Automated inspection of highway structures – Stage 2

by S McRobbie, R Lodge and A Wright

Published Project Report
PPR255

Automated inspection of highway structures – Stage 2

by S McRobbie, R Lodge and A Wright

PPR 255

PUBLISHED PROJECT REPORT

TRL Limited

PUBLISHED PROJECT REPORT PPR 255

Automated inspection of highway structures – Stage 2

Version: 1.0

by S McRobbie, R Lodge and A Wright

Prepared for: Project Record: **Automatic Inspection of Highway Structures**

Client: **Transport Research Foundation (Dr R Kimber)**

This report has been prepared for the Transport Research Foundation. The views expressed are those of the authors and not necessarily those of the Transport Research Foundation.

Published Project Reports are written primarily for the Customer rather than for a general audience and are published with the Customer's approval.

	Approvals
Project Manager	
Quality Reviewed	

This report has been produced by TRL Limited, under/as part of a Contract placed by the Transport Research Foundation. Any views expressed are not necessarily those of the Transport Research Foundation.

CONTENTS

Executive Summary

Automated inspection of highway structures – Stage 2

The condition of highway structures is determined by visual inspection. There are four main types of inspection which are undertaken at different frequencies. These are Routine, General, Principal and Special Inspections. These inspections cover a range of detail, from a cursory check for obvious defects, through to a close examination of particular areas or defects causing concern. The quality of data collected depends on the ability of inspectors to observe and accurately record details on visible defects. It has been found that the data provided by such inspections can vary significantly. Improvements to the quality of the inspections are therefore desirable.

The research described in this report has been directed at developing a system for pre-processing the images prior to delivery to the engineers. The aim is that these pre-processed images should be able to give the engineer a rough idea of the condition of the structure, and draw his attention to those parts of the structure which contain most defects, or defect like features.

The research has investigated two main areas: image collection and display; and image analysis (both manual and automatic).

The image collection and display investigations have considered the use of multiple imaging positions, single imaging positions, spherical images, and mathematically transforming images to re-project them as if they were taken perpendicular to the face of interest in order to remove the effects of parallax. Recommendations are made regarding an appropriate realistic solution to the problem of obtaining images of sufficient quality and resolution without causing too many problems, either during the image collection process, or the subsequent processing and display.

The image analysis research began by investigating the possibility of using high quality images of structures to provide a consistent and quantitative assessment of condition. Initially this analysis was performed manually, and the results of this image based assessment were compared to the onsite assessments of the structure condition. Having shown that there was sufficient information and detail in the images to perform meaningful and useful image based condition assessments the focus of the research then switched to the use of automatic image processing techniques. The methods investigated in Stage 1 showed good promise, but were plagued by false positive reports of defects and features where none were present.

Subsequent image processing research has been focussed on reducing the incidence of these false positives, and delivering results to the engineers which make their job simpler, quicker and more cost effective.

Research up to the end of Stage 2 of this project has found that there is potential in the use of images for conducting offsite inspections. The best results are obtained using an initial segmentation method, set to give good confidence that areas reported as containing nothing of interest do contain nothing of interest, followed by the use of discriminant function to determine which of the remaining parts of the bridge contain a feature or defect which is of interest to the engineer.

It is recommended that further research should be carried out to refine the system, and optimise the discriminant function. This would enable the production of a very helpful semi-automatic system for use by engineers.

The system would also allow the comparison of successive years' images. Because the images would be taken from the same locations in the same way year after year then examination of successive years' images would enable any changes in defect severity or extent to be accurately tracked, giving a good idea of how quickly something might need attention.

A semi-automatic inspection system providing guidance, but still requiring an engineer to physically visit the site is the realistic end goal of this research.

1 Introduction

The condition of highway structures is determined by visual inspection. There are four main types of inspection, which are undertaken at different frequencies. These are:

- Routine inspection. This is a cursory check for obvious defects that might lead to accidents or high maintenance costs.
- General inspection. A visual examination of all parts of a structure that can be visually inspected from the ground and deck level. These surveys are carried out not more than two years after the last general or principal inspection.
- Principal Inspection. A close examination (within touching distance) of all inspectable parts of a structure. These surveys are typically carried out every 6 years. In recent years they have included limited further testing.
- Special Inspection. A close examination of particular areas or defects causing concern. These surveys are carried out to investigate a specific problem, either found during one of the above inspections or already discovered on other similar structures. They may include extensive testing.

These inspections cover a range of detail, from a cursory check for obvious defects, through to a close examination of particular areas or defects causing concern. The quality of data collected depends on the ability of inspectors to observe and accurately record details on visible defects. This could be affected by many factors, such as the environmental conditions, and the knowledge and experience of the inspectors. Possibly for these reasons, it has been found that the data provided by such inspections can vary significantly. Improvements to the quality of the inspections are therefore desirable.

The research carried out in this project has the objective of developing an improved inspection procedure, that is more objective and repeatable than current manual inspections. The improved inspection procedure is based on the proposal that images of structures could be collected and processed (manually or automatically), to identify defects in highway structures. The research has concentrated on two key areas in the development of this procedure - image collection and display, and image analysis.

This report reviews the work carried out in Stage 1 of this project and presents the results of Stage 2. In Stage 1 a simple approach to collect images of structures was demonstrated and it was shown, using a small dataset, that it should be possible to manually assess such images to obtain a measure of the visual condition of the structure. In stage 2, work continued into the development of an approach to collect and display images of highway structures, and manually examine them to identify areas of deterioration. This work is discussed in Sections 2 and 3 of this report. Although Stage 1 proposed that it should be possible to automatically analyse images to identify defects, Stage 2 has taken this forward to develop initial methods to automatically process images to identify deterioration. This is discussed in Sections 4 and 5. Finally, Section 6 presents the conclusions and recommendations following the first two stages of this project.

2 Image collection and display systems

The automated inspection of structures must start with the collection of images suitable for both manual inspection and automated analysis, because it follows that an image that is not suitable for manual inspection is unlikely to be suitable for automatic analysis. However, it is not easy to define a requirement for images for any structure (bridge), because bridges vary in size and style across the network. Therefore an investigation was carried out to review the requirements for the collection of images of structures.

2.1 Basic Requirement

The main requirement for the images may simply be stated as follows:

> *The images must provide enough detail to enable us to recognise the features that need to be identified.*

For automated inspection there is an additional requirement that the images should not contain artefacts or anomalies that would confuse analysis software. For example, the human brain can easily distinguish details within areas of variable illumination, whereas software will often see the transition between areas of different illumination as a “feature”.

The level of detail required from the images is a question that this research must address. At the very least we could assume that to inspect a bridge using image collection will require the collection of a number of images so that each face could be examined. However, if we assume that resolution at the millimetre level is required, even the use of multi-megapixel cameras would require many images to be collected for each face to show the detail required. In order to make it easy to inspect the structure from these images (either by eye or using automated software) it would be desirable to assemble the collected imaged into a form which can be readily analysed - for example, by arranging the images in such a way that they became one large image. However, parallax and perspective make this very difficult.

It should be noted that video camera technology does not provide the level of detail needed to be able to inspect a structure, except in a very superficial way. Therefore, throughout this section it is assumed that systems using one or more digital cameras will be used to obtain the images.

2.2 Fundamentals of optics

Parallax is defined as the apparent change in the position of an object resulting from a change in position of the observer, and manifests itself as a shift in the relative positions of objects at different distances as the observer moves. Perspective is the phenomena introduced when a three dimensional scene is translated to a two dimensional image as in photography. Perspective is caused by the size of the angle within the field of view that objects occupy and is a cue to the brain that gives the impression of distance. It is most apparent when two objects of similar size are present but are at different distances from the observer. This results in the farther of the two objects occupying a smaller angle within the field of view with the result that the size of that object appears smaller.

To illustrate the effects of parallax and perspective a series of computer generated images (CGI) of a brick wall are shown in Figure 1 and Figure 2. Figure 1 shows the lower part of the wall with the camera perpendicular to the wall both horizontally and vertically. Although the image in Figure 1 looks natural, the effects of both parallax and perspective are present but are acceptable to the brain because it is similar to how the scene would be seen through the eye.

Figure 1: Camera perpendicular to wall

Maintaining the same camera position but tilting the camera up 20° gives the image shown in the left of Figure 2 and the effect of perspective, in the form of ‘converging verticals’, is very apparent.

The camera used in the right of Figure 2 is again perpendicular to the wall but the camera has been moved laterally to the left and shows the effect of parallax. Whereas in Figure 1 the right face of the left buttress is visible, in the right of Figure 2 only the left face is visible.

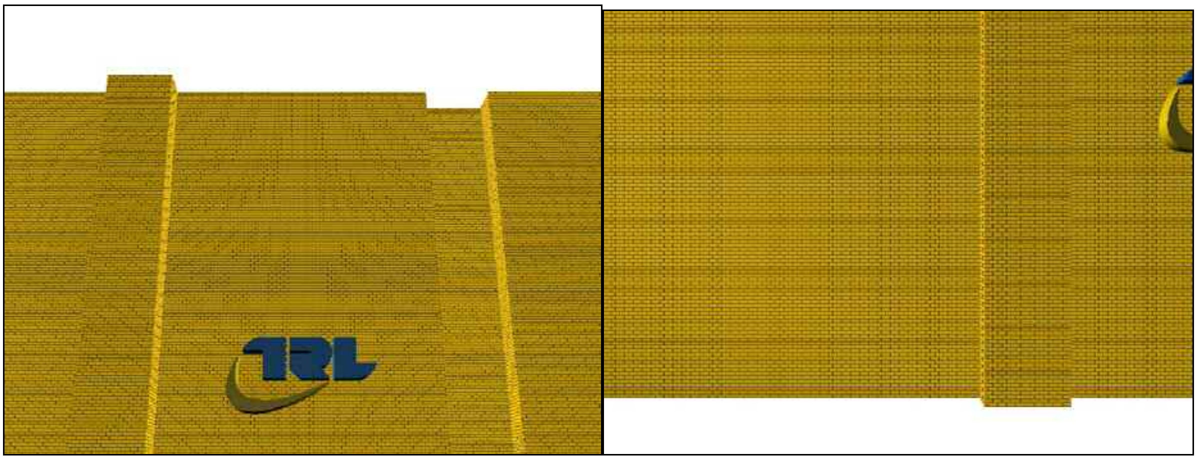

Figure 2: Camera tilted up (left image) and camera offset to left (right image)

2.3 Review of Stage 1 investigation of image collection

In stage 1, examination was undertaken of a series of images of a bridge abutment collected at equally spaced locations parallel to the abutment. At each location six pictures were taken from the roadside in the form of two columns of three images with one column to the left and the other to the right of the horizontal perpendicular to the wall. Adjacent images overlapped each other so that common points could be used for alignment. This provided a sequence of images that covered the entire abutment (this method of collection was used in the development of image analysis – see Sections 3 and 4). The objective was to determine if such a method of data collection could be used to obtain a single large image from the collage of smaller images.

Figure 3 shows the result of using twelve of the images of the bridge wall (6 from each of two locations) to create a single stitched image. This was achieved using commercially available software but took more than three hours of manual intervention to assist the software in identifying the corresponding points on adjacent images, so that the software could warp the images to align them before stitching them together. Inspection of the resulting image shows that images taken from the same location (i.e. the 6 that form the left side or 6 on the right) provide a reasonable image. In maintaining the same camera location (but twisting and tilting for each image) the same viewpoint is maintained, and hence perspective and parallax in each image are similar, making it easier for the software to the images together. However, moving the camera location means that the effects of parallax and perspective play a significant part in the image make up, making it impossible to stitch the images. This is clearly demonstrated through observation of the joints present in the soffit.

