PROJECT REPORT No. 262

DESIGN OF AN INTEGRATED MACHINE VISION SYSTEM CAPABLE OF DETECTING HIDDEN INFESTATION IN WHEAT GRAINS

OCTOBER 2001

Price £4.00

PROJECT REPORT No. 262

DESIGN OF AN INTEGRATED MACHINE VISION SYSTEM CAPABLE OF DETECTING HIDDEN INFESTATION IN WHEAT GRAINS

by

J CHAMBERS[1], C RIDGWAY[1] AND E R DAVIES[2]

[1]Central Science Laboratory, Sand Hutton, York YO41 1LZ

[2]Machine Vision Group, Department of Physics, Royal Holloway, University of London, Egham, Surrey TW20 0EX

This is the final report of a project that started in April 1998 and ended on 31 March 2000. The work was funded by a grant of £70,200 from the HGCA (project no. 1437).

Contents

ABSTRACT

Major advances have been made towards the World's first system for the rapid, automated and non-destructive determination of internal infestation within kernels of post-harvest cereals. The system will be based on the analysis of images of the kernels recorded in the near infrared (NIR) region by machine vision using highly innovative computational algorithms. These algorithms are notably more effective at analysing the images than human inspection as well as being non-subjective and operating at constant efficiency. Laboratory tests suggest that the probability of finding a 0.5% infestation of grain weevils in a batch of cereal could be around 97%. The speed of processing the images has been markedly improved such that analysis of a 3 kg sample within 3 minutes should be possible. The discovery that the method works with images recorded in the very near infrared region, at 981 nm, means that it will be possible to use a cheaper camera which is less subject to drift than had originally been thought. The method works by detecting bright patches on infested kernels which are highly correlated with the present or past location of the insect and probably result from loss of starch due to insect feeding. Repeated scanning of infested kernels suggests that the NIR effect becomes detectable around 2-3 weeks after egg laying. This is only shortly after the infestation becomes detectable by the much more expensive X-ray method. The NIR method has been optimised using wheat of Mercia variety but preliminary tests with other varieties of different reflective appearance suggest that it has general applicability to wheat. The method could share the same sample presentation arrangement as the machine vision system being developed to detect contamination external to cereal kernels but the cameras and software analysis will need to remain separate. Now that these laboratory tests have proved the potential of the method, the next steps will be to confirm that it will work with cereals other than wheat and infestations other than the grain weevil, and to construct a prototype apparatus to test under practical conditions. The implications of this project are that, provided the advances described here are exploited, the UK cereal trade will be able to benefit, in particular for both long term storage and export, by being the first to have at its disposal a system which can confirm the high quality of UK grain and which will operate at an acceptable speed and provide efficient results at an affordable cost.

1. INTRODUCTION

Previous collaborative work between CSL and RHUL, funded by the HGCA, has resulted in significant progress towards the first ever method for the rapid, automated detection of the major foreign object contaminants of cereals (Chambers *et al.*, 1997a; 1997b; 1998). The developed method is presently at laboratory stage. Images of samples of wheat are recorded in the visible region using a low-cost charge-coupled device (CCD) camera and lighting system. A package of highly innovative purpose-written computer algorithms automatically detects rodent droppings, ergot, adult and larval insects external to grain kernels (Davies *et al.*, 1998a; 1998b; 1999). This laboratory model provides the proof of concept and detailed specification from which to construct a completely integrated prototype device. It is expected that this device will be capable of scanning and classifying the contents of a 3 kg batch of grain in a total time of just 3 minutes. No existing device has a performance comparable to that expected from the system to be developed from this work. However, the detection of hidden infestation, *i.e.* kernels infested internally with insects, is currently *not* possible with this system. This would be particularly valuable for long term storage and for export, where it would enhance UK competitiveness in international trade.

Since the first report of the ability to detect internal infestation using near-infrared (NIR) spectroscopy (Chambers *et al.*, 1993), which formed part of previous work funded by the HGCA (Chambers *et al.*, 1994), work has been directed to establishing how this potentially valuable discovery could be exploited for practical benefit. Detailed understanding of the spectroscopic response (Ridgway and Chambers, 1996) led to the significant advance that internal infestation could be detected by measurement at just two wavelengths rather than a much wider spectral region (Chambers and Ridgway, 1996). This suggested that detection might be possible by NIR imaging, which was demonstrated as part of the work funded by HGCA on the detection of contaminants (Chambers *et al.*, 1998).

Wheat kernels containing larvae of the grain weevil, *Sitophilus granarius* (L.), were discovered to exhibit large and very distinct light patches when imaged at certain carefully selected NIR wavelengths. Uninfested kernels present in the same image appeared uniformly dark. The contrast between infested and clean kernels was greatest in the composite image (1202 nm - 1300 nm), obtained from separate images of the sample at each wavelength. However, the use of images at just one wavelength (1202 nm) gave an almost equally good result. Limiting the number of wavelengths to just one promises to afford a very simple and rapid NIR imaging methodology, by avoiding the need to switch between filters and subtract one image from another. This preliminary study (Ridgway and Chambers, 1998) confirmed NIR imaging as a potential basis for a simple, safe, rapid and inexpensive substitute for more costly detection by X-ray. The purpose of the present project was to explore this potential.

2. OBJECTIVES

The objectives of the present project were to:

a) ensure that NIR images can be analysed automatically by substantially improving their quality,
b) confirm that there is sufficient information in the NIR images for machine vision by using human inspection of unknown uninfested and infested kernels,
c) establish how to recognise images of hidden infestation automatically by constructing computer algorithms,
d) optimise the sensitivity of detection by confirming the origins of the differences in response from clean and infested kernels, and
e) determine how to incorporate the detection of hidden infestation facility into the existing machine vision system.

3. GENERAL

Except where stated otherwise, samples were prepared, images recorded and the quality of the images assessed at CSL, while image analysis was undertaken at RHUL. Image analysis uses a number of specific terms. A glossary of the terms of principal importance is provided as Appendix A. Two candidate wavelengths for insect detection were compared in this study. The wavelength 1202 nm was that used in previous, preliminary work. The alternative wavelength 981 nm was identified during the present study, as a result of the very near-infrared spectroscopic investigation (see Section 7).

Materials used in the images

Insects used were grain weevils, *Sitophilus granarius* (L.). Wheat used for culturing was variety Mercia. Infested wheat kernels contained large, probably final instar, larvae. These were identified to developmental stage by visual inspection of X-ray radiographs. Uninfested kernels for use as control samples were also obtained by this method at the same time and from the same culture. All kernels imaged were thus of variety Mercia, except where stated otherwise [for part of Experiment (g)].

4. IMPROVEMENTS TO QUALITY OF NIR IMAGES

Preliminary work to optimise image capture methodology at 1202nm

Various different lighting chambers, lighting arrangements and cameras were investigated in order to optimise the contrast between uninfested and infested kernels and minimise image noise. Diffuse lighting chambers consisting of a horizontal or vertical cylinder were evaluated. The horizontal chamber design was modelled on that described by Paulsen and McClure (1986). Lighting intensities ranging from 200-800W total were investigated. The criteria by which lighting quality was assessed included minimising shadow, variability in light intensity across the image and heating of the sample and camera. It was found that substituting a lighting chamber of larger dimensions to that used in the previous project gave images of similar quality. This, together with the incorporation of a telephoto lens, helped reduce the problem of camera over-heating by allowing the lighting sources to be positioned away from the camera body (Figure 1).

Cameras tested at 1202 nm included a silicon-detector CCD responsive up to 1300 nm (Hamamatsu Photonics), an indium gallium arsenide-detector CCD responsive up to 1700 nm (SU128-1.1.7RT, Sensors Unlimited) and a NIR vidicon tube camera responsive up to 1800 nm (Find-R-Scope FJW85400A, FJW Optical Systems). Image quality was assessed by considering signal-to-noise ratio (the ability to distinguish features of interest from irrelevant artefacts), resolution (image clarity) and sample size. Attempts to improve significantly on the quality of the images obtained at 1202 nm in the previous project (Ridgway and Chambers, 1998) were unsuccessful. Use of the horizontal cylinder-based lighting chamber did not remove grain shadow and did not deliver sufficient incident light to the sample. The Hamamatsu Photonics silicon-detector CCD was found to be insufficiently sensitive at 1202 nm to allow images of grain to be recorded. The indium gallium arsenide-detector CCD, included as a model for linescan cameras, was found to provide insufficient detail in the captured image.

Images recorded

Once the preliminary work to select the most appropriate sample illumination and image capture methods had been completed, a series of images was recorded to allow preliminary investigation of the machine vision process, Experiment (a). Images were recorded at 1202 nm using the NIR vidicon tube camera (Find-R-Scope FJW85400A, FJW Optical Systems) at a resolution of 256 x 256 pixels. The series comprised 24 images obtained from nine sets of three infested kernels and nine sets of three uninfested kernels, each set being imaged separately on a light background and a dark background, with and without background subtraction. Summary details of each Experiment are presented in Table 1. An example set of images from one sample is given in Figure 2.

Image analysis by machine vision: early conclusions about the process

Investigation of the 24 images from Experiment (a) proved very formative for the Machine Vision part of the work. In these and certain later experiments, it proved useful to gather the images twice, once using light and once using dark backgrounds. While the light backgrounds are valuable when enhancing the images for human observation, the dark backgrounds are needed so that the boundaries of the grains can be readily determined automatically by computer. In what follows in the Machine Vision part of the report, it will always be assumed that the images used have dark backgrounds. These images showed that NIR can provide only a very low signal-to-noise ratio, and that this would make it quite difficult to find how to tackle the internal infestation problem, let alone to obtain reliable optimised information from the images. However, noise was not the only problem. The natural variability of the grains in brightness, size and shape also became evident. In addition, the variability in the effects of the infesting larvae showed that sensitivity in detection could be elusive unless some working model of the bright patches due to the larvae could be developed. In spite of these difficulties, the small number of images available at this stage permitted useful degrees of noise elimination, understanding of the grain intensity profiles, and grain normalisation by grain location, orientation and masking. It also allowed the development of a potentially useful technique for grain intensity modelling.

This early work highlighted the natural variability of the grains and the effects of the infesting larvae: at the time it seemed that the proper way to deal with this problem

was to make the Machine Vision inspection system highly adaptive, and to use training techniques to achieve this. In particular, classifiers would have to be trained to recognise a range of uninfested and infested intensity profiles, though it was also recognised that this would be particularly difficult for infested grains, which show much higher variability than uninfested grains. Hence it was anticipated that the classifiers would have to be trained on the uninfested grains, and the limits on the variability of the latter would have to be measured carefully: the infested grains could then be recognised as those falling outside the tolerance limits for uninfested grains. It was imagined that neural networks could be used for classification, though this possibility was not considered in any detail. Another possibility was the use of a procedure which would average the intensity profiles of many grains after rigorous normalisation for parameters such as position, orientation, size, shape and background brightness. Once an idealised model for uninfested grains had been obtained in this way, it should then be straightforward to determine whether any received grain image represented an infested grain. While these ideas were very tentatively framed during this early stage, the concept of grain modelling was retained throughout the project and indeed became of central importance to the final solution.

5. INFORMATION IN THE NIR IMAGES

While the preliminary investigation of the machine vision approach was being undertaken on the images from Experiment (a), a study was conducted to establish the extent to which internal infestation could be established from such images by human inspection alone. This served two purposes: first, to confirm that the images contained sufficient information for the machine vision process to be worthwhile, and second, to provide a measure of its success. By the time this study was undertaken, the evidence from the spectroscopic investigation (Section 7) was available. Since this suggested that it might be advantageous to use a lower wavelength (981 nm) than that used so far (1202 nm), the former was used for this investigation.

Preliminary work to optimise image capture methodology at 981nm
Two lighting chambers of different height and diameter, together with lighting intensities ranging from 160-400 W total, were investigated in order to optimise the contrast between uninfested and infested kernels. For imaging at 981 nm, consistent differences between uninfested and infested kernels were not observed using the large lighting chamber and telephoto lens. However, an image capture set-up almost identical to that used in the previous project was found to be effective (Figure 3).

