

Transport Research Laboratory



Automated inspection of highway structures

Stage 3

by S McRobbie

PPR 338

PUBLISHED PROJECT REPORT



PUBLISHED PROJECT REPORT PPR 338

Automated inspection of highway structures

Stage 3

by S McRobbie (TRL)

Prepared for: Project Record: Automated Inspection of Highway Structures
Client: Transport Research Foundation,
(Rod Kimber)

Copyright Transport Research Laboratory August 2008

This Published Report has been prepared for Transport Research Foundation. Published Project Reports are written primarily for the Client rather than for a general audience and are published with the Client's approval.

The views expressed are those of the author(s) and not necessarily those of Transport Research Foundation.

	Name	Date Approved
Project Manager	Stuart McRobbie	28/08/2008
Technical Referee	Richard Woodward	15/09/2008

If this report has been received in hard copy from TRL, then in support of the company's environmental goals, it will have been printed on recycled paper, comprising 100% post-consumer waste, manufactured using a TCF (totally chlorine free) process.

Contents

Executive summary	ii
1 Introduction	1
1.1 Previous work	2
1.1.1 Image requirements	2
1.1.2 Imaging problems	2
1.1.3 Image segmentation	2
1.2 This work	3
2 Data used	4
2.1 Reference data	5
3 Image collection	7
3.1 Possible image collection set ups	7
3.1.1 Proposed image collection system – draft specifications	8
4 Reprojection / processing	10
5 Analysis / interpretation	12
5.1.1 Performance assessment	13
5.1.2 Classification	16
6 Discussion and Conclusions	18
Acknowledgements	19
References	19
Appendix A	20
Appendix B	21

Executive summary

The recent collapses of the I-35W Mississippi River bridge in Minneapolis, and Montreal's Boulevard de la Concorde overpass have highlighted the importance of regular structure inspections to ensure they remain safe and fit for purpose. Currently, in the UK, highway structures are assessed using a regime of visual inspections, performed by trained engineers. These inspections are performed at four levels: Routine, General, Principal and Special. The inspections cover a range of detail, from a cursory check for gross defects, to a close examination of all surfaces of the structure, including the use of special equipment if required. The quality of data collected depends on the ability of the inspectors to observe and objectively record details of defects. It has been found that the data provided by such inspections can vary significantly.

Research has been performed to investigate the use of images of the structure to assess its condition off-site. The aim is not to remove the engineer from the inspection process but to assist them and make their job easier.

The research has concentrated on two main areas: image collection, and image analysis.

The image collection work has investigated the practical issues involved in imaging structures. Such issues include image resolution, lighting, removal of parallax, location referencing of individual images and the development of a prototype collection system.

This prototype system makes use of distance measurement lasers and theodolites to determine the position of each image on the structure, and relative to any other image, making it easy to know precisely which parts of the structure are affected by any particular defect.

The image analysis work has attempted to segment the images so that the defects or features present on the structure can be highlighted and classified. The segmentation work has made use of a number of image processing techniques including wavelet analysis and has achieved promising results to date.

Processing work has begun to classify the segmented objects into those which should be present on the structure (cabling, drainage, lighting, etc) and those which should not be present on the structure (defects).

1 Introduction

The recent collapses of the I35-W Mississippi River bridge in Minneapolis, and the Boulevard de la Concorde overpass in Montreal, which claimed 18 lives, have tragically highlighted the importance of regular structure inspections on the highway network. Structure inspections are the first tool in the engineers' armoury for ensuring that structures remain safe and fit for purpose.

Currently, in the UK, highway structures are assessed using a regime of visual inspections. These inspections are carried out at one of four levels of detail: Superficial, General, Principal and Special. The inspections cover a range of detail, from a cursory check of the gross visual condition of the structure, to a detailed inspection and examination of all surfaces, including the use of special equipment, if required, to monitor 'hidden' elements of the structure. The quality of data collected depends on the ability of the inspectors to observe and record details of any defects. For the surveys to be useful this must be done objectively and consistently.

It has been found that the data provided by such manual visual inspections can vary significantly – both from inspector to inspector and from inspection to inspection.

