, an edge detecting filter was applied using a selection process similar to that shown in Fig. 2.13. The advantage of this method is that the boundaries of objects become enhanced and the areas therein contained become suppressed. However, without using the Hough Transform method, not only the features, but also the artefacts will remain in the image and will also have their edges emphasised. A much better result could, in theory, have been achieved by combining both methods but the probable processing time which this would have involved rendered it an impracticable alternative.

The final results of the processing are shown in Fig. 2.14. The most obvious "artefact" is a large plate-shaped "solid" object containing a number of small depressions between its rim and centre; the circle is broken in two places (giving the four "ends" referred to above). The whole object is approximately symmetrical. To the bottom-left of this object is a whorl-shaped solid and to its left are four "joined" solid blocks forming a rectangle. All these objects appeared to have a 3-dimensional morphology whereas the following objects are two-dimensional and were probably drawn by machine or hand:

- Between the centre of the plate-shaped solid and its top rim are a number of drawn dots (shown white in the diagram). It was subsequently learned that these were placed here in order to ensure that the results




Figure 2.15: The results of processing on the landscape features.

were not simply copied from “Clonehenge” — the origin of this particular overlay (see below).

- To the bottom-right of the picture there is what appears to be a parallel trackway and another at right-angles to it at the top-middle of the picture.
- At the top-left of the picture is the edge of a feature which could form a rectangular or rhomboidal corner similar to that of a vallum of a Roman fort; this is more clearly shown in Fig. 2.15.

The large oval feature — a hillfort or henge — clearly shows entrances at the bottom-left together with an probable associated trackway. The rhomboidal feature is shown top-left and, mostly obscured, is the faint trace of the bottom-right trackway mentioned above. The analysis, however, could find no indication that any part of the landscape had been “inverted”.

#### 2.4. Image no. 3: King Tut’s curse (Martijn van Leusen)

Jan Hartmann and I, both working on the SURF project (diffusion of image processing techniques to the alpha sciences) at the Faculty of Environmental Sciences, Institute of Pre- and Protohistoric Archaeology, University of Amsterdam, used the open image processing software Khoros Version 1.5, developed and distributed by the Department of Electrical Engineering and Remote Sensing of the University of New Mexico, to have a go at the test images. The software was run on an IBM RS6000 320H workstation under AIX 3.2.1. The steps that we ultimately came up with to process the third image are represented as “glyphs” in the Graphical User Interface (named cantata) of Khoros, shown here in Fig. 2.16. Some pre-processing was necessary to make the image available in Khoros’s VIFF format. This included a PCX

to PPM conversion using the PBMplus package, and a pbm2viff conversion in Khoros.

Image no. 3, showing a bright and speckled mask of Tutankhamun on a dark background of what seemed to be a forest graveyard, immediately led us to look for a separation on the basis of non-overlapping grey-scale values. A grey-scale value that resulted in a maximum separation of the constituent images (111) was interactively determined by reading off the grey-scale values of the brightest pixels in the image background. As it turned out, this was not the correct threshold value, which had been set at 128 by the authors of the competition. Had we known then how to display the image histogram, the idea of grey-scale separation would have been confirmed in a more rigorous manner. However, the software we used was relatively new to us at the time and we had not yet determined how to use all the routines.

This first processing step resulted in two images, one a foreground image (containing pixel values 112 to 255) showing the bright but speckled mask of King Tut on a black background (Fig. 2.17b), the other a background image (containing pixel values 0 to 111) showing the dark “graveyard” with black holes (zero values) obscuring a central “tombstone” where King Tut had been removed (Fig. 2.17a). Metaphorically speaking, the removal of King Tut’s mask had resulted in the curse of the black holes on the tombstone — clearly a reference to the imminent death of the desecrating archaeologist. How could we fill these holes and escape the curse?