Images recorded
Lighting was from four 40W standard pearl lightbulbs supplied by mains. Images were recorded at 981 nm using a silicon-detector CCD camera (Ikegami ICD-42E, Type F, ½ inch) at a resolution of 256 (height) x 521 (width) pixels. Materials used in the images were obtained in the same way as those used in Experiment (a). In the first part of this study, Experiment (b), each sample comprised one uninfested and one infested kernel, positioned such that they were horizontally aligned in the captured image. The positions of the infested and uninfested kernels were alternated between each sample. One image was recorded for each of 16 pairs of one uninfested and one infested kernel. An example image (contrast enhanced) obtained with the optimised set-up is given in Figure 4, for the first replicate sample of Experiment (b).

In the second part of this study, Experiment (c), the positions of the infested and uninfested kernels in each pair were randomised between each sample. One image was recorded for each of 51 pairs of one uninfested and one infested kernel.

Methods of image analysis by human inspection
The images obtained in Experiments (b) and (c) were analysed by human inspection. Images from Experiment (b) were examined by applying a binary threshold (Image Compact) to highlight any differences in brightness with infestation. In Experiment (c), the captured image was examined by a person different from the one who had randomised the positions of the kernels. The examiner was prevented from seeing the samples and did not know the relative positions of the uninfested and infested kernels under the camera. The examiner attempted to identify which of the two kernels was uninfested and which was infested on the basis of the NIR image, using a pc monitor. As an aid, images from Experiment (c) were enhanced by applying a logarithmic contrast enhancement (Image Compact). Contrast enhancement was optimised for the range of grey levels which occurred in the kernels, by first isolating the two kernels as the area of interest (AOI). This was achieved using the irregular AOI image thresholding option included in the software.

Kernels in Experiment (c) that were classified by inspection of the captured image as uninfested were placed in one container and kernels classified as infested were placed in a second container. At the end of the experiment, the contents of these two containers were inspected separately by X-ray, to determine the classification accuracy of human inspection of the NIR image.

Results of image analysis by human inspection
The thresholded images obtained from Experiment (b) are given in Figure 5. By qualitative inspection of the figure it can be seen that in the case of 11 out of the 16 samples imaged, the infested kernel displays a markedly larger area of light pixels than the uninfested kernel [Images (1)-(5), (7), (9)-(11), (13) and (15)]. In the remaining 5 images, uninfested and infested kernels appear similar.

In Experiment (c), of the 51 pairs of kernels, X-ray examination showed that in 40 the infested and uninfested kernels had been classified correctly. This classification accuracy of 78% confirms that the infested kernels appeared consistently different to the uninfested kernels when imaged at 981 nm.

6. RECOGNISING HIDDEN INFESTATION BY MACHINE VISION

On completion of the work to develop suitable image capture facilities at both 1202 and 981 nm, confirmed by the results from human inspection of certain of the images, the main investigation to assess the suitability of machine vision to detect hidden infestation was started.

Method of image capture at 1202nm
Monochrome digital images at 1202 nm, of resolution 256 x 256 pixels x 8 bit (256 grey levels), were captured using an NIR vidicon tube camera (Find-R-Scope FJW85400A, FJW Optical Systems; wavelength range to 1800 nm). This was connected to a personal computer (Scenic Pro C5, Siemens-Nixdorf; 133 MHz) fitted with framegrabber (PCImage-SG, Matrix Vision), framegrabber control software

(Image Compact version 3.0 for Windows) and image cropping software (Adobe PhotoShop version 5 for Windows). The camera was fitted with a 17.5-105 mm zoom lens (Fujinon TV lens) set at approximately 50 mm focal length. A close-up lens adapter (Close-up lens CL14058, Fujinon) was attached to the front of the lens. A spacer ring of depth 6 mm was positioned between the lens and the camera. A bandpass filter with central wavelength 1202 nm (bandwidth 10 nm at half height, 25 mm diameter, image quality; Andover Corporation) was fixed between the lens and the camera using the built-in filter holder at the back of the lens. The front of the lens was positioned approximately 30 cm above the sample. Lighting was from four 100W standard pearl lightbulbs supplied by mains. Lighting was reflected onto the sample *via* an aluminium foil-covered plastic cylinder (approximately 30.5 cm diameter x 36 cm height above sample) to give diffuse lighting. This cylinder incorporated a hole (10 cm diameter), with the bottom of the circumference level with the height of the sample. This allowed the sample to be placed in position under the camera. During image capture, the hole was blocked by attaching a sheet of card (11 x 11 cm) covered with aluminium foil. To improve signal-to-noise ratio, recorded images were the average of 100 frames. The image capture set-up is shown in Figure 1.

Images recorded at 1202 nm

Samples were presented to the camera as single kernels. Kernels were positioned with the crease away from camera and with the germ-end oriented downwards in the captured image.

For Experiment (d), the sample was placed on a 3 mm glass plate (15 x 15 cm) which was held 8 mm above a ceramic tile of the same dimensions by 4 spacers (Figure 1). The tile was covered with a sheet of glass fibre paper (Whatman GF/A). The spacers allowed the insertion of an infrared blocking filter (5 x 5 cm, KG1 heat absorbing filter, Optometrics) between the sample and the glass filter paper background. The backing material holding the sample was supported by a laboratory jack, adjusted to give the required distance between sample and camera. Two images (with and without infrared blocking filter as background) were recorded for each of 25 replicate samples of uninfested kernels and 25 replicate samples of infested kernels. Uninfested kernels were imaged alternately with infested kernels. One blank image was recorded for each of the two backgrounds (glass fibre paper or infrared blocking filter).

For Experiment (f), the sample was placed directly onto the infrared blocking filter. In this experiment, one image was recorded for each of 75 replicate samples of uninfested kernels and 75 replicate samples of infested kernels. Again, uninfested kernels were imaged alternately with infested kernels.

Method of image capture at 981 nm

Monochrome, 8 bit (256 grey levels), digital images at 981 nm were captured using a silicon-detector CCD camera (Ikegami ICD-42E, Type F, 1/2 inch) connected to the same framegrabber arrangement as for imaging at 1202 nm. Image resolution was 256 x 256 pixels. The camera was fitted with a 16 mm fixed lens (Computar TV lens, 1:1.4). A spacer ring of depth 8 mm was positioned between the lens and the camera. A bandpass filter with central wavelength 981 nm (bandwidth 9 nm at half height, 25 mm diameter, image quality; Andover Corporation) was fixed flush to the front of the

lens housing and edge-sealed (Blu-Tak). The front of the lens was positioned approximately 5 cm above the sample. Lighting was from four 40W d.c. bulbs fed by two stabilised d.c. supply units (TSV 30/5CL, Farnell Instruments) connected in parallel and each producing 25V. Lighting was reflected on to the sample *via* an aluminium foil-covered card cylinder (approximately 20 cm diameter x 20 cm height) to give diffuse lighting. To improve signal-to-noise ratio, recorded images were the average of 100 frames. The image capture set-up is shown in Figure 3.

Images recorded at 981 nm

Samples were presented to the camera as single kernels. Kernels were positioned with the crease away from camera and with the germ-end oriented downwards in the captured image. The sample was placed directly onto the infrared blocking filter.

For Experiment (e), one image was recorded for each of 75 replicate samples of uninfested kernels and 75 replicate samples of infested kernels. Uninfested kernels were imaged alternately with infested kernels.

For Experiment (g), wheat kernels of varieties other than Mercia were obtained from RHM Technology Limited and inspected by X-ray to confirm freedom from insects. One image was recorded for each of 20 replicate samples of uninfested Mercia, infested Mercia, uninfested Malacca (AU23), uninfested USA Northern Spring (AU20), uninfested Canadian, uninfested Shango and uninfested Australian White kernels. Images were recorded in rotation, by imaging one kernel of each sample type in turn.

For Experiment (h), one image was recorded for each of 20 replicate samples of uninfested kernels and 20 replicate samples of infested kernels. Uninfested kernels were imaged alternately with infested kernels. After imaging by NIR, a reference X-ray image was recorded for each infested kernel. Kernel positions were fixed by mounting the kernels on transparent adhesive plastic sheet supported by a wire frame. The resulting radiograph was digitised *via* a flatbed scanner (GT-8500, Epson). Kernel orientation in the X-ray image was approximately the same as in the NIR image.

Image analysis by machine vision: the main concept

As soon as the 50 images with the dark background from Experiment (d) became available, a large number of tests and experiments were carried out to ascertain how the principles of training, grain normalisation, grain modelling and bright patch detection could best be implemented. It rapidly became clear that the degree of variability of the grain images was so high that normalisation would be especially difficult to achieve a classification accuracy at least as good as the 78% achieved by human inspection [Experiment (c)] and with a similar consistency or better. The fault in the concept was that taking a large number of grain images and normalising them, and subsequently averaging them to determine the idealised grain model, was unachievable, as the high degree of variability meant that a large number of different idealised grain models would be needed. This could only be achieved by cluster analysis in a large parameter space, and indeed would only be successful if many hundreds and possibly thousands of grain images were available. The alternative solution was to derive an idealised model of each grain from the image of that grain itself. This approach would have the advantage that it would deal satisfactorily with

slightly different signals from other grains, however similar they might appear to be. On the other hand, this approach would not permit any averaging over grains, so the model could only be improved by adjusting internal algorithm parameters. Nevertheless, this approach did seem to have something to offer, if only the algorithm parameters could be adjusted suitably. Here, again, there would not be enough training data to institute an automatic means for optimising the parameters, so initially a pragmatic technique would have to be adopted. This pragmatic technique would not have to be totally trial-and-error, as observations could be made for all the grains, infested and uninfested, to train the eye to understand the situation, and thus guide the improvement process in a sensible way. Humans can be very incisive in reacting to small numbers of training set patterns, whereas machines are far better than humans at learning from large numbers of training set patterns. With this approach several important factors soon came to light:

1. It was consistently found that the rows of pixels in the grain images were alternately light and dark, as a result of an unknown type of acquisition problem which could not readily be overcome.
2. Searching for bright patches due to larval infestation proved to be an unreliable process around the grain boundaries. It is not entirely clear whether this was because the larvae are situated entirely within the grains, thereby preventing the bright patches from appearing at the boundaries, or whether the algorithms used to model the grain intensities are less accurate in those locations, thereby reducing sensitivity at these positions. The general conclusion is that both of these conditions applied to some extent.
3. The ends of the grains tend to be brighter than the bodies of the grains, and this meant that it was far less reliable to look for bright patches from larval infestation at the ends of the grains.
4. Modelling of the grain intensities could be achieved by applying a suitable noise averaging filter. While this would in principle be an inadequate technique, in practice a rather large filtering mask could be used for the purpose, so that not only was noise eliminated but also moderately fine grain intensity detail was cut out: in particular, the bright grain patches were largely eliminated.
5. Any bright patches on the grains would now be revealed by simple differencing against the model.
6. The brightest patch on any grain could be located by scanning over the active region of the grain (i.e. the parts encompassed by the exclusion zone defined below), the brightest patch being taken to be the part of the grain giving the most information on potential infestation – even if the grain is actually uninfested.
7. The average intensity of the brightest patch on a grain could be taken as the sole indicator of information on potential infestation, and all other information discarded.

In view of the relatively large resolution of the grain images, factor 1 was tackled quite simply: each image was sampled in alternate pixels in each direction, and this was carried out four times with horizontal and vertical displacements of 0,0; 0,1; 1,0; 1,1 – thereby obtaining four smaller images. Only one of these was used. It might be possible to use all four to reduce noise further, if higher levels of computation were

justifiable (though the occasion, and the need, never arose during the remainder of the project).