TRL has been at the forefront of automated pavement monitoring research and development for over a decade, and has considerable knowledge and experience in the use of technology in condition monitoring. Work done in this area has not only led to the development of HARRIS and HARRIS2 (Figure 1), but the implementation nationwide of automated routine condition monitoring regimes – TRACS (TRAffic-speed Condition Surveys) on Highways Agency roads, and SCANNER surveys (Surface Condition Assessment of the National Network of Roads) on local authority roads. The potential for a similar approach being applied to the area of structure condition monitoring has been identified. Therefore research has been commissioned by the Transport Research Foundation to investigate the use of images in the structure inspection process. These images would keep an accurate and complete record of the condition of the structure, would enable an engineer to perform a virtual inspection of a structure from the comfort and safety of the office, and could perhaps be automatically processed and analysed to assist the engineer in assessing the condition of the structure. The aim of the research is not to remove the human from the process, but to assist them, and make their job easier resulting in a more efficient condition monitoring / maintenance planning cycle. It must also be borne in mind that the development of any routine structure condition monitoring system is potentially a long process, which has only just begun.



Figure 1: HARRIS1 and HARRIS2 vehicles, developed and operated by TRL on behalf of the Highways Agency.

The research has been underway for three years, and good progress has been made, although thus far, only concrete structures have been considered. The target to date has been the successful segmentation of features on concrete structures. This should be

accurate enough to highlight any areas of concern to an inspector unfamiliar with the structure who can then decide whether or not they require further investigation. A secondary objective is the characterisation and classification of such segmented features. This should be able to determine which features are supposed to be there (cabling, drainage, joints etc.) and which features should not be there and are possible defects.

1.1 Previous work

Earlier stages of the research established that comparable assessments could be made of the condition of a concrete bridge by inspections performed on site and off site, using only colour images of the structure. This showed that the images contained sufficient information to be of use to an engineer. The requirements for image quality were developed, and some of the issues involved in performing such surveys in the real world were identified. Previous stages of the research have also made some progress in the segmentation of the images to automatically determine which parts of the structure require more attention, and which are sound. A summary of the previous work performed in this project is given below.

1.1.1 Image requirements

It has been established that the images used in the research must satisfy the following criteria:

- Colour;
- Minimum of 1 pixel per mm;
- Lighting should be used to ensure the images are as consistent as possible;
- Location of each image on the structure must be recorded;
- The images should cover the whole surface of the structure, with some overlap between images to ensure full coverage.

Additionally, if the images are to be assessable by an engineer, then it must be possible to join or tessellate the images. This is to enable a set of images to be viewed together, to provide some context, as individual images can be hard to interpret with no sense of the surrounding area.

1.1.2 Imaging problems

The early work showed that, unless the camera was repositioned, both horizontally and vertically, for each image then the effects of parallax would be apparent on the images. This would result in the tessellated images looking very unnatural and making it harder for engineers to interpret them. This was especially true for soffits which the early stages of the research were unable to investigate.

It was also found that the image collection process was reasonably time consuming, but not nearly as time consuming as the process of trying to line up and tessellate the images by hand.

1.1.3 Image segmentation

Image processing and segmentation routines were developed, using commercially available software (Matlab ®), which were able to process all the images of a structure, one after the other, and determine which areas of these images contained any features which were likely to be of interest to the engineer. This was done by processing the

image in a number of ways to extract features of possible interest, and then splitting each image into an array of 200mm square grid cells. If the segmented objects in any of the cells met certain criteria then the cell was flagged as being of potential interest, otherwise it was flagged as being sound. In this way a map could be built up of the surface of the structure showing those cells which were of interest and those which were not. Visual comparison with reference data showed that this approach was achieving some success, but with too many false positives to be of any real use without further refinement.

1.2 This work

The current phase of research therefore has focussed on trying to address some of the problems and issues identified in the earlier work. In particular, efforts have been made to define the requirements for an image collection system which would simplify and streamline the process, not only of collecting the images on-site, but also of reconstructing the structure and tessellating the images for later interpretation and analysis.

This has addressed some of the practical issues with collecting images in the real world, and has aimed to produce plans for a system which could collect useable images satisfying the image quality requirements as quickly and safely as possible. The goal is to enable the collection, processing, display and analysis of images with a minimum of human input.

A large part of this work has been the development of methods which can be used to reproject images in such a way as to remove the effects of parallax. This has required recording data relating to the camera bearing and elevation at each imaging position. Such reprojected images must satisfy the image quality criteria after the reprojection process, and must be suitable for use with any image segmentation or analysis approaches developed.