Hole-filling is a popular pastime among the image processing community, witness the large number of interpolation methods offered by most software packages. A number of these methods were considered by us and rejected because they would not fill the holes intelligently enough and would necessitate further processing to undo effects elsewhere in the image. Finally, a contributed tool for grey-scale morphological operations provided the means for an intelligent and non-messy closing of the holes. Morphological closing (the “vgmclose” glyph in Fig. 2.16) is a process that fills holes (represented by zero pixel values) in the image by continuously extrapolating new pixel values from the pixel values found at the hole’s edges by applying a ball structuring element created by the “vgballse” tool. Although the finer points of the algorithm used escaped us at the time, the method succeeded in filling the holes (thus hopefully saving us from the curse) and left us to perform a simple contrast enhancing operation (histogram stretching, “vhstr”) to obtain the final background image (Fig. 2.18). The algorithms used for the morphological closing and opening operations are part of the code of Khoros IP5. They are described by Gonzalez & Woods (1992). Our processing thus led to satisfactory reconstruction of what turned out to be an image of a man’s bust on a stone cross.

Having returned to his favour, we also tried to reduce the speckly appearance of the image of King Tut’s mask by applying a morphological opening filter with the same structuring element (“vgmopen”). Again this operation was successful, but it resulted in a less visually appealing image, so we decided not to submit it to the competition (Fig. 2.19). This result was probably due to our using a diameter of 5