Figure 3: Single image from twelve overlapping images taken from two positions

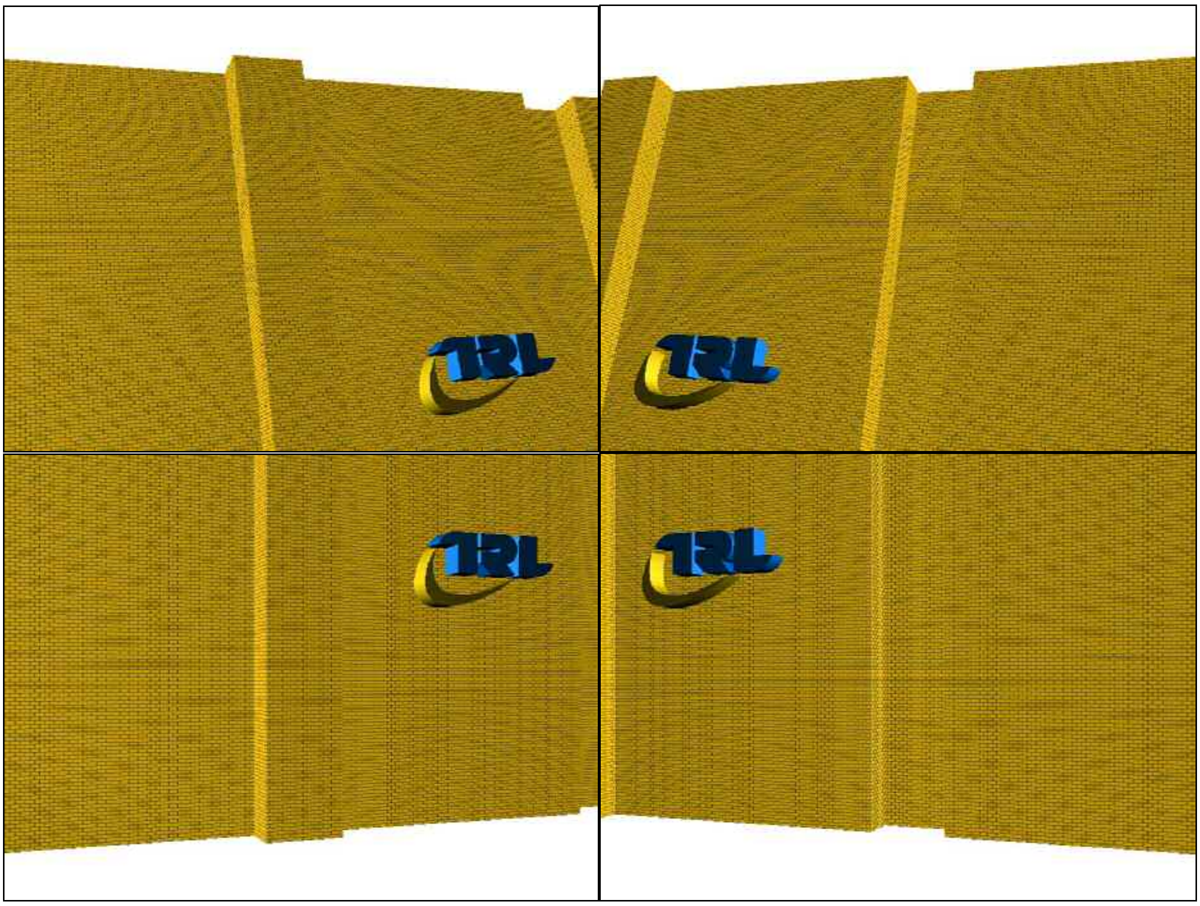

Figure 4: Four images taken from two positions with deliberate overlap

The four CGI images shown in Figure 4 give a clear indication of how difficult creating a single image from images taken at multiple camera positions and orientations would be. The lower pair of images in Figure 4 were taken horizontally with, in each case, the camera pointing 20° towards the centre of the wall. The upper pair are again taken pointing towards the centre, but are also tilted up 20°. Unlike the images taken of the real bridge, the camera location and orientation for the CGI images is known very precisely, and with this knowledge it would be mathematically possible to create a single image from the multiple images. However, when taking images of a bridge, for example, it would be a non-trivial task to measure the location and orientation of the camera to the accuracy needed to successfully stitch the images together.

2.3.1 Correcting for perspective during image collection

An investigation was undertaken to review the available optical techniques to minimise the effects introduced by parallax and perspective. A tilt and shift lens is a special type of lens that can remove the effect of perspective. It is often used in commercial photography of building facades, to remove the converging verticals when the camera needs to be tilted so as to include the whole of the building. However images collected with a tilt and shift lens were found to provide very limited improvement in the automated inspection of structures. Firstly, the amount of tilt that could be corrected was limited, with the result that it would not be possible to use the tilt adjustment at the lens to correct for all of the camera tilt needed to image the whole of a surface. Secondly, even if the tilt adjustment was large enough, it is not practical to set the amount of tilt needed for each image with the level of accuracy needed to ensure that the effects of perspective are eliminated.

Having established the issues of practicality and time required to join multiple images taken from different view points the remainder of the investigation focused on spherical images.

2.3.2 Spherical images

The term spherical image is slightly misleading in that the images are not spherical. However, the technique for obtaining and viewing the images can be imagined as creating and viewing images from a single point in space that is at the centre of a sphere with the image being on the inside surface of the sphere. The value of the spherical image is that it can be used to provide the user with a much more "realistic" viewpoint on the structure. This is because, following the creation of the spherical image, software can be used that enables the user to view the whole scene as though from a single point. The software allows the view to be panned left and right, tilted up and down as well as being able to zoom in on an area of interest – just like being present at the structure.

The nodal point of a lens is the theoretical point through which all light entering the camera passes. If a fixed optical path through the lens is assumed then all images, taken such that the nodal point of the lens remains stationary in space with the camera's orientation pivoting about this point, will possess the same perspective and will not show any effects from parallax. Images obtained in this way are required for spherical images. To create a spherical image first requires that images are obtained to cover the entire surface of the theoretical sphere with adjacent images overlapping each other by approximately 30% on each side. The images are then warped and aligned to join them together to form a rectangular image in much the same way that a map of the earth can be projected onto a piece of paper. An example of a spherical image of a bridge is shown in Figure 5. Just as a map of the earth shows a distorted image of the earth's surface when compared to the view seen from space, this image is also distorted but it is still a recognisable and useful image.

The quality of stitching for spherical images is very dependant on two functions. Firstly, the lens' optical path must be the same for all images. Changing the focus and/or the aperture can change the optical path measurably. The second function that affects the quality is the accuracy with which the camera is set up and can be moved about the nodal point.

Figure 5: A Spherical image viewed as a normal photograph

The image in Figure 5 was created from 59 photographs, each taken from a single location underneath the centre of the bridge with the camera progressively rotated horizontally through 360° and vertically through 180° about the nodal point. The photographs of the ground below the camera were excluded from the image - this is the cause of the black area at the base of the image. However, this area needs to be included in the image because the software that is used to view the image as a sphere assumes that the image covers the full lateral and vertical rotation about the nodal point.

Each photograph used to create the spherical image in Figure 5 was an 8-megapixel image and the resultant image is 153-megapixels. Commercially available software was used to create the spherical image. More than 10 hours of manual assistance was required to align the images to obtain Figure 5.

The investigation of the use of spherical images enabled us to conclude that such images would provides a significant step forward in the ability to view a structure in detail, albeit from a single point. Clearly it would be necessary to have multiple spherical images in order to ensure that surfaces hidden from one viewpoint are available from a different viewpoint. The spherical image would not only provide current status information about the structure but can be archived for future reference as an aid to monitoring deterioration over time. However, there are a number of challenges to be addressed before the spherical image becomes an everyday tool for structure inspection. In general terms these challenges fall into one of two interdependent areas:

- Collection of images;
- Creating and viewing the spherical image;

2.3.2.1 Collection of spherical images

The images collected for use in the creation of a spherical image not only have to meet the requirement that they are rotated about the nodal point but they need to have a similar level of exposure. The effects of irregular exposure can be seen in Figure 5, which was built from images taken on a sunny but windy day, with moving clouds. The majority of images were taken in sunlight, but some of the images were taken when the sun was obscured. Keeping the optical path constant meant that the exposure was constant for all images, with the result that some images were dark. This can be seen as the very significant dark area on the underside of the bridge deck just to the right of the centre of the image. It may be possible to address the exposure problems using artificial lighting to control the illumination at night.

Given the fact that the collection of spherical images requires taking many (more than one hundred) images from the same location, oriented about the nodal point it would be desirable to automate this process. This would simplify collection and make the system repeatable. This repeatability would have a positive effect on the ability to view and stitch the images. Also, automation would reduce the time taken to obtain images and minimise the disruption to road users. Manufacturers of factory robots, similar to those used for spraying cars, claim them to a high level of position repeatability.

2.3.2.2 Creating and viewing the spherical image

Commercially available software with many hours of manual intervention was used to create the spherical image of Figure 5. For a practical system the amount of user intervention must be minimal. However, a larger number of pixels than those used for this report are likely to be needed (in order to get the detail) and the amount of user intervention is likely to increase following a square law (e.g. twice as many pixels will require four times the intervention). The high number of pixels will also present problems for the viewing software. For example, the spherical image viewer used for this investigation to view the 153 mega-pixel (17500 x 8750) image was unable to display a 288 mega-pixel (24000 x 12000) image.

As discussed above automation would provide a repeatable and accurate positioning system for the camera’s orientation about the nodal point. This may enable calibration of the equipment so that the warping needed to align individual images is known. Although, it is unlikely that the images could be stitched based solely on the calibrated values, knowledge of the relative positions of the images would be useful in improving the automatic recognition of the same location on adjacent images and, therefore, reducing the amount of manual intervention required.

The review of image collection carried out in Stage 1 therefore concluded that the use of spherical images could be a powerful technique for collection and display of images of structures, but there is a need for considerable research and development of this technique to make it a simple and practical method for the inspection of structures.

2.4 Development of image collection in Stage 2

Although the above work has shown that viewing a spherical image provides an intuitive way of inspecting a structure, a number of obstacles make this solution less than satisfactory for regular inspections (in the short term). In particular, the conversion of the individual images into a single spherical image requires a significant amount of manual intervention. However, a further significant issue is the suitability of the images for the (automated) identification of features. Spherical images distort the original images in such a way that detail is lost non-uniformly (i.e. dependant on where the image is to be used within the spherical image). Therefore, in stage 2 it was proposed that an alternative approach be attempted, that retains the advantage of image collection from a single viewpoint (as used in collecting spherical images), but uses the images in such a way that the limitations of “joining” and “distortion” are removed.

The process of collecting a series of overlapping images from a single viewpoint (as used when creating a spherical image) has a distinct advantage in that the relative geometry between the images is well defined and it does not suffer from parallax. These images have a common viewpoint and, providing the camera’s optical path remains unchanged (e.g. no re-focusing or changes to aperture), they have a fixed field of view. By using mechanical indexing of the camera’s orientation, both horizontally and vertically, the pan and tilt of the camera between successive images can be monitored. This information can be used to re-project images such that they appear to have been taken from the same viewpoint, but with the camera at a different orientation to its actual orientation. Since all of the images have been taken from the same viewpoint, but with the camera at different orientations, some images will contain part of the scene that can be re-projected to appear as though it was taken as part of a much larger image with the camera at a different orientation. In this way several images can be re-projected to form a collage. This re-projection approach was investigated further, as discussed in the following section.

2.4.1 Re-projection

An image is the projection of the view of a three dimensional scene, as seen through the lens, onto a flat plane that is perpendicular to the axis of the lens. As the camera is panned and tilted the photographic plane is also panned and tilted. The result of this movement is that individual points in the viewed scene appear to change their positions relative to each other between two images taken from the same viewpoint but with a different camera orientation, in a way that is commonly referred to as the effect of perspective.

It can be shown that a given pixel on an image contains the colour of the point that is both the nearest point to the camera on the nearest object to the camera AND intersects an infinitely long vector that passes through the “node point” of the lens and the point on the image sensor that corresponds to the point on the image. For spherical images, each image has the same optical arrangement as the others in the set. The image sensor has constant dimensions and all images are taken with a common ‘node point’. Any image can be re-projected onto any other image plane providing the relative pan and tilt between the images, and optical geometry are known. Figure 6, shows two images taken from a set of spherical images with Image 2 being horizontal with (i.e. no tilt), and panned 20° to the right of, Image 1. Image 2 clearly shows perspective, in that the horizontal lines appear to converge at some point to the right of the image.

With the knowledge that Image 1 and Image 2 share a common “node point”, both have a horizontal field of view of 29.6°and the relative pan and tilt between the two images is 20° and 0° respectively, we can re-project Image 2 onto the same image plane as Image 1 and produce a re-projected Image 2. Software was developed in Stage 2 to calculate re-projected images. The output from this software when applied to the images of Figure 6 is shown in Figure 7.