Other key areas in the research have been the development and assessment of enhanced segmentation routines, and the initial development of methods for distinguishing between features and defects.

2 Data used

The bulk of the research and development performed to date has concentrated on a single bridge, although others have also been used. The main bridge used is a structure carrying the M4 over the B3030 at Winnersh, shown in Figure 2 and Figure 3. This has been found to exhibit a number of defects which are of interest to the project.



Figure 2: Map of Winnersh site M4 over B3030.



Figure 3: Winnersh site M4 over B3030 – southern view.

The bridge was of beam and slab construction with a width of 33.8m, and a span of 12.9m. The headroom under the structure was 5.05m. The bridge comprised precast post-tensioned beams (PSC), an in situ concrete slab, reinforced concrete cantilever abutments, and spread footings.

Fuller details of the bridge specifications are given in Appendix A.

Additional work has taken place at other local concrete structures, and the bridge over the TRL test track.

2.1 Reference data

Reference data was collected in order to assess the performance of the various segmentation methods tested. To produce the reference data each image (such as the one shown in Figure 4) had an array of 200mm grid squares overlaid on it. The images were then examined by an inspector who used software to record the locations (as shown in Figure 5) of any one of a range of defects (a list of recorded defects is given in Appendix B). Red cells in Figure 5 are those deemed to contain defects or features, orange cells contain other markings.



Figure 4: Original image of part of structure.




Figure 5: Reference data showing grid squares where features were noted by inspector, overlaid on image.




Figure 6: Reference data showing grid squares where features were noted by inspector.

This output feature map (Figure 6) was then kept for comparison against the automated segmented output. To assess the performance of the segmentation, the segmentation process was applied to each of the images in turn, and the output was compared, both visually and quantitatively, with the reference data for the same image.

3 Image collection

Any image collection system to be used in a real application of this approach must fit the following conditions:

- o Safe – ideally it should be able to operate from a safe place without the need for a road closure;
- o Fast – both to deploy the system in position, and to collect images;
- o Simple – should only require setting up and starting with no further human input at any location, other than moving it to the next location;
- o Reliable – Should collect position data (bearing and elevation) for each image repeatably and accurately;
- o High resolution – should collect images at high enough resolution so that they satisfy the image requirements after the reprojection process;

3.1 Possible image collection set ups

A number of possible image collection set ups were tested or considered, and the chosen system was an enhanced multiple position system. This system made use of a camera mounted on a tripod with an indexed spherical imaging head, and a pair of distance measurement lasers. The laser measurements and camera orientation data could be used to reproject the images and remove the effect of parallax. This gave some of the advantages of the simplicity of taking individual 'photograph-like' images from a static camera, but enabled more images to be taken from a single position than a simple multiple position set up. This greatly increased the speed of collection. The chosen approach seemed to be the most promising achievable with the time and budget available.

In the enhanced multiple position approach (shown in Figure 7) the camera is set up at a location (P_1), from which a number of images are taken at defined camera bearings and elevations. Laser measurements are taken to record the distance from the camera to the structure. The camera is then moved to a second location (P_2) and more images are taken. This process is followed until the entire structure has been imaged. The bearings, elevations and distances are then used to reproject the images as if they had been taken perpendicularly to the structure surface.

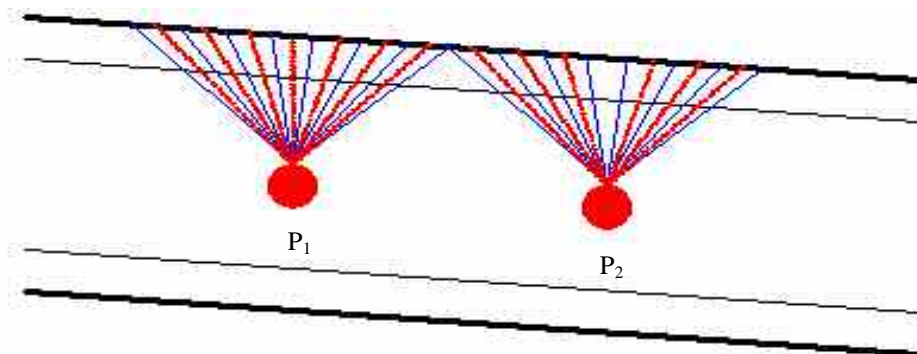


Figure 7: Illustration of enhanced multiple position.