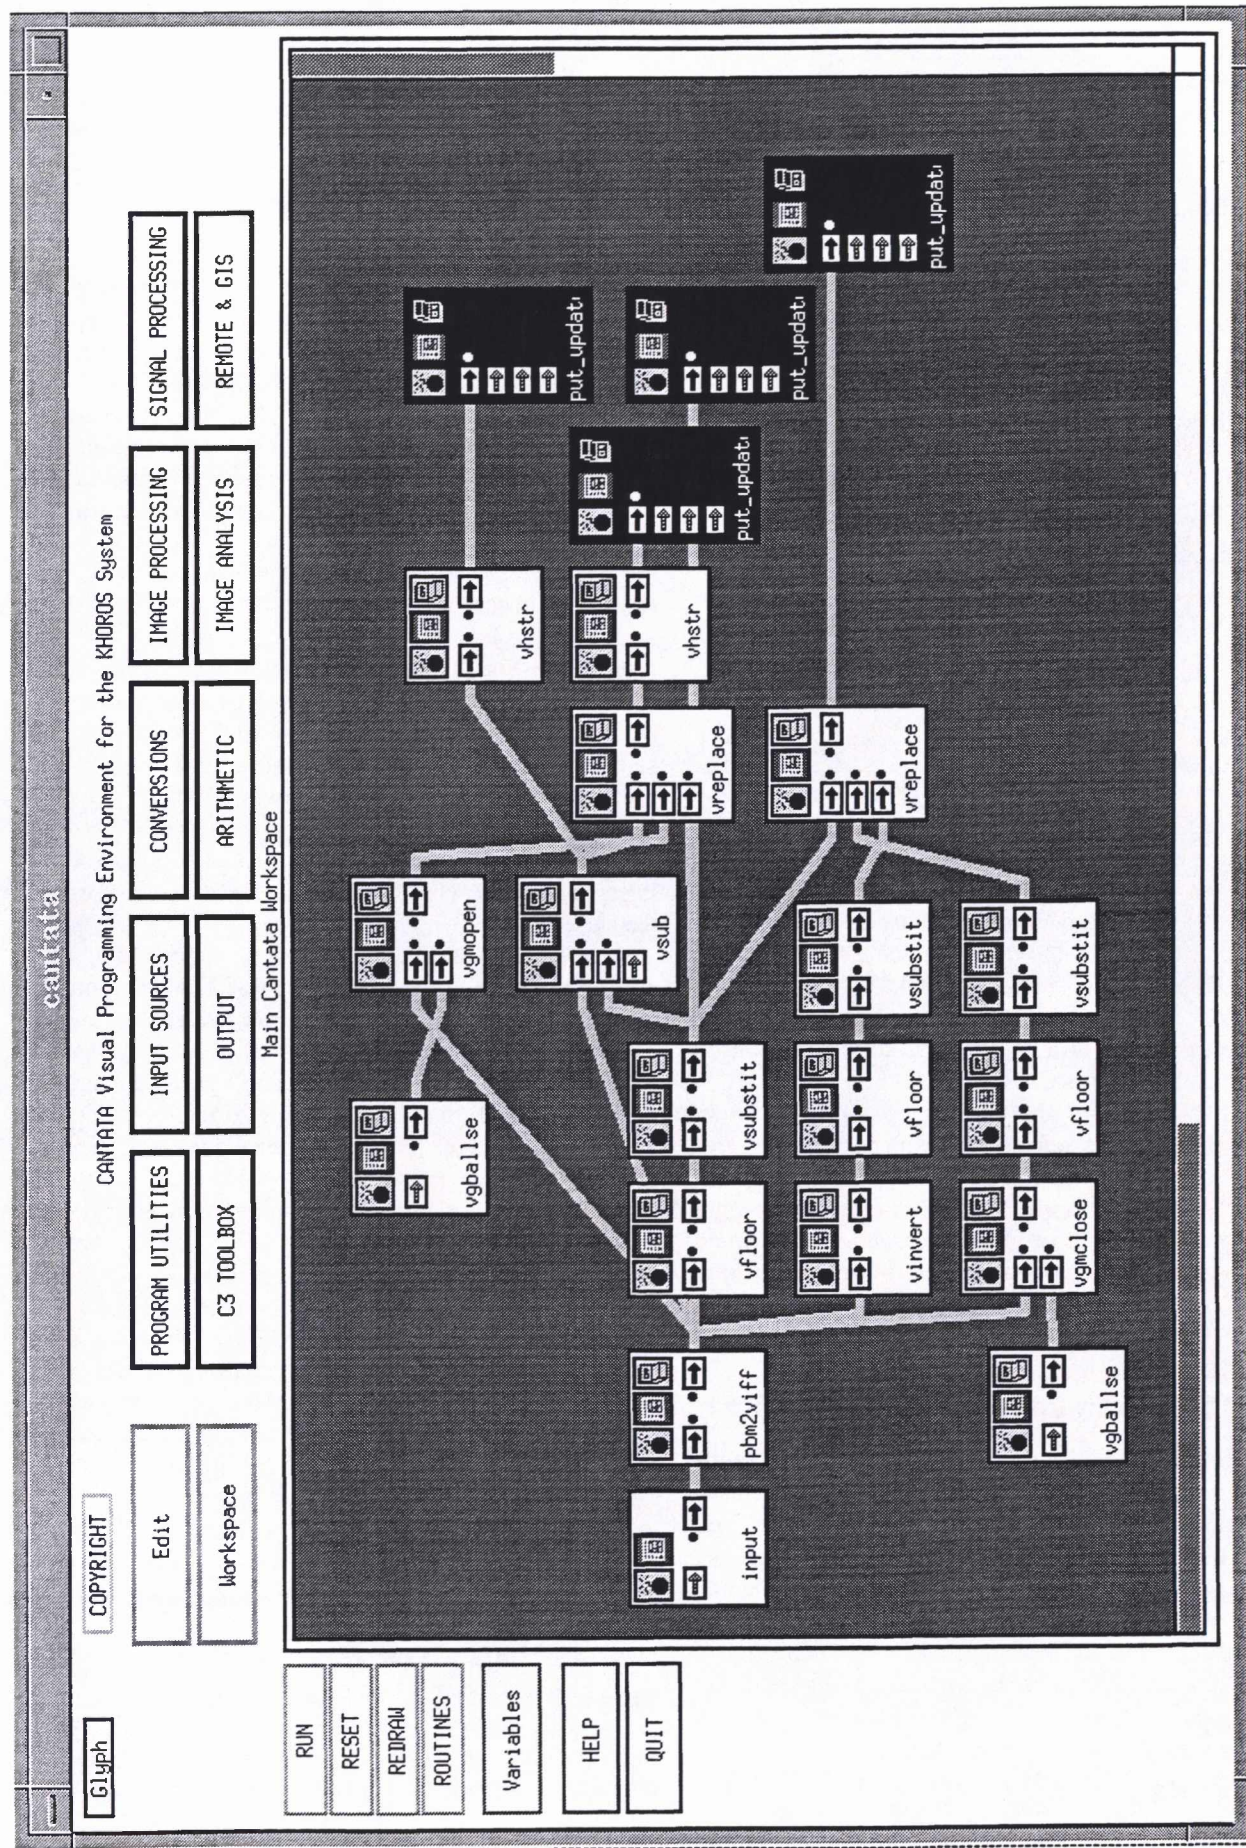


Figure 2.16: The "cantata" Graphical User Interface of Khoros, showing the tools used to convert, process, and display image no. 3.