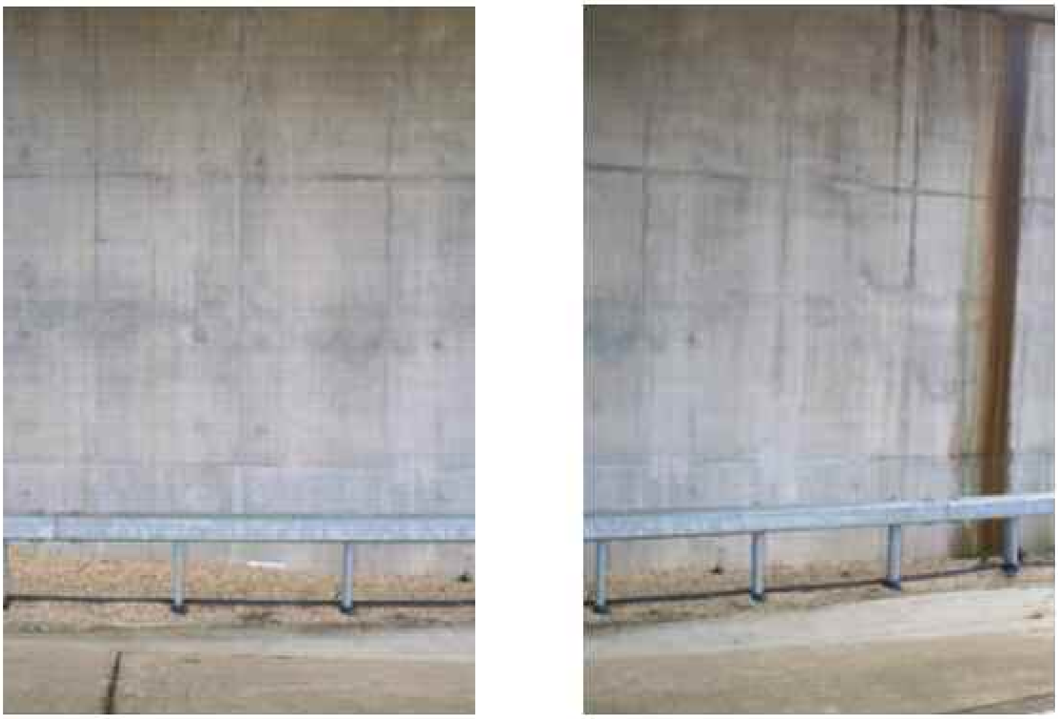

Image 1 Image 2

Figure 6: **Overlapping images from the same viewpoint – Image 2 shows perspective**

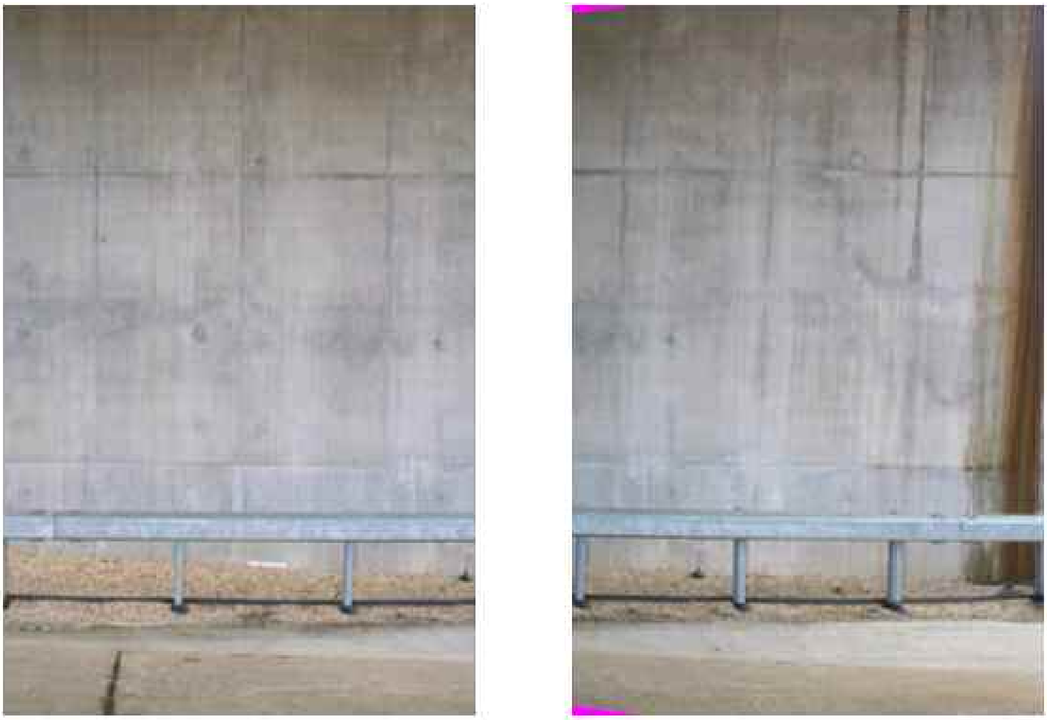

Image 1 Image 2

Figure 7: Same images as in Figure 6, but with Image 2 now re-projected onto same plane as Image 1

The result of the re-projection process is that spherical aberrations are removed, delivering images that all appear to lie on the same plane. This is an ideal approach for structures, which tend to contain planes relating the individual components (e.g. the soffits). By re-projecting all the images from a plane we can obtain an un-distorted image set that can be traversed in two dimensions by the inspector.

A limitation of this approach is that the relative orientation between the image sensor and the re-projected image determines the pixel resolution available within the re-projected image. When the image is projected onto a flat sensor, the angle subtended by two adjacent pixels becomes smaller for pixels far from the centre of the sensor. This effect also applies to the re-projection of the image, because pixels further from the centre have a smaller subtended angle with their adjacent pixels. Therefore, the angle between a pixel and its neighbours on the original image may be significantly different from that required for the equivalent pixel on the re-projected image. The result of this is that not all pixels in the re-projected image will have a exactly corresponding pixel in the original image. This implies that the original image must have a higher resolution that that required in the re-projected image, if all re-projected image pixels are to contain valid image data. Therefore, to minimise the effect of this, the images in Figure 7 have been created at a pixel scale of 1:2 relative to the original image. The scaling has removed missing pixels within the bulk of the re-projected image, but some still remain. The magenta wedges that appear in the top and bottom left corners of the re-projected Image 2 in Figure 7 show areas of the re-projected image for which no equivalent pixel is available in the original image because of the warping needed to change the projection.

2.4.2 Re-projection summary

Given the complications presented in the use of spherical images an alternative approach has been investigated and developed in Stage 2 of this project that applies a similar image collection method to that applied in the collection of spherical images, but uses automated methods to re-project the images onto a single plane suitable for subsequent analysis. This is referred to as re-projection. This approach is particularly suitable where the face of a structure is flat, because the images collected from that face are re-projected onto a plane that is parallel to the face's plane.

To obtain the images for re-projection firstly requires images taken from a fixed node point, with fixed optical (lens/sensor) geometry and known values for the pan and tilt. Also, the images need to have sufficient detail within them so that they can be re-projected at the resolution needed for analysis, but without any missing pixels. The need for 'sufficient detail' relates to the angles subtended by adjacent pixels in both the image taken and the re-projected image. Using a long focal length lens (i.e. small field of view) requires that more images are taken but ensures that more detail is available for re-projection. Also, placing the camera as far away as practical from the object of interest reduces the field of view of the collage, which in turn reduces the effect.

For practical image collection it would be appropriate to instrument the camera with equipment to automatically record the angle of the camera for each image. This data would be used in the re-projection process. Furthermore, on flat-faced structures it is suggested that a laser distance measure could be attached to the camera and aligned with the node point. This would be used to establish the distance of each image from the face. These distances would then enable us to calculate the perpendicular distance of the node point from the face, and from this the size of defects at the face could be calculated.

3 Assessing structures using image collection – manual surveys

As discussed above, this project aims to determine whether surveys of structures can be undertaken using image collection and subsequent assessment of the image data. Section 2 has discussed the approaches that were considered in Stages 1 and 2 for the collection and display of images. In this section we shall review the work carried out in determining whether the images can be used to deliver an acceptable assessment of condition. Specifically, this section presents the results of the first part of this work, which had the objective of determining (via manual analysis) whether the images contain sufficient information to provide a measure of condition comparable to those obtained by onsite inspections. The automated inspection of images is discussed in Section 4.

3.1 Image collection

The bridge shown in Figure 8 (maps showing bridge location are shown in Figure 9) was selected, inspected and photographed to enable comparisons to be made between the condition of the bridge as determined on site, and the condition determined by viewing the images. Images were collected using a 6 mega pixel digital camera, mounted on a tripod, using the Stage 1 approach described in Section 2.3. Each image covered an area of bridge approximately 2m high by 3m wide, meaning that each pixel was approximately 1mm square. The external walls of the structure were photographed, using the same camera mounted on a tripod, from a single point on the side of the road, approximately 10m away from the structure. During image collection the location of each image relative to the face being imaged was recorded (e.g. x,y position relative to the left bottom point on the face). This enabled the images to be collected together for display and manual analysis.

Figure 8: Structure at Winnersh (M4 over B3030).

Unfortunately, the problems associated with perspective and parallax resulted in the images of the soffit being unusable (images of the soffit are included in Appendix A to demonstrate these problems). However, the images of the internal walls were considered acceptable for this investigation. Figure 10 shows an example of an image of the west internal wall. Close-ups of the boxes highlighted in Figure 10 are shown in Figure 11 (red box) and Figure 12 (yellow box). The close-up in Figure 11 shows highlights the level of detail available within these images, and Figure 12 shows the maximum level of detail available in the images used in this investigation.

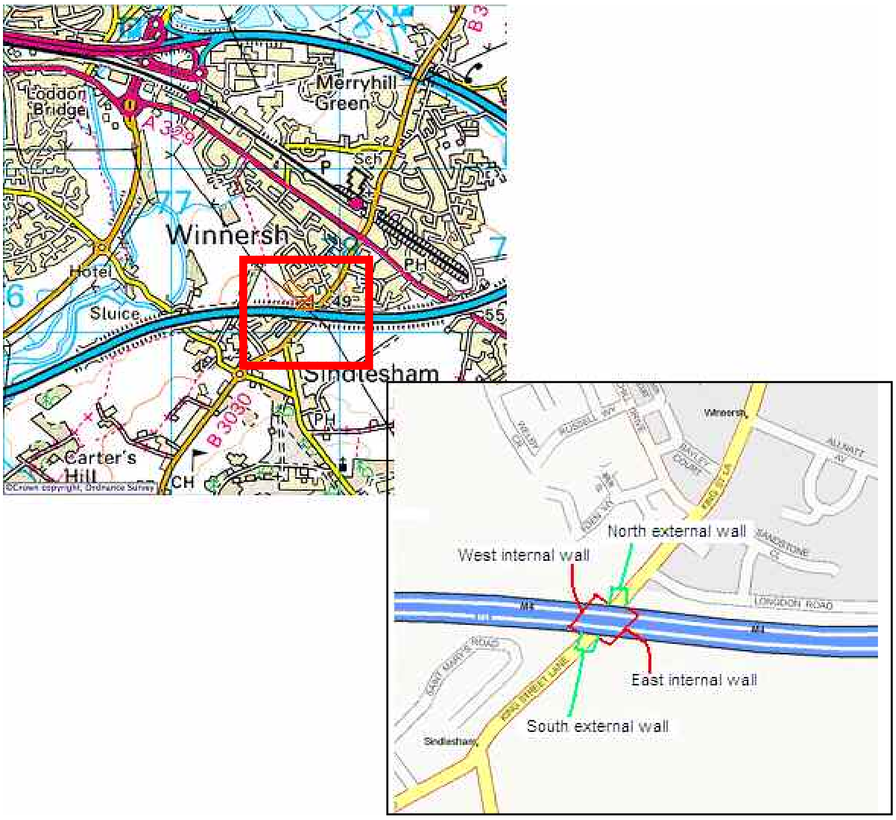

Figure 9: Map of Winnersh site (M4 over B3030)