Because of factors 2 and 3, an 'exclusion zone' scheme was devised. First, a mask of constant width was engineered around, but within, the boundary of the grain (Figure 6). Second, the ends of the grain were excluded by further modifying the mask so that it would not extend outside a circular region centred at the centroid of the grain. The radius was determined as a factor beta times the radius of a circle of area equal to that of the grain being considered, beta being one of the parameters to be optimised for sensitivity (see below). (It turns out that the centroid location and area of a grain can be measured accurately and rapidly in these types of image.)

Items 4–7 lead to the main numerical information on each grain permitting it to be classified as infested or uninfested. This numerical information also permits the classifier to be trained for the purpose (Table 2).

Summary of results from Experiment (d)
The degree of variability of the grain images was so high that normalisation would be especially difficult to achieve to the required level of consistency and accuracy. Hence each grain had to be normalised individually to derive an idealised model of each grain from the image of that grain itself. In addition, each grain had to be masked to exclude the outermost parts which might be misleading for estimating infestation. After this had been achieved, each grain could be compared with its model, and the most significant bright region in the difference image located: its mean intensity could then be used as an indicator of infestation.

This basic schema was sufficiently powerful that it was retained in one form or another right up to the end of the project.

Image analysis by machine vision: training considerations
The maximum patch brightness value was the main piece of information that arose from every image. In principle, this only needs to be thresholded in order to determine whether a grain is infested or not. If the brightness is above a certain threshold, the grain is taken to be infested, and otherwise it is taken to be uninfested. In practice the level of threshold was adjusted by a recognition procedure which was adjusted for optimality on the given training set, and then the algorithm was ready for testing on new data.

In statistical pattern recognition applications, it is taken as axiomatic that the system's accuracy should not be tested on the training set, as the system is taken to be over-adapted to the training set, giving a higher apparent accuracy that could be illusory. In fact, there is a further problem, as optimisation can lead to the test set being over-used, so the system becomes over-adapted to the test set too – with the result that the final quoted accuracy could be unrealistic. To overcome this problem, it is necessary, at minimum, to divide the initial set into a training set, a test set and a validation set. In our case we took the 150 images in each major experiment and divided it equally into training, test and validation sets respectively. The validation set was kept separate and used only for testing once the method had been optimised by training and testing using the other two sets. To improve accuracy further, our final test used the first 100 images as training data and the final set of 50 images as test data.

Unfortunately, this sequential approach is liable to error if there is some trend in the data, such as the camera drifting during acquisition. In an effort to detect this possibility, we also carried out a final experiment in which the first 50 and last 50 images were used for training and the second 50 for testing. This would indicate whether drift had happened but would not be able fully to allow for it, as the second set of 50 images had already been used in previous training and could not be taken as a proper validation set (see above).

Image analysis by machine vision: refining the concept

With Experiment (e), a great many more images became available with which to train and test the schema already devised. Before work could begin, it was necessary to allow for the different type of data in these images. The sizes of the grains in these images were considerably increased, and the striations in the images now appeared as both horizontal and vertical lines, and were no longer restricted to alternate lines, but could be of higher and indeed variable periodicity. The fact that the grains now appeared larger permitted some reduction in size of the images, accompanied by significant averaging: this was found to eliminate the worst effects of the striations. This measure made the quality of the derived images similar to those for Experiment (a), and the effects of the striations became lower than the effects of ordinary noise in these intrinsically noisy images. It is instructive to examine the difference images that occur in these cases after computation of the idealised grain models and subtraction from the received images – especially for uninfested grains which do not exhibit bright patches due to larvae (see Figure 6). In all cases the difference images appear as pure noise (*i.e.* entirely random intensities) containing some bright patches, plus higher intensity patches near the ends of the grains except where these are obscured by the specially devised masks. Thus the signal itself consists of noise, and any averaging to eliminate noise immediately results in loss of signal: it was some time before this was properly realised and measures to reduce noise were abandoned.

As indicated above, the greater number of images available from Experiment (e) gave considerable potential for increasing the reliability of grain classification via additional training and testing. It proved possible to formulate a set of standard parameter adjustments that need to be optimised for any new set of grain images. These parameters were:

1. The threshold value for initial detection of the grain against its dark background.
2. The amount by which the grain mask is shrunk to form the basic exclusion zone.
3. The radial distance parameter beta used in refining the exclusion zone.
4. The length of the filtering template used in generating the grain model.
5. The width of the filtering template used in generating the grain model.
6. The amount by which the grain size is expanded to help eliminate model boundary effects.
7. The threshold for truncation of the difference signal.
8. The multiplicative factor needed to maintain adequate precision in the patch brightness.
9. The bright patch size parameter.

All nine of these parameters were adjusted by pragmatic methods to optimise the classification accuracy (the percentage of kernels classified correctly whether uninfested or infested): all parameters were needed to guarantee the capability for

optimisation, but the only way of optimising their values was to test for improvements in classification accuracy (Tables 3a–d).

In addition to the optimisation of these parameters, numerous tests were made to determine whether other optimisations were possible. For example, a number of tests of circular and elliptical filtering templates (for generating the grain model) were made, in place of the normal square and rectangular templates, but no gain in classification accuracy was obtained: in fact, the accuracy was decreased by such measures. This result is counter-intuitive, but there is a clear explanation: the striations in the images were all horizontally or vertically orientated, and having rectangular templates aligned along these directions would almost certainly help to eliminate residual effects from the striations.

Further tests were also made to determine whether totally different strategies might be more effective. In particular, an adaptive thresholding strategy was tried, as this is often successful in other applications where dark or light patches have to be detected. In our previous project on rapid automated detection (Chambers *et al.*., 1998) a refined version of this approach was used with great success. However, in the present project it was found to be a considerably less effective approach, and testing on it was soon discontinued. Conceivably it was confused by the fact that bright patches are often situated near grain boundaries, and therefore its sensitivity was artificially limited.

Summary of results from Experiment (e)
After training on the first 50 images and testing on the second set of 50 images, optimisation led to adjustment of nine important parameters which permitted the algorithm to be adapted to the grain image data. As a result, it was found possible to bring the classification accuracy up to about 81% on the training and test data (the first 100 images) and to achieve an even higher accuracy of some 84% on the validation set (the last 50 images). Further tests using the validation set for the different purpose of checking on camera drift revealed that the latter was not especially significant.

Image analysis by machine vision: re-adapting the concept to 1202 nm data
Experiment (f) (recorded at 1202 nm) provided the same number of images as Experiment (e) (981 nm), and again these could be used to perform more rigorous training and testing of the algorithm, and optimisation of the nine major parameters (Tables 4a–c). However, the software first had to be readjusted to the conditions pertaining for Experiment (d), namely, eliminating alternate rows and columns of the input images, to remove the specific type of striations that occur with this type of image. While some differences in the parameter values arose because of the need to adapt to the different sizes of the grains in the two sets of image, the algorithm required very little modification, and can really be regarded as essentially the same algorithm. At this stage, the acquisition modes at the two different wavelengths merely became corroborative of each other, and it became a question of convenience which would be chosen for any final implementation of the algorithm. At the same time, the fact that the algorithm worked equally well (within the limits shown in Tables 4a–c) on two sets of data acquired at totally different times on different cameras added confidence that the system would work on any similar setup in

different conditions. Thus the solution was not a freak success but rather one which was scientifically reproducible under different conditions.

Summary of results from Experiment (f)

In this case, optimisation was carried out on just six parameters, and very similar classification accuracies (in the region of 80%) to those for the Experiment (e) images were obtained on the first 100 images. However, significantly poorer performance was achieved with the validation set. Further tests indicated that this can be explained as being due to significant camera drift occurring by the time the validation set was obtained. Camera drift will be most likely when a vidicon camera is used with strong tungsten lighting which may warm and even overheat the camera.

Image analysis by machine vision: testing the concept on other varieties of wheat

At this stage, the opportunity was taken to examine the success of the method with varieties of wheat other than Mercia, Experiment (g). This took the form of preliminary investigations with five samples of wheat differing widely in their physical and reflective appearances. Twenty kernels from each sample, all uninfested, were imaged to determine whether the algorithm would work. It was realised that variations in size, shape or brightness profile might prevent the algorithm from working properly. (In many situations in Machine Vision, minor changes in available image data such as additional noise can prevent shapes from being defined properly.) In fact no particular problems arose in applying the algorithm to these new images. The next task was to consider its capability for accurately distinguishing infested grains of these other varieties.

The first real test was to train the algorithm on a set of 20 uninfested + 20 infested Mercia wheat grains, whose images had been gathered at the same time as those of the other five varieties, and then to test the trained algorithm on these other varieties. The procedure worked but with somewhat disappointing results in some cases (Table 5, row A). The reason is that these other varieties were sufficiently different in size, shape and intensity profile that each needed to be trained individually on sets of uninfested and infested grains of the same type. Thus it had proved insufficient to train on one variety and to test on another.

With no infested training data for the five sets of images, there seemed no possibility of taking the investigation further. However, inspection of the images showed that the five varieties were somewhat narrower than the Mercia variety: it was therefore decided to test them on a version of the algorithm which had been modified to allow for the difference in width. This markedly raised the success rate for these other varieties but still left the performance disappointingly below that for Mercia wheat (the only case for which full training data were available). It was therefore decided to do a further test. This was a simulation which involved placing a bright spot on all the grains and determining whether the algorithm would be able to detect it. The intensity and size of the spot were determined by adding the spot to uninfested grains of the Mercia variety, and adjusting their values to give similar performance (as far as could be judged) to that obtained by training and testing sets of uninfested and infested grains. These tests were applied for both settings of the nine parameters used in the immediately previous tests. Both of these tests gave much more encouraging results, and the one which used parameter values appropriate to narrower grains gave the best results of all (Table 5, row E). It appeared that all the varieties should be able

to respond similarly to the Mercia variety, with success rates of 80% or more, and thus the algorithm concept seemed to have general applicability for wheat.

In spite of these successes, it has to be emphasised that these latter tests were only simulations, and the accuracy of the figures still depends on the validity of taking a bright spot of the type that appears on an infested Mercia grain and transferring it to these other varieties. There is no absolute guarantee – just a very strong likelihood – that bright patches that would appear on infested grains of these other varieties would have the same degree of visibility and detectability. However, no reason is known why this should not be so: we have a high degree of confidence that the algorithm is soundly based and is equally applicable to all these types of wheat grain.

Image analysis by machine vision: investigating the meaningfulness of the concept
In Experiment (h), a set of 20 uninfested + 20 infested grains was examined by NIR, and 20 X-ray images were obtained of the infested grains. We proceeded to train and test the algorithm on the 20 uninfested + 20 infested grains, using both the normal and the narrow grain parameter settings (with so few samples it was impossible to carry out a full optimisation of the algorithm). The narrow grain settings seemed to be more appropriate. However, on examining the difference images produced by the algorithm, the signal-to-noise ratio was quite low in some cases, and there was some doubt as to which of the bright patches indicated the most relevant position on the grain. Accordingly, we marked some of the bright patch locations as "reliable position indicators" and others as "unreliable position indicators" in a subjective judgement. We judged that out of the 20 images of infested kernels, the location of the bright patches was reliable in 12 cases. On examining these locations on the X-ray images, we found that 7 were well within the boundaries of the larvae within the grains, 4 were within the cavity once occupied by the larvae, and one was in a position that seemed to bear no relation to the larval cavity (Table 6). This last result is puzzling, though 100% correlation between the insect and the effect it produces on the surface of the grain (as seen by NIR imaging of the surface) is hardly to be expected. (It should also be noted that noise itself can generate a bright patch, and this may be the mechanism by which uninfested grains sometimes appear to be infested.) What is more relevant is that there is a very high correlation between the location of the bright patch in the NIR image of the kernel and the present or past location of the larva as revealed by X-ray. This poses interesting questions about the exact mechanism for detection in the NIR images. These results are largely consistent with the conclusion from the spectroscopic study (Section 7) that the origin of the observed NIR effect is the loss of starch.