Figure 2.17a: Background image, separated at grey level 111, with contrast stretched for clarity.



Figure 2.17b: Foreground image, separated at grey level 111.

pixels for the ball structuring element. A smaller diameter might have resulted in a less blurry image. High quality prints of the results were obtained by loading the images in *Xview*, another public domain package (Bradley 1989) used for colour editing and reformatting, and printing them on a 300 dpi Tektronix Phaser II thermal wax printer.

## 2.5 Concluding remarks

The happy winners of the competition are portrayed in Fig. 2.20. Some of the results of the competition were quite surprising. For one thing, no relation was apparent between the contestants' success in processing one image, and their processing another image. The competition's winners, i.e. those persons most successful in processing one of the images, each failed miserably in segmenting the other two images. For another, the processing of the test images was steered by the contestant's interpretation of their contents, leading to "blind spots" on the one hand and to "artefacts" on the other. The aims of the competition received different



Figure 2.18: Morphologically restored and histogram equalised background image.




Figure 2.19: Morphologically "restored" foreground image.

interpretations accordingly — from the focused enhancement of "interesting" parts of a test image by Mike Canter to the complete separation and restoration of the constituent pictures of an image by Martijn van Leusen. Third and lastly, the availability and quality of the hardware and software used



Figure 2.20: The winners of the image processing competition, left to right: Martijn van Leusen, Stephen Bullas and Mike Canter, with their posters and prizes.

in image processing clearly influenced the processing paths taken by the contestants.

### 2.5.1. Interpretations

As in serious archaeological image processing, the contestant's interpretation of the test images determines the direction of subsequent processing steps. Obviously such interpretation is necessary to extract whatever information one is interested in — concentrating on the horse while disregarding the furrows in image no. 1 — but there is a real danger of misinterpretation. Thus one may become waylaid by a failure to recognise what is in the picture (e.g. the fact that Maes Howe had been turned upside down, and the fact that the “background” to Tutankhamun contains interesting features) or by focusing on irrelevant features (the “bent man reading a book” seen by one contestant in image no. 1). These interpretations determined the goals set by the contestants, the methods used to enhance the images, and ultimately the results of their efforts. A second suite of interpretative errors can be caused by the processing itself. Many processing steps introduce spurious patterns in an image, and inexperienced users of image processing software are prone to interpret these artefacts as real (if enhanced) information. Some of the authors' comments on these subjects are listed below.

- Subsequent to the submission of results to the competition's organisers, Stephen Bullas' attention was drawn to an article on the mathematical model “Clonehenge” (Fletcher & Spicer 1988). The article describes how “three-dimensional” pseudo-objects may be mathematically created to form the original image no. 2. A comparison of the original objects' descriptions against the results given in the present paper emphasises quite clearly just how subjective the process of visual interpretation of “solid” objects can be. In particular, the whorl-shaped artefact and its central hole had originally been intended to represent a robbed-out barrow, and the large “plate” to represent the earthworks of a henge. Had an indica-

tion of scale been supplied with the test images, then the mistakes in subsequent human interpretation would almost certainly not have been made. The results of the two-dimensional analysis — in particular, the Roman ditch — were more precise, but, again, walls and trackways were easily confused with one another.

- Martijn van Leusen notes that he and his colleague Jan Hartmann processed the images “by ear”, displaying them on the computer screen and looking for features that would allow segmentation or enhancement. Such a strategy makes little or no use of the image information that the computer can provide (of which the grey scale histogram is the most important), and depends instead on the limited range of grey values and amount of detail distinguished by the human eye. Had we started our processing by displaying the image histograms, we would have understood how test images no. 2 and no. 3 were created, and this would have given us a handle for further processing. Had we zoomed in on the images until we could actually see all the information (pixels) in them, we might have recognised the chequer-board nature of image no. 2. Of course, these test images were composed artificially, and “real” archaeological imagery cannot be enhanced so easily (as illustrated by image no. 1). However, our assumption that we would be able to see whatever we needed to see without consulting image statistics is clearly unwarranted even with “real” imagery, and we hope to better our lives accordingly.