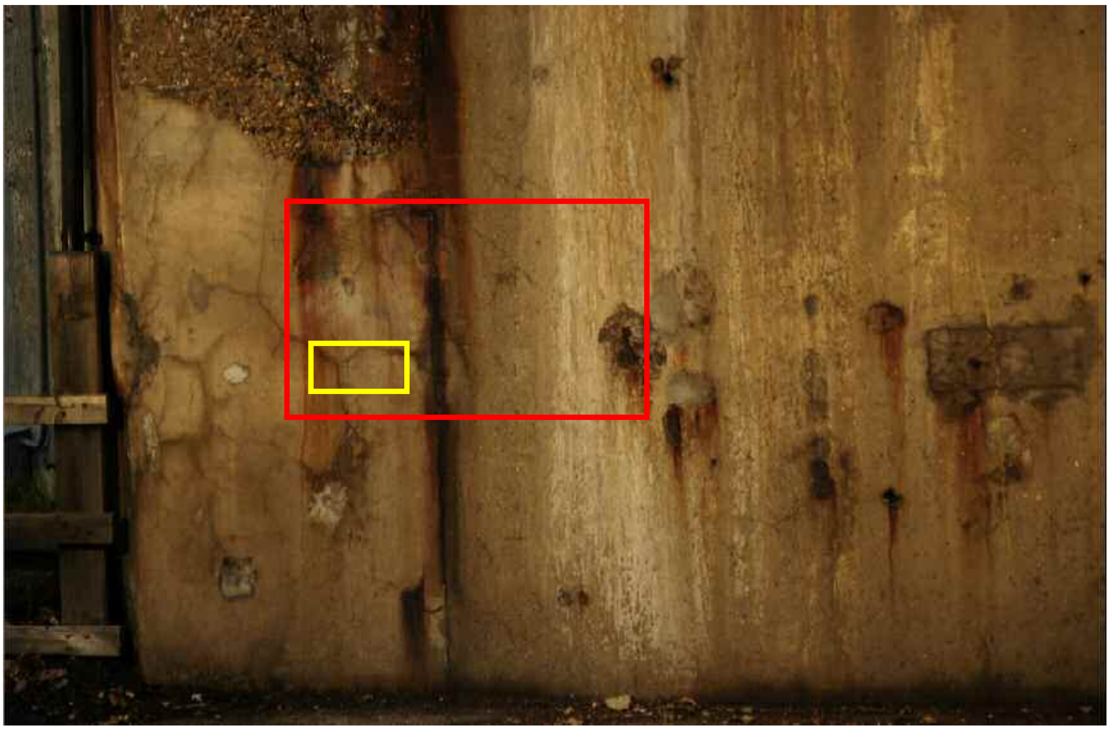

Figure 10: Full image of part of west internal wall

Figure 11: Zoomed in image of region highlighted in red in Figure 10.

Figure 12: Zoomed in image of region highlighted in yellow in Figure 10 and Figure 11, showing maximum level of detail achievable without distortion.

Figure 13 shows a composite image made up of all the individual photographs taken of the structure. The images are collected together in this mosaic using the location data recorded during the collection of the imaged. The composite image is a useful way of gauging the overall condition of the structure as it allows the whole structure or surface to be seen at once and considered as a whole. The white dots in Figure 13 are an artefact of the image display software (used in generating the mosaic of images).

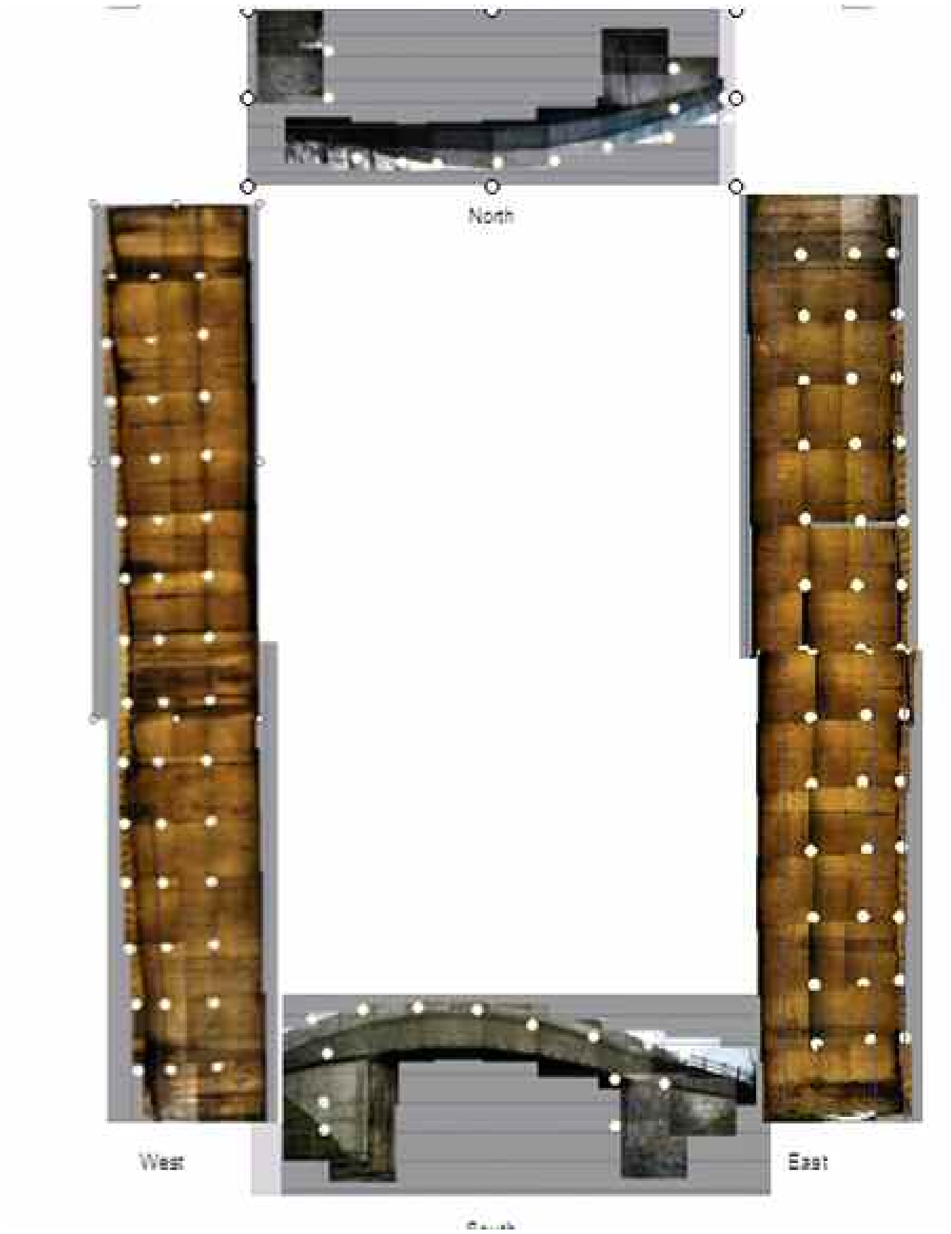

Figure 13: Composite image of whole structure

3.2 Onsite assessment

An onsite assessment of the bridge was carried out to determine the nature, severity and extent of any defects present, and to note the locations of these defects. The assessment was only concerned with the visible condition of the structure.

This assessment was carried out at two levels of detail providing a **general assessment** of the condition of the bridge, and the individual elements comprising the bridge, and an in depth quantitative **detailed assessment** giving details of any individual defects observed.

General assessment:

The general assessment of the structure was performed to determine the overall condition of structure. The condition of each component (west internal wall, east internal wall, north external wall, south external wall and soffit) was also recorded, using a system developed for this work, detailed in Table 1.

Category	Description
6 – Sound	No maintenance needed at moment, and no visible signs of degradation which will lead to maintenance within 5 years
5 – Good	No maintenance needed at moment, but one or two possible signs of degradation which may lead to maintenance being required within 5 years
4 – OK	No maintenance needed at moment, but several possible signs of degradation which may lead to maintenance being required within 5 years
3 – Poor	No maintenance needed at moment, but signs of degradation suggesting maintenance needed within 2 years
2 – Very Poor	Single area requires attention within 6 months
1 – Deteriorated	Multiple areas require attention within 6 months
0 – Severely Deteriorated	Requires attention urgently

Table 1: General condition assessment categories

Detailed assessment of defects:

A more detailed assessment of the condition of the bridge and its separate components was performed by noting the positions and extents of a specified range of defects. The locations of the following defects were recorded, and the defects were photographed for reference:

- Cracks
    - The number, nature and extent of cracking were recorded in as much detail as was practical.
- Wet or damp surfaces
    - The positions of any wet or damp surfaces were noted, along with any relevant information the inspector could give.
- Spalling
    - The positions of any areas of spalling were noted, along with any relevant information the inspector can give, especially where the underlying reinforcement is visible.

- Rust
  - The positions of any areas where rust or signs of rust are visible were noted, along with any relevant information the inspector can give. This included the engineer's opinion on whether the rust is superficial or a sign of a potential problem.

- Other obvious defects or features
  - Anything else which the inspector, in their professional opinion, felt to be worth recording was noted down with relevant measurements and locations.

The codes shown in Table 2 and Table 3 were used to describe the Extent and Severity levels respectively of each defect, where:

Extent:	The area, length or number (as appropriate) of the bridge element affected by the defect/damage.
Severity:	The degree to which the defect/damage affects the function of the element or other elements on the bridge.

Code	Description
A	No significant defect
B	Slight, not more than 5% of surface area/length/number
C	Moderate, 5% - 20% of surface area/length/number
D	Wide: 20% - 50% of surface area/length/number
E	Extensive, more than 50% of surface area/length/number

Table 2: Extent Codes

Code	Description
1	As new condition or defect has no significant effect on the element (visually or functionally).
2	Early signs of deterioration, minor defect/damage, no reduction in functionality of element.
3	Moderate defect/damage, some loss of functionality could be expected
4	Severe defect/damage, significant loss of functionality and/or element is close to failure/collapse
5	The element is non-functional/failed

Table 3: Generic Severity Descriptions

3.2.1 Image assessment procedure

General assessment

The overall condition of the structure as a whole, and of its individual components, were assessed using the image display software. Where possible and appropriate this general assessment used the same methods and scales (see Table 1) as were used in the onsite reference inspection. This delivered a condition rating (0-6) for each component which was directly comparable with that obtained in the onsite inspection (Section 3.2). This assessment was performed without zooming in too closely to the images - the display level for this assessment was set such that it showed at least 2m in the vertical direction at all times.

Detailed quantitative assessment:

A more quantitative assessment of the bridge condition was made using the viewing software. This inspection involved marking the images where any of the defects or features listed in Table 4 were visible. This created a data file referred to as a defect "map" containing the locations, and extents of the various defects. This detailed assessment was performed with a higher level of zoom than the previous assessment. The display level for this assessment was set such that it showed between 1m and 2m in the vertical direction at all times on a 15 inch monitor.

Edge of Structure being considered
Cracking
Spalling
Rust Staining
Wet/damp surfaces
Delamination
Patch Repairs
Construction Joints
Previous examinations - cores / powder drillings
Other surface features (conduits, etc.)
Obscured / not visible

Table 4: Defects / features recorded in detailed image based inspection

The results of this assessment, and the notes made by the inspector while performing it, were also used to produce a condition score for the structure and its component parts using the scales defined in Table 2 and Table 3. These allow codes to be assigned to each surface, and the overall structure which show how much of each surface is affected by any defect (the extent) and how serious the defects are (the severity).

3.3 Comparison of onsite surveys with manual inspection of images

General assessment

The results of the general assessments made onsite and using images were compared to see whether or not both inspection methods produced similar overall impressions of the condition of the bridge.

The general condition of the structure, and of the individual components, were assessed using the scale shown in Table 1. In this scale a structure found to be in 'Sound' condition was represented by a score of 6, and 'Severely Deteriorated' surfaces or structures were denoted by a score of 0. Table 5 shows the results of the onsite inspection and image inspection. We can see that in general the differences between the assessment of the bridge and its component parts made onsite, and by looking at images are very small – on average there is a difference of less than 1 in the reported condition.

Face	Onsite	Images	Difference
North External	5	5	0
South External	4	3	1
East Internal	5	5/6	0/1
West Internal	4 (2 at drainage)	4	0/2
Overall	4	5	1

Table 5: General condition assessment results

The largest difference seen is on the west Internal wall of the bridge. Here the onsite survey made two separate assessments, one for the overall condition of the surface, and a second for the drainage system at the south west corner. The drainage was seen to be in 'Very Poor' condition, and should, in the inspectors opinion, be replaced within 6 months to prevent further deterioration to the structure. The rest of the surface was thought to be in 'OK' condition, an opinion which was shared by the image assessment.