Image analysis by machine vision: estimation of batch infestation
Here we consider the practical consequences of the strategy adopted in this work leading to classification accuracies in the range 80–85%. With such an indirect relation between internal infestation and its detection from properties of the grain in noisy NIR images, it is difficult to see how the classification of any individual grain could be made significantly higher. Even if some new algorithm, optimisation tool or huge training set became available, reliable classification of more than 90% of grains would appear unachievable using this approach. However, in the practical situation, the detection of infestation in a *batch* of grain is more important than the classification of *individual* grains. This raises the question of how infestation of a batch is to be defined. Taking the basic specification as the requirement to inspect 3 kg of grain in

3 minutes (Chambers *et al.,* 1998), we can refine this specification to assessing whether the level of infestation is serious, e.g. ~0.5%. This level of infestation corresponds to a 3 kg batch of grain in which each of three gravid female weevils were to infest 100 kernels. Such a figure is reasonable given that the female grain weevil is capable of laying up to 254 eggs (Mallis, 1997). It is relevant to take 3 gravid females as corresponding to the number of insects which could have gone unnoticed by a visible light inspection system. Thus we envisage a system in which a visible light inspection system and an NIR inspection system work together, and the former detects external insects while the latter detects internal insects. If the visible light inspection system finds insects, there is no need to take account of the NIR inspection system. The latter is needed when no insects are detected by the visible light inspection system to confirm the absence of a developing infestation in which all the insects are hidden.

Given that 3 kg corresponds to about 60,000 grains, the NIR inspection system needs to be capable of detecting infestation when some 300 infested grains are present. This will be taken as our definition of *an infested batch* or *an infestation*. In what follows we concentrate on detecting such an infestation.

For practical reasons, the numbers of grain images in both of our main experiments [Experiments (e) and (f)] are far below 60,000. In addition, the level of infestation in these images is substantially different from 0.5%. These factors mean that the distributions of patch intensity have to be analysed carefully, and theoretical estimates made of the relevant probabilities for individual grain detection and batch detection.

If one grain is analysed there will be ~80% probability of accurately detecting infestation but if several grains are analysed the probability of accurately detecting infestation will be increased. For example, if one infested grain is examined, the chance of *missing* an infestation is around 0.2; if two such grains are measured the chance of missing an infestation drops to around 0.2^2; and for n infested grains the chance of missing an infestation drops to around 0.2^n. This approach offers the opportunity to reduce the chance of missing an infestation to arbitrarily low levels (especially if n ~ 60,000). However, the situation is complicated by the fact that false alarms (false positives) appear as well as missed insects (false negatives), and a full analysis will have to take account of both. We have carried out analyses based on the respective per grain false positive and true positive probabilities, P_1 and P_2, and how they vary with intensity thresholds. The simplicity of the 0.2^n calculation then breaks down, and we are left with probabilities P_1,P_2 coupled with variances proportional to $P_1(1 - P_1)$, $P_2(1 - P_2)$. To get observable results from these parameters, we must allow for the different numbers of grains in the two distributions, by multiplying in the two cases by 59,700 and 300 (59,700 = 60,000 – 300). What is especially important is that these analyses are probabilistic, the probabilities having to be calculated using the variances of the distributions. For the values of P_1, P_2 found experimentally, and assuming Gaussian error distributions truncated at the 3σ limits, we find that it is not possible to get sufficient information from 60,000 grains to guarantee detecting 0.5% infestation. However, as an alternative, we can calculate the probability of missing a 0.5% infestation in 60,000 grains: thus we find that if the probability of a false positive infestation is taken to be equal to the probability of a false negative infestation, the latter probability is around 5.3%. If on the other hand, the probability of a false positive infestation is permitted to rise to around 13%, the

probability of missing an infestation can be made to fall to around 3%. Thus there is definite possibility of missing an infestation, but the probability of this is reduced to quite low levels. Further analysis of the trade-off is difficult as it has to be carried out numerically by integrating error functions, and no simple closed formula applies.

Summary of results from Experiment (h)
While classification accuracy for each individual grain was normally around 80%, it seemed unlikely that the approach could lead to performance above 85-90%. However, individual grain assessment is of lesser importance than assessment of the batch of grain. This work has shown that an overall batch performance in the 95% category could be achieved by consideration of fairly large numbers of grains (typically 60,000) in a sample of 3 kg. The 95% probability of finding a 0.5% infestation results from the case when the numbers of false positives and false negatives are equal. In the practical situation, where a false alarm is much less important than a missed insect, the number of false positives could be increased, allowing the probability of finding a 0.5% infestation to rise to around 97%. Given the potential for this detection of internal infestation to be entirely non-subjective and automated, such a result is remarkable and unparalleled by any other method.

Image analysis by machine vision: improving the speed of the algorithms
When initially devised, the algorithm took some 70 ms to process a single grain image of the type obtained from Experiment (f). The algorithm was timed on a 350 MHz Pentium II, and timing accuracy was boosted by timing 1000 repetitions of the individual sub-processes. This works out at 70 minutes for a sample containing 60,000 grains, and is 20–30 times slower than the specified figure of 3 minutes. In addition, it should be remarked that this takes no account of the image acquisition time. However, the latter can essentially be ignored for the present analysis, as any practical system will have a separate framestore which will operate in parallel with the computer, so acquisition will not impose any direct load on the computer.

To overcome the speed problem, the algorithm was divided into a number of modules whose speeds were analysed individually. It was soon found that most of the execution times could be halved and that some could be improved further, by quite simple measures. (The original form of the algorithm was dictated more by scientific effectiveness and accuracy of classification than by speed considerations: indeed, the latter were deliberately suppressed during the development stage.) The simple measures adopted involved including more in-line code, preliminary calculation of relevant parameters rather than repeated estimation when this was not absolutely vital, and rapidly by-passing background regions of the images. These measures reduced the overall execution time to 28 ms. It was also deduced that the preliminary stage of averaging for noise (specifically striation) reduction and accompanying reduction of the sizes of the input images would not be required in any eventual practical system, as (a) the images would then be acquired at the correct resolution, and (b) better framestores would be used which would not give rise to striations. (The latter constitute an artefact in which some faulty electronic timing mechanism on the framestore causes alternating lines of high and low intensity, and a vendor would be able to overcome such problems.) In addition, display of the final image would not be required, as logging and classification of the output data would not require the images themselves to be retained; similarly, output zooming of the images would not be required. Elimination of these initial and final processes would cut out another 5 + 1

= 6 ms, leaving an overall execution time of 22 ms. This left two processes which were especially slow because they operated within relatively large windows: these were the grain modelling filter and the bright patch scanning filter, at 16 and 4 ms respectively. Special attention to these processes permitted reduction in their respective execution times to 3 and 2 ms. This required re-engineering the overall design of the algorithm to permit a preliminary reduced size mask to be produced to restrict the region of application of the final scanning filter. Such changes of strategy in algorithm design are often vital to the process of improving timing of complex inspection procedures. On the other hand, it should be remarked that such improvements may only be available if some compromise is taken over performance, and a tradeoff between performance (i.e. classification accuracy) and speed commonly exists. Here, however, there appeared to be no such penalty, and the final execution time of the algorithm was 7 ms. There appeared to be no way to cut the software speed by more than a few percent from this figure, and indeed a great deal more effort would be needed to cut it to 6 ms, and in that case the algorithm would become a lot more special purpose and ability to adapt to new conditions would probably be lost.

Overall, these considerations brought the execution time down to practically useful levels. The actual time needed to process 60,000 grains would thereby be reduced to some 7 minutes on a 350 MHz Pentium II. Further speedup depends on having a faster processor, or using several processors. Much depends on when an implementation is to be pursued. If this is to occur in the next six months, the best option will be to use two 450 MHz Pentiums. However, if it is to occur in the autumn of 2001, by then the best option may well be a single 800 MHz Pentium, or its equivalent.

Integration of the algorithm with the system previously developed for detection of contaminants in grain images using visible light adds a certain constraint. In particular, a connected components analysis of all the objects in an image must be performed, thereby adding an estimated 1 ms to the total times listed above. Thus two 500 MHz Pentium II computers would currently be required, or a single 1000 MHz Pentium, or its equivalent, at some future date.

7. SENSITIVITY AND ORIGIN OF RESPONSE TO INFESTATION

Introduction

To ensure that the differences observed in the NIR images between infested and uninfested kernels are due to the presence of insects rather than some artefact, a full NIR spectroscopic study was undertaken. This had the additional purpose of helping optimise sensitivity of detection.

Method

Single uninfested and infested kernels (20 replicates each), for both grain weevil large larvae and pupae, were scanned by reflectance over the whole wavelength range 700nm - 2500nm at 2 nm intervals on an NIRSystems 6500 spectrometer.

An additional study was carried out by repeatedly monitoring kernels at intervals of two or three days as the insects developed. For this purpose, three replicate samples of kernels infested with grain weevil larvae, about 1.5-2.5 weeks after egg-laying,

were prepared, along with two replicate samples of single uninfested kernels from the same culture. The infested kernels selected from the culture were those with signs of larval presence in their X-radiographs which were just discernible to the naked eye. Scanning by NIR commenced on the same day as the examination by X-ray.

Results from NIR spectroscopy on single kernels

Wavelength pairs, consisting of a measurement wavelength (to detect changes due to infestation) and a correction wavelength (to correct for variation in kernel size) were identified which gave very impressive differences between the uninfested and infested kernels. Larval and pupal data were similar so they were combined, giving n=80 total. Of the very near-infrared wavelength models evaluated, the greatest sensitivity for spectroscopic detection of internal infestation was achieved by measuring *log1/R(982 nm)-log1/R(1014 nm)*. This model discriminated between uninfested kernels and kernels infested with either large larvae or pupae with 96% classification accuracy. One other model, *log1/R(972 nm)-log1/R(1032 nm),* gave equal detection performance, but the wavelength 972 nm corresponds to a peak associated with water. Hence this latter model is likely to be prone to interference from changes in grain moisture content. Moving to the very near-infrared improves sensitivity compared to using a model based on the wavelength 1200 nm identified in previous preliminary work. This model, *log1/R(1200 nm)-log1/R(1330 nm),* gives a classification accuracy by spectroscopy of 93%. This is an important advance because imaging below 1100nm could be carried out using an inexpensive silicon detector-based CCD camera.

The wavelengths 982 nm and 1014 nm both lie within a broad NIR band which coincides with a published starch band centred at 990 nm. This band decreased with infestation. When examining kernel spectra over the full wavelength range used (700-2500 nm), starch was implicated as the origin of the response to infestation in 5 out of the 6 wavelength models found. These observations strongly suggest that at both 982 nm and 1200 nm, discrimination between uninfested and infested kernels originates from detection of the loss of kernel starch as a consequence of insect feeding.

Almost immediately after the start of the repeat scanning experiment, when the culture was 2–3 weeks after egg-laying, consistent differences between the infested and uninfested kernels were observed by NIR. This suggests that the spectroscopic method has a sensitivity approaching that obtained using X-ray inspection.

Both the work to develop correlation models and the study of change in NIR response with age of internal infestation have been written up in full and published (Ridgway *et al.*, 1999).

8. INCORPORATION OF FACILITY TO DETECT HIDDEN INFESTATION INTO EXISTING MACHINE VISION SYSTEM FOR CONTAMINANTS EXTERNAL TO GRAIN KERNELS

While it is attractive to attempt to produce an integrated machine which will examine the grain sample at visible (for external contaminants) and NIR wavelengths (for hidden infestation), this will still require separate cameras and software analysis systems for the two cases. There would appear to be no saving in computer software

or hardware from combining the two machine vision systems, as they operate on images of different types and different resolutions. Additionally, a significant degree of parallel processing is required to solve both of the computational tasks, so there is anyway a benefit from employing more than one computer. Nevertheless, similar mechanical arrangements can probably be used to form the basis for a combined system of this type. For example, a single vibratory (or other) feeder would feed a single conveyor, and the grains would travel along this, and may also be projected and photographed in mid-air at the end of the conveyor.