### 2.5.2. Influence of the software and hardware used

From the descriptions of the processing above it is already clear that the contestants used widely different sets of processing software and hardware, ranging from simple PC based systems with limited memory and processing capacity and 16-colour screens running software not especially written for image processing tasks to workstations with large 24-bit screens running special software. Processing time ranged from hours to minutes. Printing hardware ranged from the lowly dot matrix printer to a 300dpi thermal colour wax printer. There are obvious limits to what the simpler combinations of hardware and software can do — but there are also more subtle influences:

- The limited capability of humans to see grey level differences is paralleled by the limited capability of the computer screens to display these grey levels. In an interesting quirk of fate, his system being limited to 16 colours, Mike Canter was forced to look at all 256 levels in 16-level clumps. Therefore he actually saw more information than Martijn van Leusen and Jan Hartmann who were able to display (but not distinguish) all grey levels at once.
- Stephen Bullas noted that there are major physical differences between images printed on paper and the same images displayed on a computer screen. In general, printed output of grey-scale images is created by a process known as “dithering”. This process causes the various shades between absolute black and absolute white to be represented by a series of very fine dots. The human eye is then misled into interpreting these dots as shades of grey. However, as each printer hardware manufac-

turer uses a different set of algorithms to position the dots, the same image printed on differing manufacturers' machines will tend to have different parts of the image slightly enhanced or suppressed. The prints will often suffer from "banding" — artificial linear marks introduced during the printing process. The advent of Microsoft® Windows™ has gone some way in standardising this situation but has not fully cured the problem, since its printer software drivers are still primarily sourced from the individual hardware manufacturers. A second factor leading to inconsistency of the printed image is the hardware "resolution" of the printer. Most laser printers, such as the Hewlett Packard® LaserJet® series (and its compatibles), have traditionally operated at a resolution of 300 dpi (dots per inch) although the LaserJet® 4, which operates at 600 dpi, is now becoming more common. Operating at one or other of these resolutions has important impact, viz.: no matter how much better the resolution of the final publication may be, it will nevertheless still be confined by the maximum of 300 or 600 dpi of the original. On the other hand, the image seen on a modern graphics screen, especially on SuperVGA and beyond, will normally have completely different resolution to that of the printer and will also be more naturally displayed (i.e. without dithering), having at least 16 "natural" grey-scales, sometimes substantially more. The combined effect of the above considerations is, therefore, that the pictures referred to in this paper can at best be only poor approximations of the view that the user sees on the screen.

### Bibliography

- BRADLEY, J. 1989. *Xv Version 2.21*. University of Pennsylvania.
- EASTMAN, J. R. 1992. *IDRISI manual*. Clark University, Massachusetts.
- FLETCHER, M. & D. SPICER 1988. "Clonehenge C-Coded!", *Archaeological Computing Newsletter* 16: 4–10.
- GONZALEZ, R. C. & R. E. WOODS 1992. *Digital image processing*. Addison Wesley.
- RIMMER, S. 1992. *The Graphic File Toolkit*. Addison Wesley.

Stephen G. Bullas  
SDL -Strategic Decisions Ltd.  
Geophysical Software Division  
Pilgrim House  
The Sycamores  
Fetcham  
GB-KT22 9EX Leatherhead, Surrey  
CompuServe 100111,2665

Michael J. T. Canter  
Brinna House  
GB-ZE2 9JF Bigton, Shetland

Michael Fletcher  
Staffordshire University  
School of Computing  
Leek Road  
GB-ST4 2DF Stoke-on-Trent  
cmtrmf@staffs.ac.uk

P. Martijn van Leusen  
Instituut voor Pre- en Protohistorische Archeologie  
Nieuwe Prinsengracht 130  
NL-1018 VZ Amsterdam  
martijn@scanner.frw.uva.nl

Richard D. Spicer  
Staffordshire University  
School of Computing  
The Octagon  
Beaconside  
GB-ST18 0AD Stafford  
cmtrds@staffs.ac.uk