Overall, we can see in Table 5, the onsite inspection thought the bridge condition was OK, while the image based assessment rated it slightly better, as being in Good condition. The difference between OK, and Good condition assessments as laid out in Table 1 is that a Good structure has no signs of maintenance being required at the moment, but one or two signs that some maintenance may be needed within 5 years; an OK surface has no need for maintenance at the moment, but exhibits several signs of maintenance perhaps being needed within 5 years.

Detailed quantitative assessment:

Comparisons of the results of the detailed inspections carried out onsite with those carried out remotely (using the manual analysis of images) were performed by transforming both the onsite inspection results and the image based inspection results into the same format. This format was based on a grid, where the structure was split into an array of 200mm square boxes. Each box was assigned a code number depending on the defect present within it:

- For the onsite assessment this involved interpreting the notes and photographs taken during the inspection, and assuming that if a particular part of the bridge was not mentioned, then it must have been sound, and contained no defects (hence assigning a value of zero to the grid square). Unfortunately the notes recorded by the onsite inspector did not directly provide scores for the condition. Therefore interpretation of the inspection results was carried out, along with review of the photographs taken during the inspection, to produced the set of scores, shown in Table 6.
- For the images, software was produced which enabled the images to be displayed on the screen of a PC and an assessor recorded the presence of defects within any particular grid square using a mouse

As can be seen in Table 6, the reported condition of most of the surfaces, and the overall assessment of the condition of the bridge, differs between the two inspection methods. However, closer inspection shows that the majority of the differences are quite small. The differences are all (apart from the drainage system discussed above) associated with the extent rather than the severity of the defective areas, with the image based assessment tending to report a slightly larger area affected by the defect. This may be because the image based assessment provided information for the whole of the structure, whereas the onsite assessment tended to only report the more serious and obvious defects, and did not report every crack, or the full extent of some defects.

Face	Onsite	Images
North External	B2	C2
South External	C3	D3
East Internal	C2	B2
West Internal	C3 (E4 at drainage)	C3
Overall	B2	C2

Table 6: More detailed quantitative condition assessment results

Figure 14 shows the results of the onsite inspection after being converted into a defect map. The grey lines show the outline of the structure being surveyed (this outline has been directly copied from the image based assessment so as to enable a like-for-like comparison), and any obscured areas of structure. Green lines correspond to features which were recorded in the inspection, but which do not relate to defects (e.g. cables). Other features, such as cracking, spalling or rust are shown in redder tones.

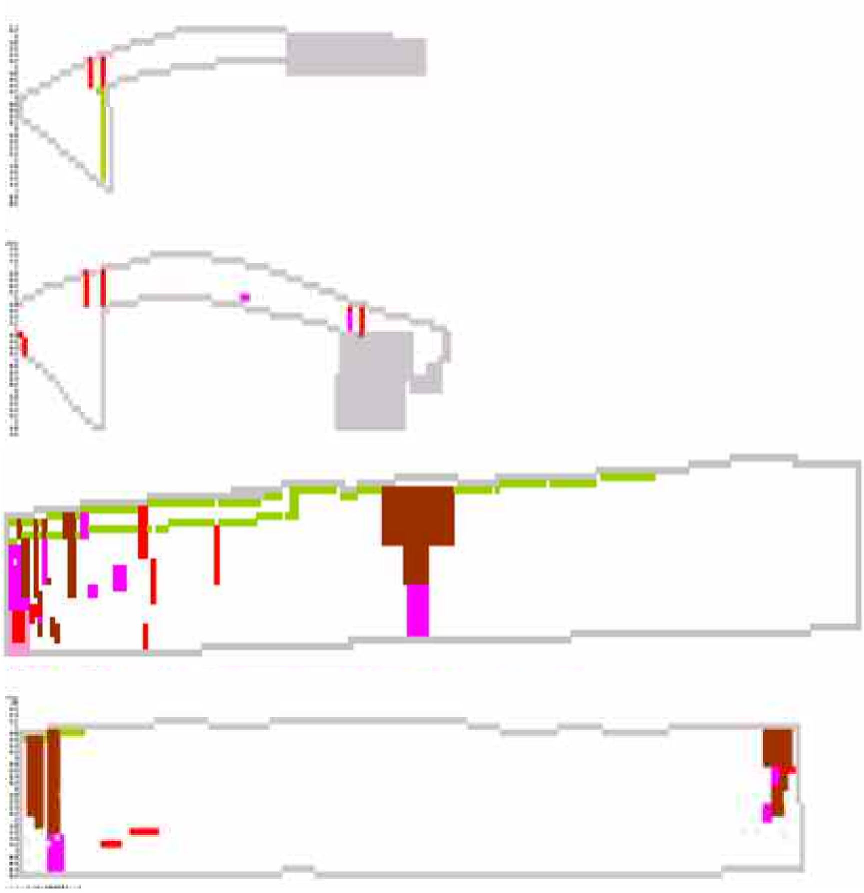

Figure 14: Results of onsite inspection translated into a defect map (top to bottom – north external wall, south external wall, west internal wall, east internal wall). For key see text.

Figure 15 shows the defect map obtained from the manual inspection of the images. As above, the outline of the surface being considered is shown in grey; non-defective features are shown in green; defects are shown in reddish colours. Note that parts of the image were not able to be imaged because they were obscured (e.g. by trees). These grid squares are shown in black. Blue lines represent areas of the structure which were seen to be wet or damp in the images.

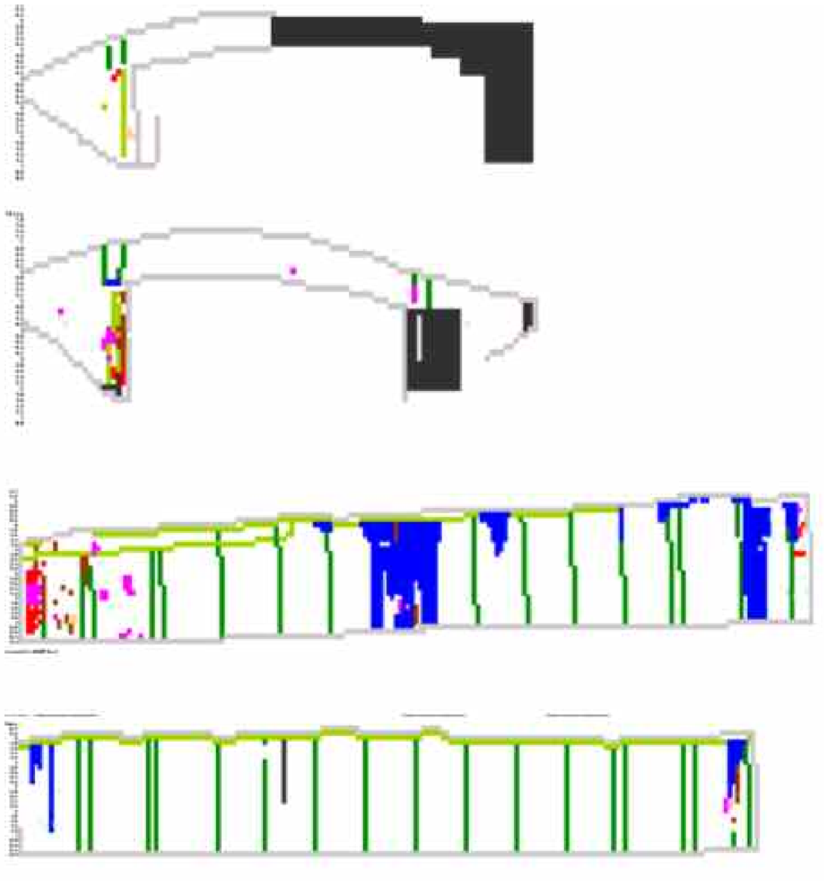

Figure 15: Results of image based inspection translated into ChartCrack defect map (top to bottom – north external wall, south external wall, west internal wall, east internal wall).

Comparison of Figure 14 and Figure 15 shows that the manual inspection of the images generally identified the areas of the structure which were affected by defects, and it therefore follows that this inspection generally identified those parts of the structure where no defects were present.

However, there are some differences between the results of the two surveys. Noticeably, the onsite assessment has not reported features such as the regularly spaced vertical joints running along the length of the internal abutments. These features will have been seen and inspected by the onsite inspector, but, as there were no defects associated with them, they have not been mentioned in the report, or photographed. Therefore they were not included in the translation of the inspection results to the defect map. The onsite assessment therefore reports on general areas of defects, and specific features of concern, but does not accurately record the extent or position of features or minor defects in detail, and does not mention features felt to be of no significant concern to the engineer. In contrast to this, the manual inspection of the images has recorded information for the entire surface in significant detail. This is because of the way in which the inspections of the images were carried out. This inspection examines every part of each image, and marks any defect or feature seen.

However, both methods of inspection highlighted the same general areas as having the most defects, and the image based assessment correctly matched the type, location, severity and extent of most of the defects recorded during the onsite assessment.

3.4 Manual inspection of images of structures – Conclusions

The investigation showed that there is sufficient detail in images collected of a bridge structure to enable a meaningful inspection of the bridge to be performed offsite. The results of inspections performed this way compare well with the results of onsite inspections, although there are clearly some differences, such as the onsite inspector being able to have multiple viewpoints or perspectives on any areas of interest.

Inspections based on the manual inspection of images, such as those used in this section, would be very useful in enabling the engineer to have a first look at a bridge before going out on site. This would familiarise them with the site and the structure, and give them some forewarning of the condition of the bridge, letting them know which areas needed particular attention and which were less critical.

The collection of images in successive years would enable engineers to store a database of images, all taken in the same way, and covering the entire structure each year. Any particular defect or feature can be seen and positively identified in successive years' images. This would be a very useful tool, giving the engineer the ability to track any changes in severity and extent of a defect very accurately from year to year.

4 Image Processing and Automation of object detection

Having successfully demonstrated that enough information existed in the images for a meaningful inspection of the structure to be performed using the images alone, attention turned to the use of image processing techniques to automatically identify defects. The goal of this initial stage of image processing was to develop a simple method to distinguish between those areas of the images that contained only the sound bridge surface (“background”), and those which contained a feature or defect (“feature”). This process is referred to as segmentation of the image into background and feature segments.

4.1 Data used – bridges

The development of automated techniques utilised images from the structure at Winnersh (used in the manual assessment of images and shown in Figure 8 and Figure 9), and two additional bridges. These bridges were located in Eton Wick (Figure 16 and Figure 17), and between Frilsham and Hampstead Norreys (Figure 18 and Figure 19).

Although these additional structures were constructed from concrete, similarly to the structure at Winnersh, these bridges differed in a number of areas. In particular, the Eton Wick structure has a grid like appearance on the supporting walls (Figure 17) and has a surface containing larger aggregate sizes visible on the surface, giving it a more roughly textured appearance than the other two smoother bridges.

Figure 16: Map of Eton Wick site. A332 over B3022.

Figure 17: Eton Wick site. A332 over B3022.