The NIR inspection system requires a high uniform rate of grain presentation, whereas the visible light system largely ignores the grains and does not need to scrutinise them at a steady rate. Hence there are still some open questions on how best to combine the inspection processes on a single mechanical conveyor system, though there is also considerable freedom in how it is to be achieved. Thus much will depend on the inclinations and capabilities of the vendor who takes on the task of marketing the overall system.

9. CONCLUSIONS

1. For detecting internal infestation in grain kernels by NIR imaging at 1202 nm, the best lighting arrangement uses 4 x 100W lightbulbs above a vertical cylinder of 31 cm diameter with a distance of 30 cm between sample and camera. An NIR vidicon camera gave better quality images than CCD cameras using either silicon or indium gallium arsenide detectors.

2. To detect internal infestation by machine vision assessment of bright patches on the kernels, it was necessary to overcome the difficulties resulting from the very low signal-to-noise ratio in the NIR images and the natural variability of the grains in their brightness, size and shape, by adopting an approach based on modelling grain intensity profiles.

3. For detecting internal infestation in grain kernels by NIR imaging at 981 nm, the best lighting arrangement uses 4 x 40W lightbulbs above a vertical cylinder of 20 cm diameter and with a distance of 5 cm between the sample and camera. At this wavelength, it was possible to use a silicon-detector CCD camera, which is cheaper and potentially more reliable than the vidicon camera required for imaging at 1202 nm.

4. Human inspection of NIR images recorded at 981 nm from pairs of Mercia variety kernels resulted in successful classification of the infested kernels in each pair in 78% of the 51 cases, demonstrating that there are consistent differences due to internal infestation and providing a measure by which the success of the machine vision approach could be judged.

5. The high degree of variability between the images of grain kernels meant that the machine vision approach could not average them to determine an idealised grain model. Instead it was necessary to derive an idealised model of each grain from the image of that grain and improve the model only by adjusting internal algorithm parameters.

6. Searching for bright patches due to insect infestation proved to be unreliable around the kernel boundaries and particularly so at its ends. However, such searching was successful, once these problematic areas had been excluded, by comparing the average intensity of the brightest patch of each kernel with a modelled intensity achieved by applying a noise filter.

7. The large (150) set of images at 981 nm allowed optimisation of the nine internal algorithm parameters covering brightness threshold values and the sizes of the exclusion zone, filtering template and bright patch. As a result, classification accuracy of about 81% was achieved on the test and training sets (the first 100 images), and 84% on the validation set (remaining images). There was no evidence for camera drift during the recording of these sets. This is an improvement on human inspection, where similar classification accuracy was obtained but the rate of infestation was known.

8. With the large (150) set of images at 1202 nm, optimisation of just six parameters gave similar classification accuracies of around 80% for the training and test sets but a poorer performance with the validation set was probably due to camera drift caused by overheating from the lighting.

9. Preliminary investigations have been undertaken of five samples of wheat varieties other than Mercia differing in physical and reflective appearances. Despite their differences in size and grain intensity profiles, classification rates of 80% or higher were obtained between images of these uninfested kernels and images of kernels which were simulated as infested by the artificial addition of a bright spot similar to those seen on infested Mercia kernels. Thus the algorithm concept seems to have general applicability to wheat.

10. There is a very high correlation between the location of the bright patch in the NIR image of the kernel and the present or past location of the larva as revealed by X-ray.

11. The probability of finding an infestation of 0.5% of the kernels in a batch of 3 kg could be around 97% which, given that this method is entirely non-subjective and could be automated, is remarkable and unparalleled by any other method.

12. The speed of the machine vision process was increased by a factor of ten by re-engineering the overall design of the algorithm, spending less time on background regions of images and eliminating processes which would not be needed in the practical application. This reduced the time which would be needed to process the 3 kg sample to around 7 minutes. Further reduction to achieve the target time of less than 3 minutes should be achievable by using greater processing power.

13. A study over the full NIR spectroscopic range of 700-2500 nm strongly suggests that at both 982 nm and 1200 nm, discrimination between uninfested and infested kernels originates from detecting the loss of starch in the kernel as a consequence of insect feeding.

14. Repeated scanning of kernels shows that consistent differences are observable between those which are uninfested and those in which grain weevil eggs were laid 2-3 weeks previously. This is almost as soon as such infestation would be detectable by X-ray.

15. The system developed in this project to detect hidden infestation could share mechanical arrangements relating to sample presentation with those to be used in the system for detecting external contamination. However, the two systems use different cameras and software analysis and differ in that one scrutinises the kernels while the other largely ignores them. There will be few technical constraints to limit the vendor who develops this invention to a marketable product bringing its major benefits to the cereal trade.

10. NEEDS FOR FUTURE RESEARCH AND DEVELOPMENT

- Improved (striation-free) image acquisition‡
- Mechanical grain orientation, or its equivalent in software
- Full tests with larger quantities of data, including various varieties of wheat and other species of grain
- Connected components analysis to cope with multiple objects in any image
- Mechanical integration with visible light inspection system.

11. PUBLICITY

The achievements described in this report represent significant advances in intellectual property. This must be protected to encourage commercial exploitation of the system to be based on it. Nevertheless, the investigators have been active in the promotion of their work and have taken care to do this without compromising the protection. Once the intellectual property has been protected, the work will be further promoted in trade journals and primary scientific journals of repute. The paper describing work on the origin of the observed effect and optimising the sensitivity has already been published (Ridgway *et al.*, 1999).

12. ACKNOWLEDGEMENTS

We are grateful to Dr S. C. Hook of RHM Technology for supplying the samples of five wheat varieties other than Mercia; to Mr. D. Charles of RHUL for providing the programming shell which was used as a front end for the Machine Vision test system; and to Mr I. A. Cowe of Foss Electric Development (UK) Ltd (now Foss Tecator R&D) for access to the NIRSystems spectrometer and considerable help with the spectroscopic study.

‡ It will be necessary to ensure that the images are sufficiently noise-free without repeated grabbing and averaging (though an alternative may be to use several cameras operating together).

13. REFERENCES

Chambers J., and Ridgway C. (1996). Rapid detection of contaminants in cereals. In *Near Infrared Spectroscopy: The Future Waves*. Eds. A. M. C. Davies and Phil Williams, NIR Publications, Chichester, UK, 484-489.

Chambers J. , Davies E. R., Ridgway C., Mason D. R. and Bateman M. W. (1997a). Detection of contaminants in cereals by image analysis. Poster at Image Analysis for Automating Agricultural Processes meeting, Society of Chemical Industry, 28 January 1997.

Chambers J. , Davies E. R., Ridgway C., Mason D. R. and Bateman M. W. (1997b). Detection of contaminants in cereals. Poster at Cereals 97, 11-12 June 1997.

Chambers J., Ridgway C., Baker C. W. and Barnes R. J. (1993a). Improving quality by rapid pest detection. *Aspects of Applied Biology 36*, Cereals Quality III, 465-470.

Chambers J., Ridgway C., Davies E. R., Mason D. R. and Bateman M. W. (1998). Rapid automated detection of insects and certain other contaminants in cereals. HGCA Project Report No. 156, 64pp.

Chambers J., Ridgway C., van Wyk C.B., Baker C.W. and Barnes R.J. (1994). The potential of near-infrared spectroscopy for the rapid detection of pests in stored grain. Home-Grown Cereals Authority, London Project Report No.92, 51pp.

Davies E. R., Mason D. R., Bateman M. W., Chambers J. and Ridgway C. (1998a). Linear feature detectors and their application to cereal inspection. Proc. EUSIPCO'98, Rhodes, Greece, 8-11 Sept., pp 2561-2564.

Davies E. R., Bateman M. W., Chambers J. and Ridgway C. (1998b). Hybrid non-linear filters for locating speckled contaminants in grain. Non-linear signal and image processing, IEE Colloquium, 22 May 1998, pp 12/1-12/5.

Davies E. R., Bateman M. W., Mason D. R., Chambers J. and Ridgway C. (1999). Detecting insects and other dark line features using isotropic masks. IEE Conference on Image Processing and its Applications, Manchester, 12-15 July 1999, 5pp.

Davies E.R. (1997). Machine Vision: Theory, Algorithms, Practicalities, Academic Press (2nd edition), pp. xxxi + 750.

Mallis A. (1997). Handbook of Pest Control, 8th Edition.

Paulsen M. R. and McClure W. F. (1986) Illumination for computer systems. Transactions of the ASAE, 29(5), 1398.

Ridgway C. and Chambers J. (1996). Detection of external and internal insect infestation in wheat by near-infrared reflectance spectroscopy. Journal of the Science of Food and Agriculture, 71, 251-264.

Ridgway C. and Chambers J. (1998) Detection of insects inside wheat kernels by NIR imaging. Journal of Near Infrared Spectroscopy, 6, 115-119.

Ridgway C., Chambers J. and Cowe I.A. (1999). Detection of grain weevils inside single wheat kernels by a very near infrared two-wavelength model. Journal of Near Infrared Spectroscopy, 7 (4), 213-221.

APPENDIX A: GLOSSARY OF TERMS

For a source book on further information on image analysis and inspection issues see Davies (1997).

Adaptive thresholding	Thresholding in which the threshold value is automatically adjusted to match the local conditions in an image
Algorithm	Sequence of rules to conduct a mathematical process
Background subtraction	Subtraction of the image of the background-only from the image containing the sample, to improve image quality
Centroid	The centre of gravity of a shape (if the shape were suspended from this point, it would remain horizontal)
Cluster analysis	Analysis of the various samples in a feature space to group them so that similar samples appear within the same cluster and dissimilar samples appear in different clusters
False negative	Instance of a failure to locate an object or contaminant
False positive	False alarm found when searching for an object or contaminant
Neural networks	Networks of neurones to which input signals are fed and from which processed outputs emerge that in some way help to classify the input data. The basic concept is to mimic the processing of signals by the neurones in the brain
Pixel	Short form of 'PICture CELL'. The smallest resolved area in a digital image, normally taken to be square
Signal-to-noise ratio	The factor by which the signal under investigation is larger than the background noise level. High signal-to-noise ratios make signals easier to recognise. In difficult cases the ratio may be less than unity.
Template	An idealised form of an object which is used for matching purposes in images.
Thresholding	The process of converting a grey-scale image into a binary image by applying a threshold: during conversion, any pixel which is darker than the threshold becomes black, and any other pixel becomes white

APPENDIX B: PRELIMINARY SPECIFICATION FOR SYSTEM DEVELOPED IN THIS PROJECT

System description
Kernel feeder: *input capacity of 3 kg batches to present to machine as 6600 portions each of about 0.45 g, including sieve to protect against blockage*
Image capture: *monochrome or linescan equivalent at 256 x 256 pixels, 8 bit precision*
Computer: *two off (or twin) Pentium 500 MHz, or equivalent*
Framestore: *two frames of 256 x 256 pixels, 8 bit precision, with double buffered capability, possibly with linescan interface*
Output devices: *VDU, printer, light and audible alarm*

Sample handling
Typical samples: *wheat, barley, oats*
Capacity: *batches of 3 kg in 3 minutes*
Dimensions of individual portions imaged: *2.4 x 2.4 cm*
Size of sample: *9 ± 2 cereal kernels; no item to be less than 15 pixels from image boundary, or from adjacent kernels*
Image capture rate: *37 images per second.*

Performance requirements
Types of contaminant: *larval insects internal to kernels*
Minimum limit of detection: *60 larval insects in 60,000*
Mean false positive and false negative rates: *better than 0.3% for batch infestation*
Presentation of results: *detection of likely contaminants to result in incrementation of prominent contaminant counter, printout of numbers of contaminant detected; likely contaminants to be separated from main batch for subsequent inspection*.