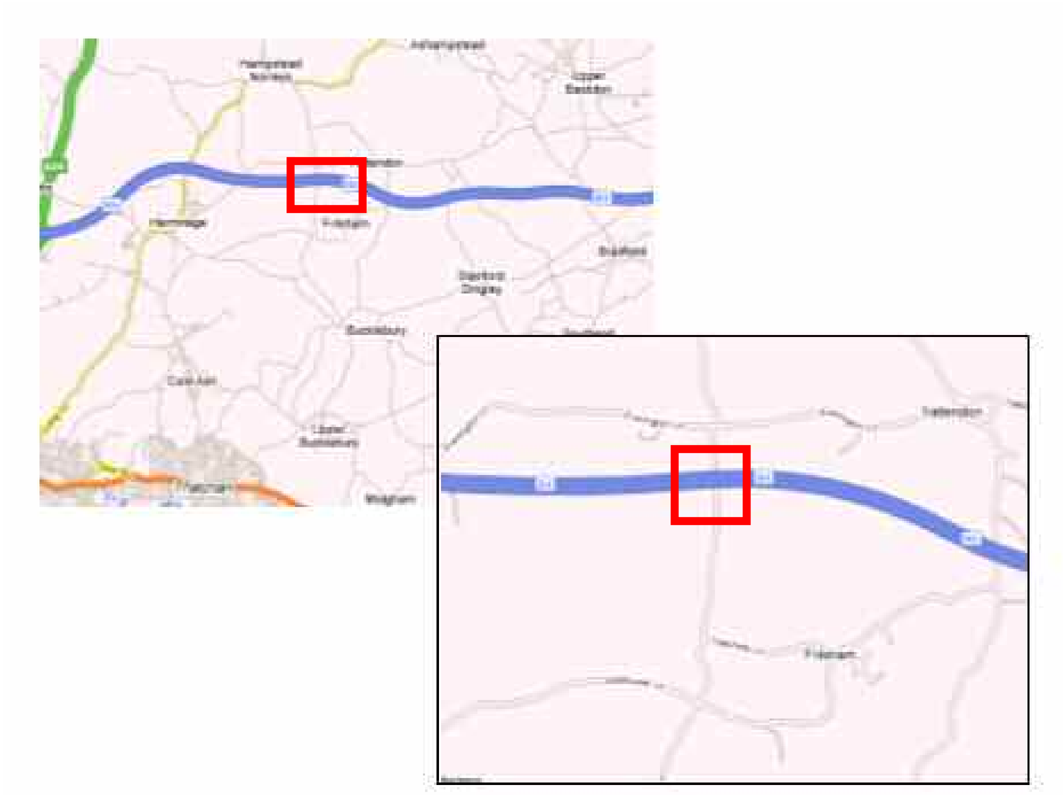

Figure 18: Map of Frilsham site. M4 over minor local road.

Figure 19: Frilsham site. M4 over minor local road.

4.2 Development of algorithms for segmentation of images

The development of image processing algorithms to segment images into background and feature segments was carried out in the Matlab development environment. The approach drew on the development of existing image processing tools, and considered how these could be combined to deliver efficient discrimination between grid squares considered as background and those considered to contain a feature.

4.2.1 Initial segmentation

An initial investigation made use of image processing algorithms to calculate the local entropy of the image, and identify edges within the images using the Canny edge detector. These methods are described in Appendix B. Both of these methods are targeted at the identification of areas containing features, as features are likely to contain edges or texture. The two methods were applied, in turn to the original images, and the results were then considered as a series of 200mm grid squares, which had the values 0 or 1, depending on whether or not a 'feature' had been detected within each grid square. These binary results, indicating the presence or otherwise of detected features were able to be displayed for direct comparison with the results of the manual inspection of the images. Note that the methods developed made no attempt to try to distinguish between defects and non-defect features, only between 'any feature' and 'background'.

4.2.2 Improved segmentation of features

To improve the segmentation between parts of the images containing 'background' and those containing features, and to reduce the false positives (i.e. the grid squares reported to contain a feature that actually do not contain a feature) reported using the initial segmentation methods, a number of additional techniques were considered. Each technique was applied to each image in turn, and assessed using a cell based, block processing approach.

In this block processing approach a value for each technique was calculated for each 200mm grid square within the image. This resulted in the creation of a series of 150 ([15 200mm columns] x [10 200mm rows in each image]) n-dimensional vectors for each image. These results were then compared with the the results of the manual image based assessment, which had already classified each 200mm grid square box to contain 'background' or 'any feature'. Two groups of data were created; one containing the results for the techniques when applied to grid squares which contained only background, and one containing the results for the techniques when applied to grid squares that contained features.

4.2.2.1 Approach to, and selection of, the discriminant function

A total of 15 different techniques (referred to as metrics – list given in Appendix C) were tested, none of which were able to reliably distinguish between the two classes of image segment on their own. The correlations between the various metrics were assessed to see which would complement each other most effectively in the improved identification of features and background. Following this, and the selection of the metrics most likely to work well in a classification function, a stepwise discriminant analysis was performed to calculate the optimum settings for use in a discriminant function.

The metrics tested included the quadtree decomposition of the image, a selection of different edge detection methods, and the standard deviation and mean of the pixel values within each cell, in 2 of the colour channels.

For example, Figure 20 illustrates the principle behind the quadtree decomposition. In this approach the image is split into 4 squares, each square is then assessed to determine whether or not it satisfies a defined homogeneity criterion. If it does it is left intact, if it does not then it is split into 4 more squares, each of which is then assessed for homogeneity itself. Regions of the image which show sharp changes in pixel values, or which differ from neighbouring pixels, are therefore split into more squares than regions of the image which show little or no changes.

By applying this process to the images, and then examining each cell of the image, it is possible to determine which parts of the image were more or less homogeneous than other parts.

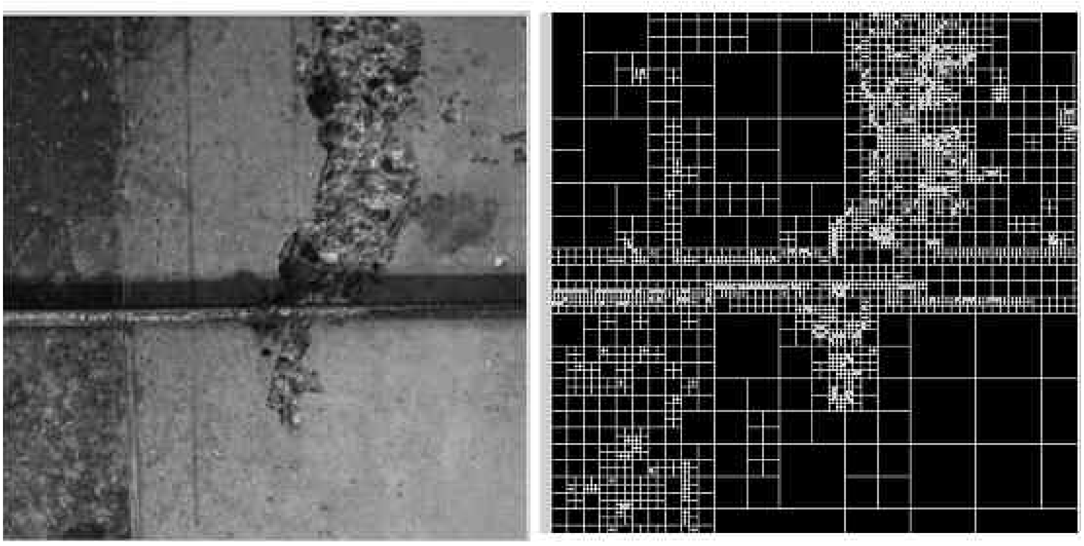

Figure 20: Example of a part of a structure image and a visual representation of its Quadtree Decomposition

4.2.2.2 Determination of key dimensions

To determine which of the fifteen metrics would combine most successfully into a single discriminating function, the correlations between them were tested. In investigating independence of these variables, several extremely strong correlations were uncovered. Similarly, the correlations between a number of the parameters were found to be quite weak. These were then taken forward to be included in a stepwise discrimination analysis.

This produced the resulting discriminant function:

$$\text{Discriminant} = (2.31 * \sqrt{(\mathbf{S})}) - (1.889 * \ln(\mathbf{Q})) - (1.470 * \mathbf{M}) - (35.092 * \mathbf{E}) - 9.237$$

Where:

S = standard deviation of pixel values (red channel);

Q = result of quadtree decomposition of image;

M = mean pixel value of cell (red channel);

E = entropy of cell (calculated differently to the entropy parameter used in the initial segmentation).

If the discriminant score for a given cell was above a threshold value then that cell was classed as including something of interest to the engineer (feature or defect), if it was below, it was classed as image background.

4.3 Results of Image processing and automation of object detection

4.3.1 Initial Segmentation

Figure 21 shows the reference data obtained by the manual inspection of the images taken of the east internal surface of the Winnersh structure. It can be seen that this surface is in generally good condition, with very few defects, with only a number of areas of damp at either end of the wall considered to be of likely concern.

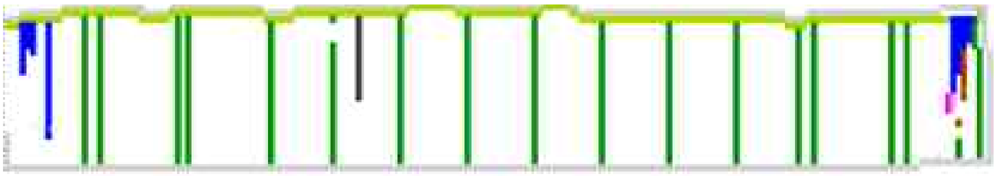

Figure 21: Reference image inspection results – east wall, Winnersh bridge

Figure 22 shows the results of the two initial methods of image segmentation: Entropy (top), and Canny edge detection (middle). All of the grid squares reported to contain a feature by the algorithms are filled in with either green or red. Where the image processing has detected a feature at the same place as the reference data showed the presence of a feature or defect the grid square has been coloured green. False positives are coloured red. Grey grid squares indicate grid squares where a feature or defect has been missed by the image processing methods.

The bottom part of Figure 22 shows the results of the combination of the two methods. In this display the red is where both the two image processing approaches have detected the presence of a feature or defect, the blue is where only one of them has detected something, and again, the grey shows the location of features or defects which have been missed by the image processing methods.

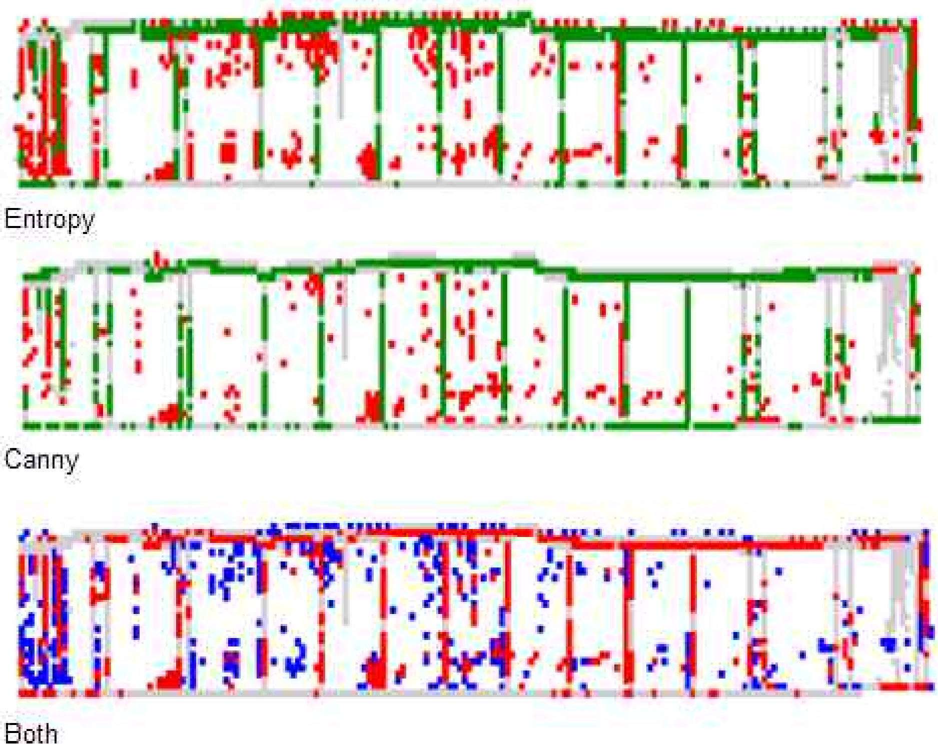

Figure 22: Initial image segmentation results – east wall, Winnersh bridge

The general shape of the bridge is clearly visible, including the regularly spaced vertical joints in the concrete. The distribution of the red areas would suggest that the areas of concern in this bridge are the left hand end (which agrees with the reference inspection), and a couple of areas towards the bottom of the structure approximately 20 and 40% along the bridge (nothing is seen here in the reference data). The reference data also indicated that the right hand edge of the bridge would require inspection as it showed the presence of water, and a few small defects. These have been missed by the initial segmentation methods. There are also a number of isolated red patches visible, which do not match the reference data in any way, and even more blue patches, where only one of the segmentation approaches reported a feature.