Operation data
Image analysis procedure: *object recognition, logging, interface to delivery/rejection/sample storage grain handling mechanisms*
Calibration mode: *preliminary screening of small samples of grain without contaminants; continual updating during operation*
Operating system: *Windows NT or other robust system for factory use*
Interface: *via controlling computer; possible printer output*
Operating temperature: *0 - 40 C (preferred less than 30 C to avoid electronic failures)*

Technical data
Power specifications: *for image analysis expected to be about 240 V, 2-3 kW*
Dimensions and weight: *for image analysis expected to be about 70 cm (h) x 30 cm (l) x 30 cm (w)*

All technical specifications will be subject to agreement between vendor and IPR holders, to ensure that variations (such as use of linescan camera, choice of framestore, alternative operating systems, etc.) do not interfere with the basic algorithm concept. It must be noted that the system developed by the IPR holders is a laboratory development system, which is not itself suitable for field use: thus the vendor would be *expected* to suggest his preferred adaptations to the specification laid down by the IPR holders.

Delivery to vendor: It is intended that the vendor will receive the following items under licence from the IPR holders:

1. Relevant know-how and experience.
2. Relevant source code image analysis routines which pass image data and interpretation data via allocated memory areas.

In addition, the following will be available:

3. Further advice on how to handle unforeseen cases in the input images.

However, item 3 cannot be delivered without charge, as only items 1 and 2 have been paid for by the HGCA-funded project.

The IPR holders will *not* deliver executable code, as they will not be familiar with the vendor's marketable systems. In particular, *all responsibility for handling windows-type (image display) routines will rest on the vendor*: this reflects the 'proof of concept' approach adopted by the IPR holders.

NOTE: The above specification represents our present scientific and technological thinking. It has been written in good faith **but has not yet been vetted or approved by the RHUL External Relations Department or the CSL staff with similar responsibility**: hence no formal guarantees can be given over it. It is intended solely to clarify our present intentions with regard to development and marketing of the system.

Table 1: Sets of Images Recorded

Expt	Wavelength	Purpose	Samples
a	1202 nm	To improve image quality and provide early data for machine vision	27 uninfested and 27 infested kernels imaged in total
b	981 nm	To improve image quality	16 uninfested and 16 infested kernels imaged in total
c	981 nm	To determine whether information in the image is sufficient	51 uninfested and 51 infested kernels imaged in total
d	1202 nm	To develop main concept in machine vision	25 uninfested and 25 infested kernels imaged in total
e	981 nm	To refine machine vision concept	75 uninfested and 75 infested kernels imaged in total
f	1202 nm	To re-adapt machine vision concept to 1202 nm data	75 uninfested and 75 infested kernels imaged in total
g	981 nm	To test machine vision concept on wheat varieties other than Mercia	20 uninfested (Mercia), 20 infested (Mercia), 20 uninfested (Malacca), 20 uninfested (USA Northern Spring), 20 uninfested (Canadian), 20 uninfested (Shango) and 20 uninfested (Australian White) imaged in total
h	981 nm	To test the meaningfulness of the machine vision concept	

Table 2 Results for 46 Images, Experiment (d)

grain	patch value	actual class	result class	error
1	137	0	0	1
2	179	1	1	1
3	96	0	0	1
4	158	1	1	1
5	134	0	0	1
6	157	1	1	1
7	105	0	0	1
8	168	1	1	1
9	126	0	0	1
10	158	1	1	1
11	119	0	0	1
12	206	1	1	1
13	170	0	1	0
14	147	1	1	1
15	144	0	1	0
16	203	1	1	1
17	110	0	0	1
18	118	1	0	0
19	144	0	1	0
20	173	1	1	1
21	140	0	0	1
22	218	1	1	1
23	140	0	0	1
24	218	1	1	1
25	123	0	0	1
26	158	1	1	1
27	115	0	0	1
28	215	1	1	1
29	110	0	0	1
30	195	1	1	1
31	121	0	0	1
32	179	1	1	1
33	124	0	0	1
34	159	1	1	1
35	127	0	0	1
36	136	1	0	0
37	136	0	0	1
38	140	1	0	0
39	145	0	1	0
40	176	1	1	1
41	137	0	0	1
42	157	1	1	1
43	135	0	0	1
44	157	1	1	1
45	135	0	0	1
46	200	1	1	1

This table shows the results of applying the algorithm to the individual grains for the images of Experiment (d). The 46 grains are alternately uninfested (0) and infested

(1), as indicated in column 3. The class determined by the algorithm is shown in column 4 and any errors are indicated by zeros in column 5.

The algorithm was trained on these images using only the information in columns 2 and 3. The training procedure computes a threshold on the patch brightness value: in this case the threshold value was 145. Higher values lead to the 'infested' (1) classification, and lower values lead to the 'uninfested' (0) classification. Classification was carried out on all 46 grains, though classification on the training set is liable to appear too accurate because of over-adaptation to the training set.

The overall accuracy for classification on the training set was 0.848 (39 of the 46 grains classified correctly).

Table 3a Parameter Optimisation

parameter	value	accuracy	
		I (train and test)	II (test only)
1	-1	0.900	0.720
	0	0.920	0.760
	1	0.820	0.780
2	–	–	–
	0	0.920	0.760
	1	0.840	0.760
3	-1	0.820	0.720
	0	0.920	0.760
	1	0.780	0.760
4	-1	0.900	0.740
	0	0.920	0.760
	1	0.840	0.680
5	-1	0.840	0.760
	0	0.920	0.760
	1	0.800	0.760
6	-1	0.900	0.760
	0	0.920	0.760
	1	0.900	0.760
7	-1	0.880	0.740
	0	0.920	0.760
	1	0.860	0.760
8	-1	0.920	0.760
	0	0.920	0.760
	1	0.800	0.800
9	-1	0.900	0.760
	0	0.920	0.760
	1	0.860	0.780

In this table, the nine parameters have been adjusted relative to their central value by a scale factor times the value listed in column 2. Accuracy I refers to the first 50 images; II refers to the second 50 images.

For parameter 2, no reduction was possible, as a negative value would be meaningless.

Table 3b Parameter Optimisation

parameter	value	accuracy	
		II (train and test)	I (test only)
1	-1	0.740	0.820
	0	0.780	0.840
	1	0.820	0.780
2	–	–	–
	0	0.780	0.840
	1	0.760	0.860
3	-1	0.760	0.760
	0	0.780	0.840
	1	0.780	0.720
4	-1	0.740	0.900
	0	0.780	0.840
	1	0.820	0.780
5	-1	0.740	0.760
	0	0.780	0.840
	1	0.760	0.840
6	-1	0.780	0.840
	0	0.780	0.840
	1	0.780	0.840
7	-1	0.820	0.800
	0	0.780	0.840
	1	0.800	0.840
8	-1	0.780	0.840
	0	0.780	0.840
	1	0.800	0.820
9	-1	0.760	0.860
	0	0.780	0.840
	1	0.820	0.840

Accuracy I refers to the first 50 images; II refers to the second 50 images.

Table 3c Results for First 100 Images, Experiment (e)

grain	patch value	actual class	result class	error
1	184	0	0	1
2	217	1	1	1
3	203	0	0	1
4	233	1	1	1
5	152	0	0	1
6	237	1	1	1
7	174	0	0	1
8	207	1	1	1
9	191	0	0	1
10	241	1	1	1
11	168	0	0	1
12	228	1	1	1
13	172	0	0	1
14	248	1	1	1
15	179	0	0	1
16	238	1	1	1
17	169	0	0	1
18	204	1	0	0
19	183	0	0	1
20	228	1	1	1
21	231	0	1	0
22	217	1	1	1
23	178	0	0	1
24	222	1	1	1
25	184	0	0	1
26	216	1	1	1
27	166	0	0	1
28	212	1	1	1
29	240	0	1	0
30	211	1	1	1
31	204	0	0	1
32	255	1	1	1
33	188	0	0	1
34	255	1	1	1
35	205	0	0	1
36	214	1	1	1
37	194	0	0	1
38	232	1	1	1
39	178	0	0	1
40	244	1	1	1
41	185	0	0	1
42	215	1	1	1
43	202	0	0	1
44	253	1	1	1
45	205	0	0	1
46	255	1	1	1
47	182	0	0	1
48	242	1	1	1
49	190	0	0	1
50	190	1	0	0

grain	patch value	actual class	result class	error
51	182	0	0	1
52	199	1	0	0
53	134	0	0	1
54	198	1	0	0
55	180	0	0	1
56	247	1	1	1
57	163	0	0	1
58	177	1	0	0
59	178	0	0	1
60	235	1	1	1
61	182	0	0	1
62	238	1	1	1
63	169	0	0	1
64	210	1	1	1
65	188	0	0	1
66	209	1	1	1
67	172	0	0	1
68	204	1	0	0
69	145	0	0	1
70	210	1	1	1
71	210	0	1	0
72	198	1	0	0
73	174	0	0	1
74	208	1	1	1
75	210	0	1	0
76	247	1	1	1
77	185	0	0	1
78	208	1	1	1
79	191	0	0	1
80	254	1	1	1
81	236	0	1	0
82	255	1	1	1
83	197	0	0	1
84	196	1	0	0
85	190	0	0	1
86	160	1	0	0
87	188	0	0	1
88	220	1	1	1
89	215	0	1	0
90	241	1	1	1
91	172	0	0	1
92	255	1	1	1
93	177	0	0	1
94	239	1	1	1
95	175	0	0	1
96	228	1	1	1
97	214	0	1	0
98	213	1	1	1
99	199	0	0	1
100	236	1	1	1

This table shows the results of applying the algorithm to the individual grains for the images of Experiment (e). The 100 grains are alternately uninfested (0) and infested (1), as indicated in column 3. The class determined by the algorithm is shown in column 4 and any errors are indicated by zeros in column 5.

The algorithm was trained on the first 50 images using only the information in columns 2 and 3. The training procedure computes a threshold on the patch brightness value: in this case the threshold value was 205.7. Higher values lead to the 'infested' (1) classification, and lower values lead to the 'uninfested' (0)classification. Classification is carried out on all 100 grains, though classification on the training set is liable to appear too accurate because of over-adaptation to the training set.

The overall accuracy for classification on the training set is 0.92, and for classification on the test set is 0.76. Taking a weighting of 1:2 between these figures gives a result of about 0.81. (Taking the test set result is unnecessarily harsh, as indicated on examining the result of training on the second 50 images and testing on the first 50 (see Tables 3a,b), which gives the respective figures 0.780, 0.840.)

Table 3d Summary of Performance

		accuracy	
training set	test set	training	testing
I	II	0.920	0.760
II	I	0.780	0.840
I+II	III	0.820	0.854
I+III	II	0.850	0.771

I refers to the first 50 images; II refers to the second 50 images. Some faulty images had to be eliminated from the analysis, so the third set III contained only 48 images.

It is gratifying to note the high performance when training on I+II and testing on III – the validation set which had not been looked at earlier.

Table 4a Parameter Optimisation

parameter	value	accuracy	
		I (train and test)	II (test only)
1	-1	0.740	0.700
	0	0.760	0.820
	1	0.820	0.800
2	-1	0.760	0.840
	0	0.760	0.820
	1	0.800	0.800
3	-1	0.800	0.820
	0	0.760	0.820
	1	0.740	0.740
4	-1	0.760	0.860
	0	0.760	0.820
	1	0.740	0.760
5	-1	0.780	0.760
	0	0.760	0.820
	1	0.760	0.820
9	-1	0.780	0.800
	0	0.760	0.820
	1	0.780	0.820

In this table, the nine parameters have been adjusted relative to their central value by a scale factor times the value listed in column 2. Accuracy I refers to the first 50 images; II refers to the second 50 images.

Note that the rigorous optimisation of the Experiment (d) data led to reduced need for optimisation of the Experiment (f) data: in particular, parameters 6–8 did not have to be optimised further in this case.

Table 4b Parameter Optimisation

parameter	value	accuracy	
		II (train and test)	I (test only)
1	-1	0.740	0.760
	0	0.860	0.760
	1	0.800	0.820
2	-1	0.840	0.760
	0	0.860	0.760
	1	0.780	0.800
3	-1	0.820	0.800
	0	0.860	0.760
	1	0.780	0.780
4	-1	0.820	0.760
	0	0.860	0.760
	1	0.780	0.720
5	-1	0.780	0.740
	0	0.860	0.760
	1	0.840	0.760
9	-1	0.800	0.760
	0	0.860	0.760
	1	0.840	0.740

Accuracy I refers to the first 50 images; II refers to the second 50 images.

Table 4c Summary of Performance

		accuracy	
training set	test set	training	testing
I	II	0.760	0.820
I+II	III	0.810	0.652
I+III	II	0.720	0.761
I+II+III	–	0.740	–

I refers to the first 50 images; II refers to the second 50 images. Some faulty images had to be eliminated from the analysis, so the third set III contained only 46 images.

It is believed that the poor performance when training on I+III and I+II+III indicates that significant camera drift occurred by the time III was obtained, and that this explains the poor test performance when training on I+II and testing on III.

Table 5 Results of Tests on the Other Varieties of Wheat

	0	1	2	3	4	5
A	0.800	0.800	0.550	0.500	0.800	0.400
B	0.700	0.800	0.800	0.500	0.750	0.700
C	0.750	0.850	0.700	0.650	0.825	0.625
D	0.700	0.875	0.800	0.800	0.850	0.825
C'	0.800	0.900	0.750	0.700	0.875	0.675
E	0.800	0.900	0.800	0.800	0.875	0.825

Key: 0 = Mercia (uninfested and infested), 1 = Malacca (AU23), 2 = USA Northern Spring (AU20), 3 = Canadian, 4 = Shango, 5 = Australian White.

(a) Columns 0–5 refer respectively to the Mercia control and to the five alternative varieties of wheat.

(b) A and B present the results of the preliminary test on the six varieties, A using the standard parameter settings, and B using the narrow grain settings.

(c) C and D present the results of the 'bright spot' test on the six varieties, C using the standard parameter settings, and D using the narrow grain settings.

(d) In A and B the algorithm was trained on the control sample, for which uninfested and infested grains were available.

(e) In C and D the algorithm was trained individually on each of the six varieties, using uninfested grains and taking the uninfested grains with added bright spot as the infested grains. (The true infested grains available for the control sample were excluded from these tests.)

(f) In all cases the results were obtained by testing on the training set because of lack of data. Thus the results may be slightly higher than they should be: more realistic figures may be obtained by subtracting 0.025 from each figure. The figure of 0.025 was derived from Experiments (e) and (f) by comparing the accuracies from test sets on which training had been undertaken with those from test sets which had not been used for training.

(g) The bright spot was assigned an intensity of 20 in C and 10 in D, as the value 20 would have given absurdly high success rates in D. To provide a better comparison, the bright spot intensity values should have been adjusted to give the same results for the control sample in cases A and C (they already give the same values for B and D). C' shows the results for C modified in this way.

(h) Ideally, when examining the results in B and D, the control sample results should be taken as those in A and C, as the latter reflect the proper training for the wider grains of the control sample. (If C' is used instead of C, this problem does not arise.)

(i) Following notes (g) and (h), E shows the results obtained by taking the best parameter setting between C' and D. It is seen that all the varieties give similar results, and in no case is any worse than the (normalised) control sample: the available evidence indicates that the algorithm will perform well with all varieties of wheat

Table 6 Locations of the Bright Patch for Infested Grains, Experiment (h)

item	x	y	r	reliability	position
2	-8	1	8	low	
4	0	11	11	good	on larva
6	-8	6	10	low	
8	0	10	10	low	
10	0	-11	11	uncertain	
12	0	11	11	good	on larva
14	2	10	10	good	on larva
16	1	7	7	good	on cavity
18	-4	-10	10	good	on larva
20	6	-2	6	uncertain	
22	-1	-10	10	good	on cavity
24	6	-6	8	good	on larva
26	9	5	10	good	no relation
28	-5	-7	8	good	on larva
30	7	-2	7	uncertain	
32	-6	8	10	low	
34	-4	-9	9	good	on cavity
36	-2	9	9	uncertain	
38	6	-2	6	good	on larva
40	0	10	10	good	on cavity

In this table, the items are the infested grains, which are the even ones in the set of 40. x and y are the coordinates of the centre of the brightest patch found in the difference image, relative to the centroid of the grain. r is the distance from the centroid, accurate to the nearest pixel. Reliability was ascertained by checking visually whether the patch was prominent or likely to have arisen by chance placement of a few high intensity pixels. This final column indicates whether the bright patch corresponded to the position of the larva or the hole left behind by it, as seen in the X-ray images: this check was only made for the images in which good reliability of the bright patch was recorded.

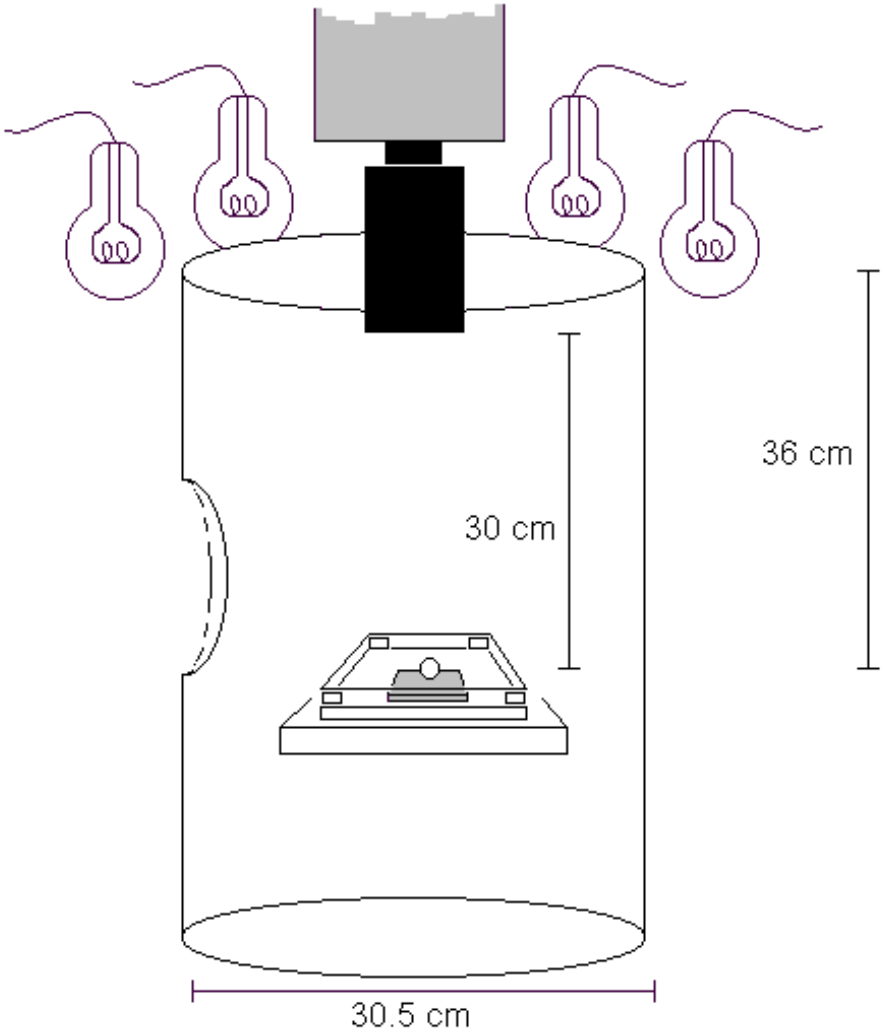

Figure 1: Image capture set-up, 1202 nm

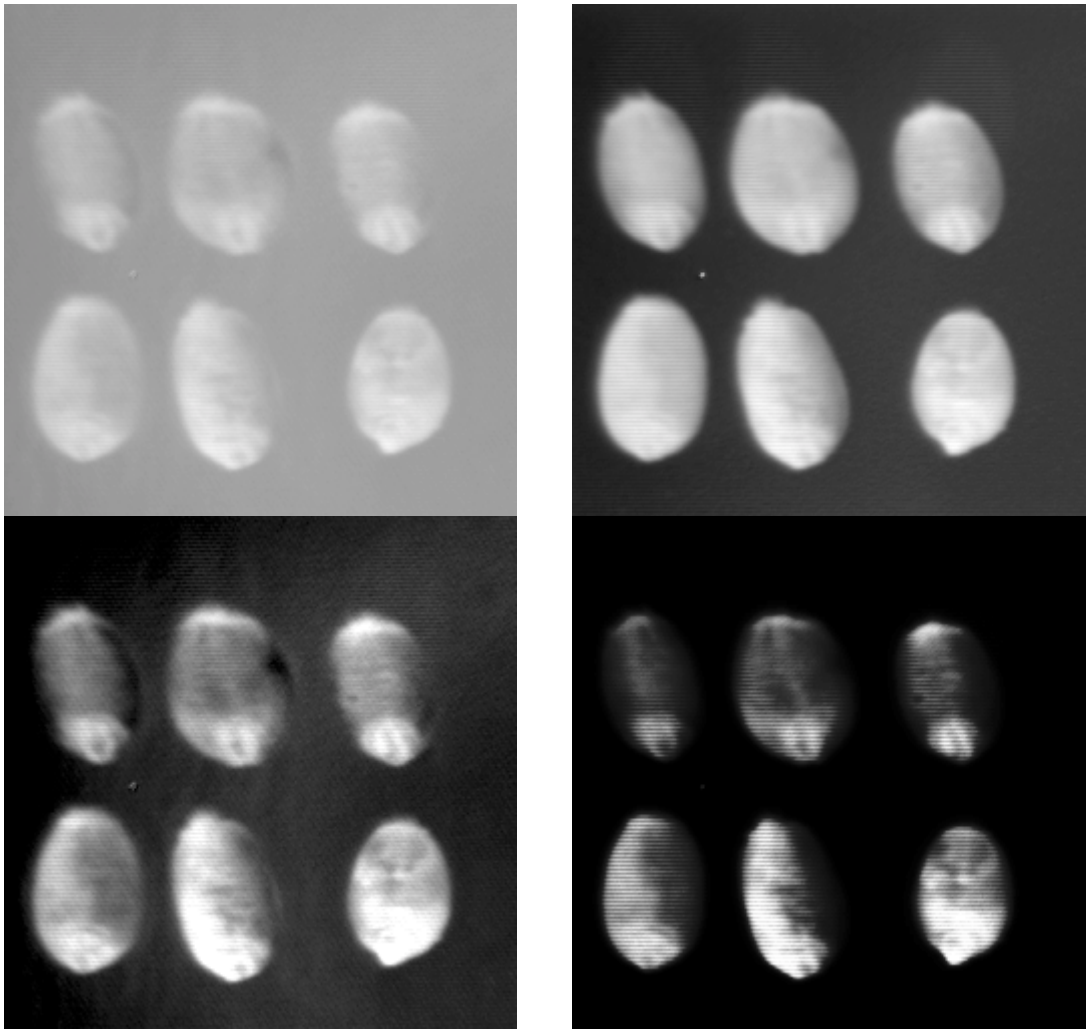

Figure 2: Example images from Experiment (a), to demonstrate image quality. Top row of each image: three uninfested kernels. Bottom row of each image: three infested kernels. Images, clockwise from top left: glass fibre paper background (background-subtracted), infrared blocking filter background (background-subtracted), glass fibre paper background (background-subtracted; contrast-enhanced), infrared blocking filter background (background-subtracted; contrast-enhanced)