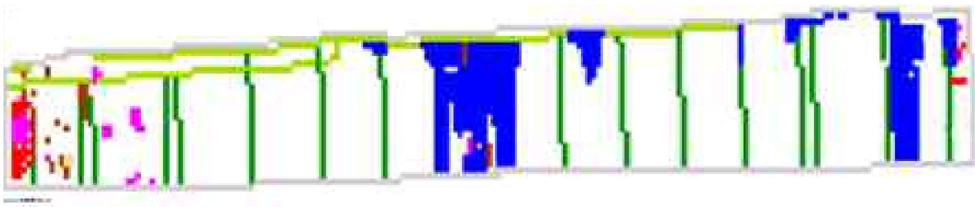

Figure 23: Reference image inspection results – west wall, Winnersh bridge

Figure 23 shows the reference data obtained by the manual inspection of the images taken of the west internal surface of the Winnersh structure. As for the east wall, this wall also shows two large areas of damp. However, in contrast to the east wall, this surface exhibits more defects, and a wider variety of defect types, indicated by the different colours present.

Figure 24 shows the results of the combination of the two initial methods. Here red grid squares show where both image processing approaches have detected the presence of a feature or defect, and blue squares show where only one of them has detected something. Grey squares show features that appear in the reference that have been missed by both image processing methods.

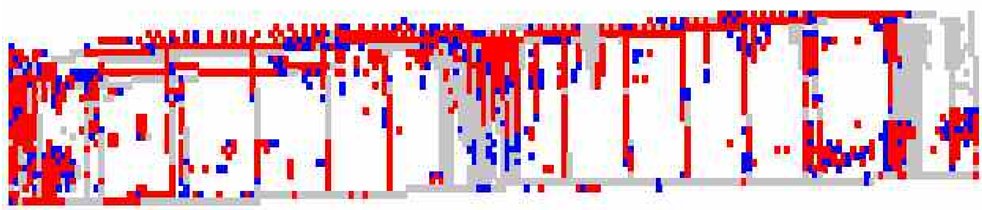

Figure 24: Initial image segmentation results – west wall, Winnersh bridge

Again in Figure 24 the general shape of the bridge is clearly visible, including the regularly spaced vertical joints in the concrete, and the cabling running horizontally along the top of the left third of the surface (light green line running along the top of Figure 23). The distribution of red grid squares would suggest that the areas of concern in this bridge are at the left hand end (which agrees with the reference inspection), a couple of patches on the wall midway up the second slab from the left (corresponding to areas of spalling seen in the reference data), the centre of the bridge (which is a large wet area), and the right hand edge of the surface (which is again wet, with some cracking). Nevertheless, the reference reported the presence of water, and a few small defects, at the right hand end of the bridge that have been missed by the initial segmentation methods, and the extent of the wet area in the middle of the bridge has been underestimated. However, Figure 24 shows far fewer false positives than were visible on the East wall of the same structure.

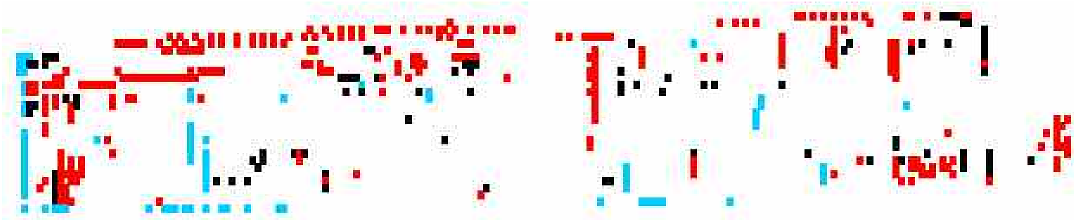

Figure 25: False positive reports using initial segmentation methods: Red = feature or other obvious cause; Blue = adjacent to feature / defect, possible reference data alignment issue; Black = no obvious cause

The causes of any false positive reports of 'something' on the bridge surface were investigated. It was found that most of the false positives were a result of shadows, dirty marks or textural anomalies on the surface of the structure, but a substantial minority were unexplained. Figure 25 shows the explanations for the false positives found on the west face of the Winnersh structure.

4.3.2 Improved segmentation of features

Following the development of the initial methods, work continued (in Stage 2 of the project) to improve the segmentation, using additional techniques for discrimination between background and features, as described above. The performance achieved with these additional techniques was again assessed against the reference. However, to obtain a better handle on the reasons for the poor performance the reference data was re-defined on the basis of four key types:

- 'Background' (bridge surface in sound condition);
- 'Features' (objects such as joints, cabling, drainage systems which were supposed to be there);
- 'Defects' (features visible on the surface of the bridge which were not supposed to be there);
- 'Other' (neither a feature nor a defect, but something which may confuse an automated system).

Figure 26 shows the reference data for a face of one of the sample structures (Winnersh structure) when defined in terms of these types. Defects are shown as red, features that are supposed to be there as green, and other marks or features that are not defects, but are not part of the structure as yellow. It can be seen that there are large concentrations of defects at either end of the structure, and in the middle, with isolated patches of defects elsewhere, and a horizontal feature running along the top left hand third of the bridge. The joints between the concrete slabs of which the bridge is constructed are clearly visible as regularly spaced vertical green lines.

Figure 26: Manual image analysis reference data (individual images, Winnersh bridge, west face). Red = 'Defects'; Green = 'Features'; Yellow = 'Other'

Figure 27 shows the results of the initial segmentation methods following application to the same images used in the manual assessment. In this display red indicates places where both initial methods reported a feature, amber indicates a feature reported by only one of the methods. As in Figure 26, the display shows heavier concentrations of reported features at either end, and in the middle of the structure. The vertical joints between the concrete slabs are clearly discernible in the display, as is the horizontal defect running along the top of the left hand third of the bridge. Also visible in Figure 27 are a lot of false positives, particularly ones reported by one or other of the two methods.

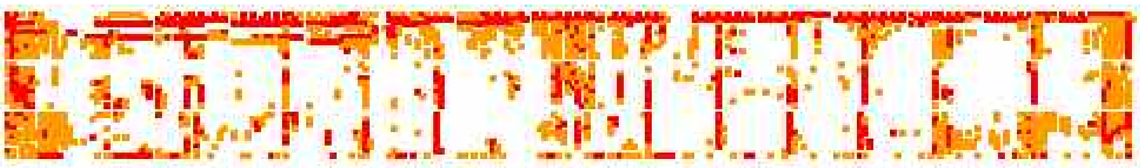

Figure 27: Initial segmentation results (individual images, Winnersh bridge, west face). Red = Reported by two methods; Amber = reported by one method

Figure 28 shows the results of the discriminant function proposed in Section 4.2.2.1. This uses a simple red/blank display, with red showing those parts of the bridge which the discriminant has classed as a feature. The concentrations of features seen at either end and the middle are clearly visible, as are the vertical joints, and the horizontal feature. Once again, though, there are a lot of false positives.

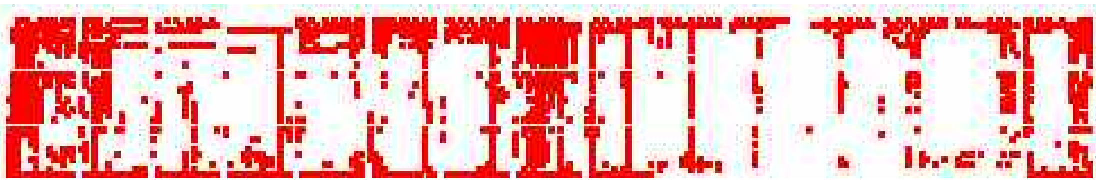

Figure 28: Discriminant function results (individual images, Winnersh bridge, west face). Red = classed as a defect / feature by discriminant

Figure 29 shows the results of combining the initial segmentation methods with the discriminant function. In this display pink shows the results of the discriminant function alone, and red shows the locations where the discriminant function and both the initial segmentation methods reported a feature. Here, the vertical joints show up quite clearly, and the left and middle of the bridge show slightly more red than the rest of the structure, but overall the results are disappointing.

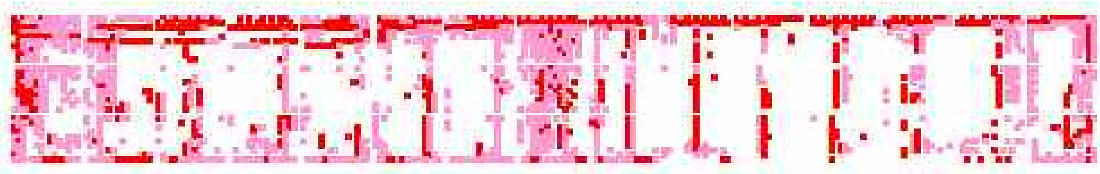

Figure 29: Combination of discriminant function results and initial segmentation. (individual images, Winnersh bridge, west face). Pink = classed as a defect / feature by discriminant; Red = reported by both initial segmentation methods.

Figure 30 shows another display of the initial segmentation and the discriminant, but this time the initial segmentation has been used to rule out parts of the bridge which do not contain anything of interest. In this display light blue shows those parts of the bridge which did not generate any response from either of the initial segmentation methods, which are now assumed to not contain anything of interest. Red areas now show the areas of the structure which were classed as containing something of interest by the discriminant function, and were not ruled out by the initial segmentation.

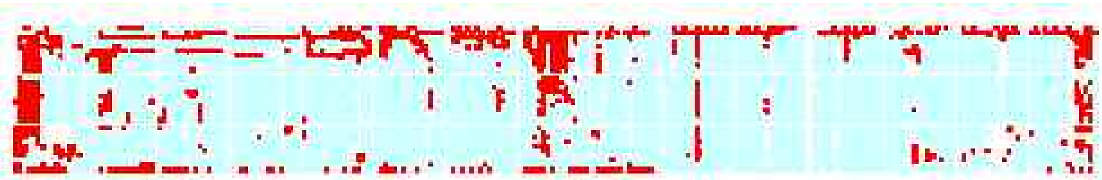

Figure 30: Combination of discriminant function results and initial segmentation. (individual images, Winnersh bridge, west face). Blue = zero reported by both initial segmentation methods; Red = classed as a defect / feature by discriminant, and not ruled out by initial segmentation.

Figure 31: Manual image analysis reference data (individual images, Winnersh bridge, west face). Red = 'Defects'; Green = 'Features'; Yellow = 'Other'

Figure 31 shows the reference data for the structure again (as in Figure 26), for comparison with Figure 30. It can be seen in Figure 30 and Figure 31 that the process of using the initial segmentation to detect areas thought to contain nothing of interest, together with the discriminant function to detect features of interest gives encouraging results, in terms of defects and features. It must be noted however, that the vertical joint features, clear in the reference data, are not detected in the automatic analysis shown in Figure 30.

5 Additional areas for investigation

The research has shown some encouraging results in the identification of features on the structures using image processing techniques, with the discriminant function suggested in stage 2 assisting in the reduction of false positives, whilst increasing the confidence that a part of the bridge reported as containing a feature, does in fact report locations of interest to the engineer. However, the system is still far from complete, and a number of areas remain for investigation.

5.1 Increased spatial resolution

By reducing the size of the grid cell used in the cellular approach to processing the images it may be possible to improve the spatial resolution of the processed results. However, care would have to be taken to ensure that this did not just result in an increase in the 'noise' of the results, or the computation time.

5.2 More complex metrics

The use of more complex metrics may enable improvements in two areas: firstly it may allow the discriminant function, which determines whether or not a grid square contains a feature, to be improved. Secondly, it may enable greater distinction to be made between those features which are accepted (supposed to be there), and those which relate to defects. This would be a key step in the development of a system which genuinely aided the engineers in their efforts to inspect the bridges.

Possible new metrics for investigation include a measure of the number of features within a cell. Currently the system simply considers the proportion of the cell, which is segmented by any of the metrics, but it is very likely that a cell containing a single large segmented object, with a total area of, for example, 450 pixels, contains something very different from another cell containing 45 small objects, each of 10 pixels.

In a similar vein, there is currently no metric relating to the shape of the segmented object, and no effort is made to find straight lines. Information regarding both of these aspects would be very useful in distinguishing features which are supposed to be on the bridge (joints, cables, pipes) from those which are not (defects).

5.3 Haar transform

Abdel-Qader, et al. (2003), who have also investigated defects within images of structures, have suggested implementing a cellular based approach similar to the one used in this research and in a further paper (Abdel-Qader, et al. 2003) tested four different approaches for detecting cracks in a small sample of bridge images. The four approaches used were the Sobel edge detector, the Canny edge detector (both of which were tested as part of this research), the Fourier transform and the Haar transform. The report suggested that by far the best results were obtained using the Haar transform. It would be appropriate to review the Haar transform in further work in this research.

5.4 Crack detection

TRL have undertaken previous work to assess cracking in road pavement. The algorithms developed for the detection of cracking could be used in the assessment of images of structures.

6 Conclusions and Recommendations

Research up to the end of Stage 2 of this project has found that there is potential for the use of images for conducting offsite inspections of structures to a level of accuracy that is comparable with that achieved on-site. Stage 1 found that sufficient detail was contained in the images to enable an image based inspection to detect the areas of concern to the engineer, and showed some promise using two methods for performing an initial segmentation of the bridge into grid squares containing 'features and areas containing 'background'.

Stage 2 has investigated the use of more techniques, or metrics, for classifying grid squares as being of interest to the engineer or not. These new metrics were used to create a discriminant function which helped classify the cell. By combining the results of the initial segmentation with the additional information from the discriminatory function the incidence of false positives in the results was reduced, while maintaining the pattern of real defects and features detected.

Such a system shows great potential for use by an engineer. However, this work has not investigated how this approach would be viewed by engineers and therefore it is recommended that a number of engineers be consulted regarding the system, and how and when they would make use of it, in order to ensure that any further development is directed appropriately.

At this stage in the project the most promising results have been obtained through the use of initial segmentation (set such that any grid-square which is not reported to contain a feature can be confidently ruled out), followed by the use of discriminant function to determine which of the remaining grid squares can be classed as something of interest. Further research should be carried out to refine the system, and optimise the discriminant function.

If suitable levels of accuracy can be achieved in the segmentation then it would be appropriate to develop classification methods to distinguish between those segmented features which are accepted part of the structure, and those which are defects. This would enable the production of a very helpful semi-automatic system for use by the engineers. Such a system would be able to process images of a structure, and provide an engineer with an indication of which areas require close attention. It is felt that such a semi-automatic inspection system, which provides guidance but still requires an engineer to confirm the results, would be an appropriate first target for this research.

7 Acknowledgements

The work described in this report was carried out in the Technology Development Group of TRL Limited. The authors are grateful to Richard Woodward who carried out the quality review and auditing of this report.

8 References

I Abdel-Qader, S Pashaie-Rad, K Ahmed and O Abudayyeh, '*A local PCA algorithm for inspection of concrete bridges'*, Document name, 2003

I Abdel-Qader, O Abudayyeh and M E Kelly '*Analysis of edge-detection techniques for crack identification in bridges'*, Journal of Computing in Civil Engineering, Vol 17, No 4, 2003

Appendix A. Soffit Images

Figure 32 shows a series of six images of the soffit of the Winnersh structure which have been arranged as a mosaic designed to show a larger area of the soffit than is available in any single image. This composite image clearly illustrates the problems which were encountered in trying to align the soffit images.

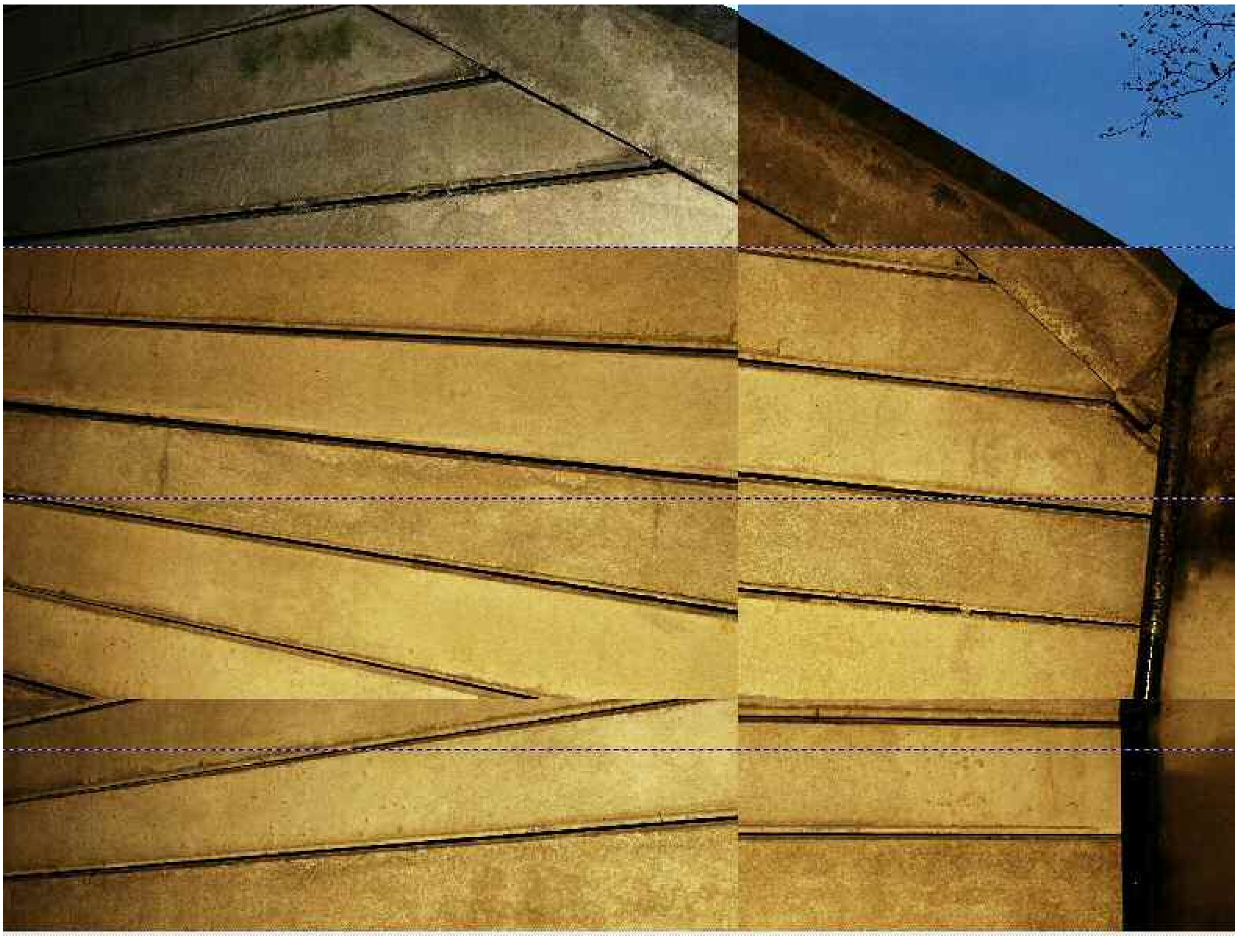

Figure 32: Illustration of problems encountered when attempting to align images of soffit

Appendix B. Image processing and statistical techniques

B.1 Entropy

Entropy is a statistical measure of randomness that can be used to characterise the texture of the input image. Entropy is defined as:

$$-\sum(p.*\log(p))$$

Where ‘p’ contains the histogram counts of pixel values.

B.2 Std dev

$$s = \left(\frac{1}{n-1} \sum_{i=1}^{n} (x_i - \bar{x})^2 \right)^{\frac{1}{2}}$$

B.3 Canny

The Canny operator was designed to be an optimal edge detector (according to particular criteria - there are other detectors around that also claim to be optimal with respect to slightly different criteria).

The Canny operator works in a multi-stage process. First of all the image is smoothed by Gaussian convolution. Then a simple 2-D first derivative operator is applied to the smoothed image to highlight regions of the image with high first spatial derivatives. Edges give rise to ridges in the gradient magnitude image. The algorithm then tracks along the top of these ridges and sets to zero all pixels that are not actually on the ridge top so as to give a thin line in the output, a process known as *non-maximal suppression*. The tracking process exhibits hysteresis controlled by two thresholds: *T1* and *T2* with *T1* > *T2*. Tracking can only begin at a point on a ridge higher than *T1*. Tracking then continues in both directions out from that point until the height of the ridge falls below *T2*. This hysteresis helps to ensure that noisy edges are not broken up into multiple edge fragments.

B.4 Mean

The mean pixel value is the simple arithmetic mean of the pixels within the region of interest.

B.5 Quadtree

Quadtree decomposition is an analysis technique that involves subdividing an image into blocks that are more homogeneous than the image itself. This technique reveals information about the structure of the image. It is useful as the first step in adaptive compression algorithms.

This function works by dividing a square image into four equal-sized square blocks, and then testing each block to see if it meets some criterion of homogeneity (e.g., if all the pixels in the block are within a specific dynamic range). If a block meets the criterion, it is not divided any further. If it does not meet the criterion, it is subdivided again into four blocks, and the test criterion is applied to those blocks. This process is repeated iteratively until each block meets the criterion. The result might have blocks of several different sizes.

B.5.1 Image and a Representation of its Quadtree Decomposition

The following figure shows an image and a representation of its quadtree decomposition. Each black square represents a homogeneous block, and the white lines represent the boundaries between blocks. Notice how the blocks are smaller in areas corresponding to large changes in intensity in the image.

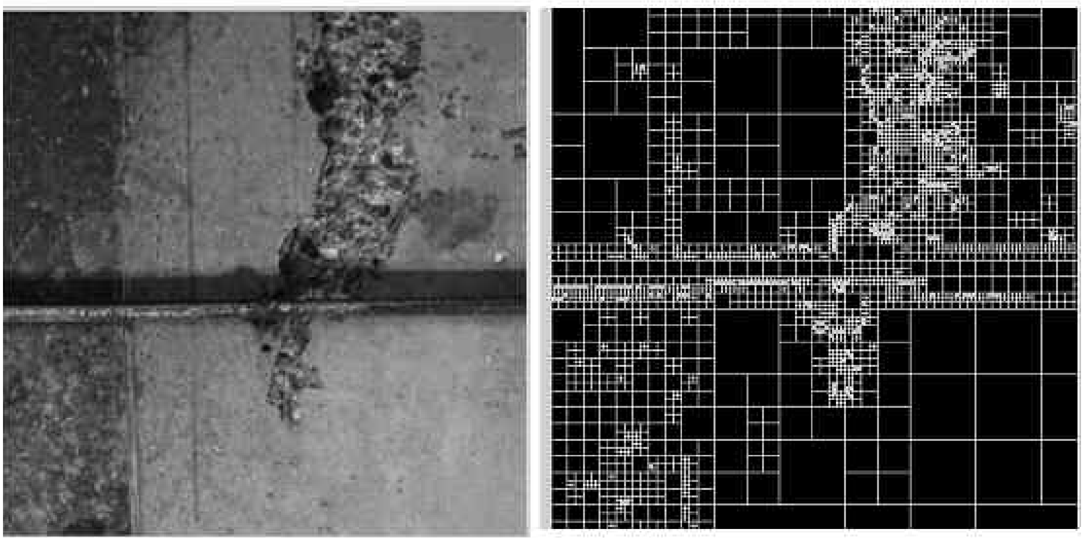

Figure 33: Example of a part of a structure image and a visual representation of its Quadtree Decomposition

Appendix C. Metrics investigated

- Entropy of area surrounding pixel in Red channel;
- Entropy of area surrounding pixel in Green channel;
- Standard deviation of area surrounding pixel in Red Channel;
- Standard deviation of area surrounding pixel in Green Channel;
- Mean value of area surrounding pixel in Red channel;
- Mean value of area surrounding pixel in Green channel;
- Quadtree decomposition;
- 6 Edge detection methods:
    - Canny
    - Sobel
    - Roberts
    - Prewitt
    - Zero Crossing
    - Laplacian of Gaussian

Abstract

The quality of data provided by visual structure inspections can vary significantly from inspector to inspector and from inspection to inspection. Improvements to the quality of the inspections are therefore desirable.

Research has been undertaken with a view to developing a system for collecting a series of images covering the entire surface of the structure, and then pre-processing the images prior to delivery to the engineers. The aim is that these pre-processed images should be able to give the engineer an overall view of the condition of the structure, and draw attention to those parts of the structure which contain most defects, or defect like features.

The research has investigated two main areas: image collection and display; and image analysis.

The image collection and display investigation considered the use of multiple imaging positions, single imaging positions, spherical images, and mathematically transforming images to re-project them as if they were taken perpendicular to the face of interest in order to remove the effects of parallax.

The image analysis research considered the possibility of using high quality images of structures to provide a consistent and quantitative assessment of condition. Sufficient information and detail was seen in the images to perform meaningful and useful image based condition assessments. The focus of the research then switched to the use of automatic image processing techniques. The research in this area has focussed on reducing the incidence of false positive reports of defects, and delivering results to the engineers which make their job simpler, quicker and more cost effective.