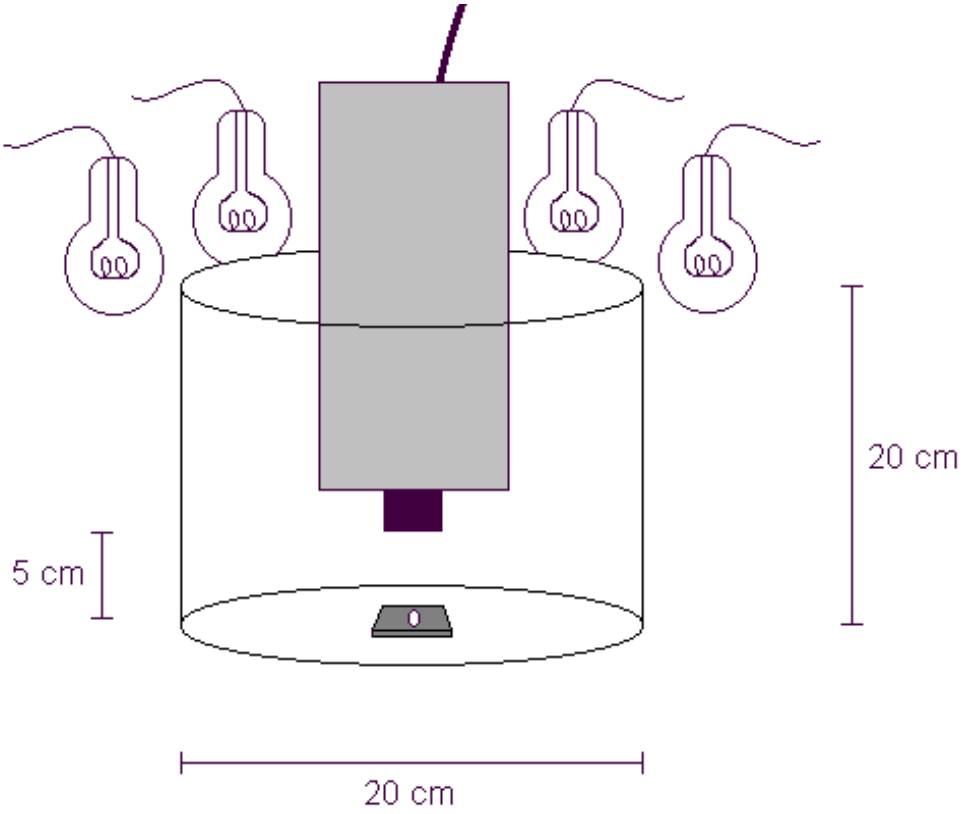

Figure 3: Image capture set-up, 981 nm

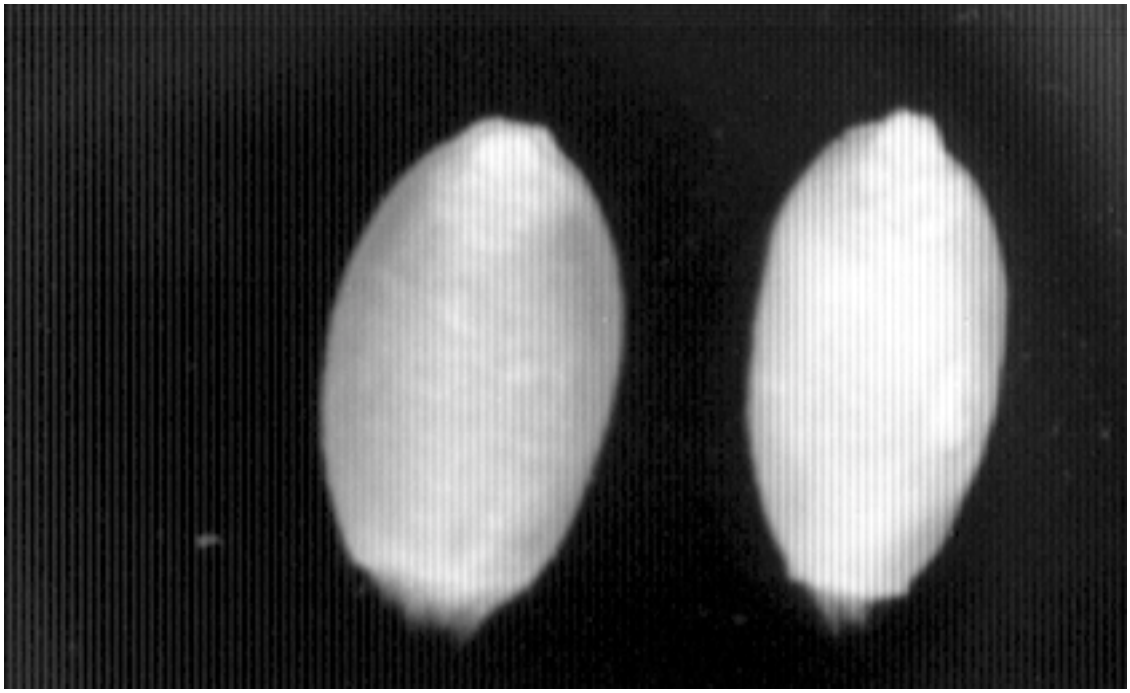

Figure 4: Image obtained from first sample [uninfested kernel (left) and infested kernel (right)], Experiment (c), as an example of image quality at 981 nm

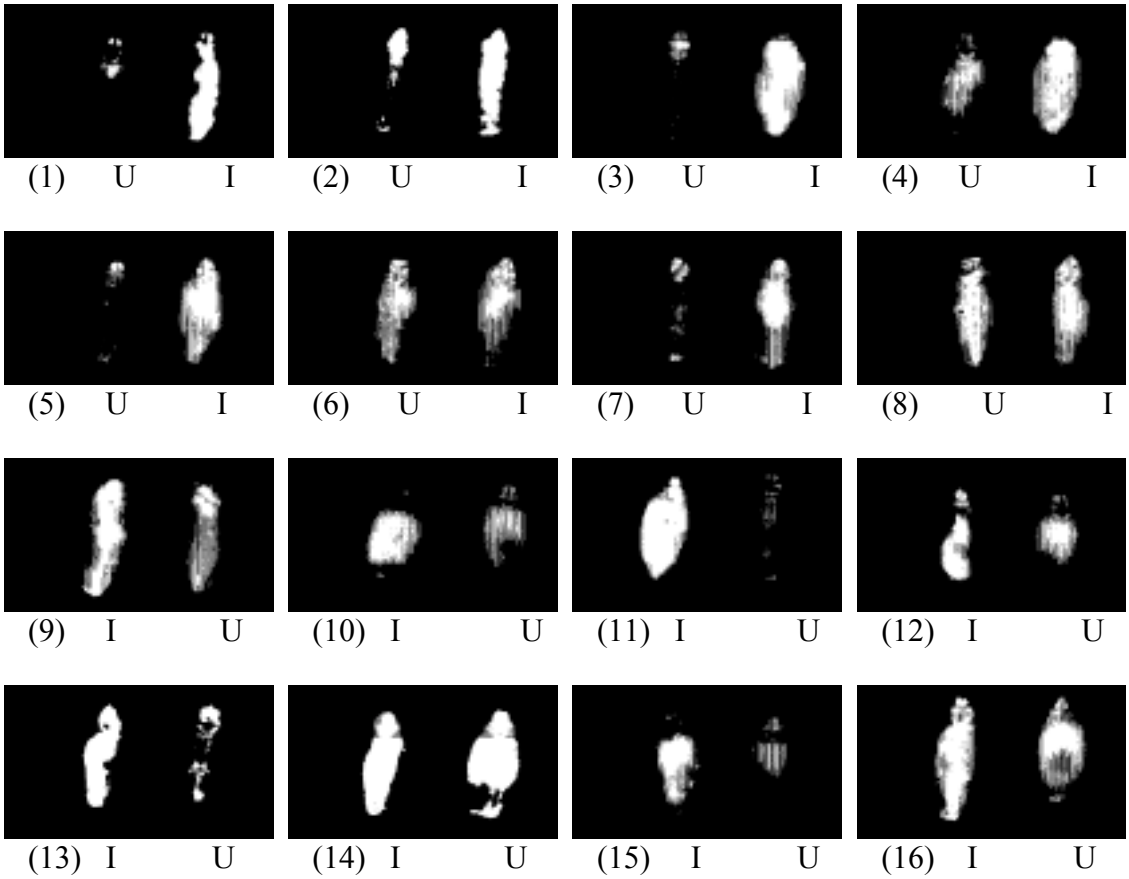

Figure 5: Thresholded images from Experiment (c). U=uninfested, I=infested. Image numbers are arbitrary. See text for discussion of individual images.

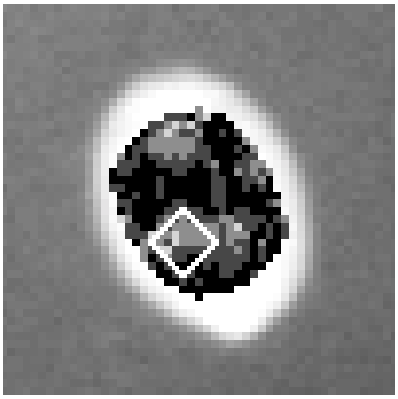

(1)

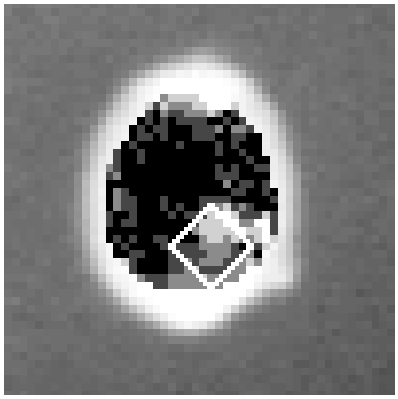

(2)

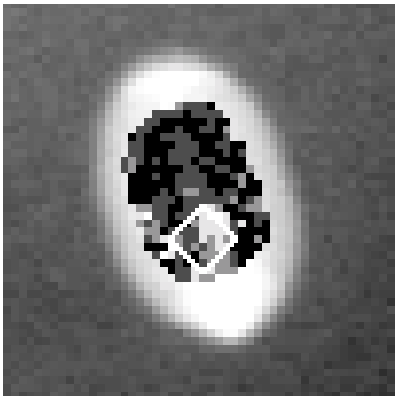

(3)

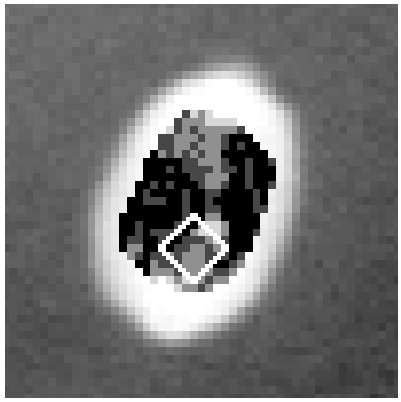

(4)

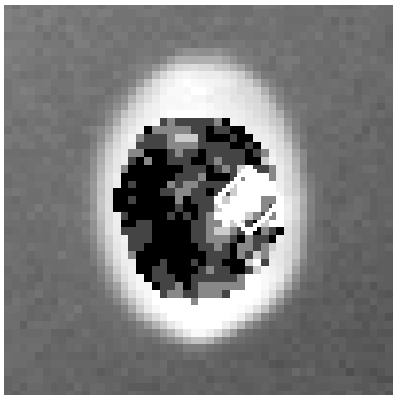

(5)

Figure 6. Samples of grains from Experiment (d) processed by the new algorithm: (1) a normal grain, correctly interpreted, (2) an infested grain, correctly interpreted, (3) a normal grain, incorrectly interpreted, (4) an infested grain, incorrectly interpreted and (5) an infested rather extreme grain, correctly interpreted. Note the shape and position of the automatically inserted exclusion zone for each grain. For clarity, the difference image is enhanced within the exclusion zone, while the part of the original image outside the exclusion zone has reduced brightness. The four dots arranged in the form of a diamond show the region of most likely infestation, irrespective of whether any infestation was present.