3.1.1 Proposed image collection system – draft specifications

Figure 8 shows the imaging head used to collect images for this research. The image shows a camera mounted on a spherical imaging head, on a tripod, which is set to enable the focal point of the lens to remain in the same location in space even when the camera is rotated or elevated. Also visible are the twin distance measurement lasers (above and below the lens) used to ensure accurate camera – structure distance recording, and the flash used to minimise lighting variability.




Figure 8: Prototype image collection system used in work

Table 1 shows details of the image collection system used, and shows places where future automations could be made to streamline and simplify the process.

Table 1: Details of image collection system and areas for future automation

Feature	Current implementation	Future automation?
Camera	Nikon D30	Image taking could be automated by linking camera shutter control to system, image could be recorded as soon as feedback from pan-tilt unit indicates correct position has been reached.
Lens	105mm	The 105mm lens was suitable for imaging the Winnersh structure. Larger structures, where the camera had to be located further away from the structure could be imaged by using a longer focal length lens, and taking more images.
Resolution	3872 x 2592	Recent trends in digital imaging suggest that cameras capable of recording more and more pixels will become increasingly available and cheaper in the future.
Tripod	Spherical Imaging Head	
Position	Theodolite – record position relative to a known point	Total station – record absolute position of camera. Link this information with laser distance measurements and camera orientation data to establish where imaged surface belongs.
Imaging position	Click stopped horizontally, friction vertical, manual operation	Indexed, automated pan-tilt unit, possibly linked to camera shutter control to enable automatic image taking.
Image distance measurement	Lasers and manual data recording	Lasers and Bluetooth device to transmit and store measurement data
Lighting	Photographic flash	

4 Reprojection / processing

By making use of the imaging position data (bearing and elevation) for each image it is possible to mathematically resample the image data and reproject it to display the image as it would have looked had it been taken perpendicularly to the surface of the structure (Figure 9).



Figure 9: Example of images collected on site (left) and reprojected to appear perpendicular (right).

This produces images which are much easier for a human to look at and interpret as they are not full of discontinuities and misalignments. This is especially worthwhile on soffits as shown below in Figure 10. The top part of Figure 10 shows an illustration of problems encountered when lining up soffit images which have not been reprojected. As can be seen this gives very poor results. The same image data has been used, following reprojection, to produce the aligned image shown in the bottom part of Figure 10. This is much easier to work with and interpret.




Figure 10: Example of soffit images collected on site (top) and reprojected to appear perpendicular (bottom).

In order for the approach to be workable and for the images to be useful they must be collected at very high resolution so that, even following the reprojection process, they still show a minimum of 1mm per pixel.

5 Analysis / interpretation

The image analysis stage has been designed to assist the engineers identify areas of the structure which may need further investigation. This has been done on an image by image basis, using a grid based cellular approach. Reference data was produced as explained in Section 2.1. The segmentation processes were then applied to the images, and the results compared with the reference.

Of the two main segmentation approaches investigated, one was based on the entropy of the image, while the other was based on the use of wavelets – specifically the Haar transform, which had been identified in the literature (Abdel-Qader et al, 2003) as an approach which showed potential for detecting cracks in concrete structures.

Figure 11 shows the results of the entropy-based segmentation process on the same portion of structure shown in Figure 4. Figure 12 shows the results of the wavelet based segmentation method. The red and orange blocks show the reference data categorisation of each image cell: red areas have been deemed to contain defects or features in the reference data; orange areas contain marks which may confuse the segmentation processes.




Figure 11: Reference data and objects segmented using entropy based approach.




Figure 12: Reference data and objects segmented using wavelet based approach.

As shown in Figure 11 and Figure 12 neither approach gave perfect results (both missed some of the reference features, and gave a few false positives, but the two approaches picked up the major features very successfully. More detailed segmentation performance assessments are given in Section 5.1.1.

As well as the segmentation development, some attention was also given to the characterisation and subsequent classification of any segmented features. The primary goal was to be able to differentiate between any features present on the structure which were supposed to be there (cables, drainpipes, lighting, etc) and those which were not (defects). Section 5.1.2 discusses the performance of the characterisation methods developed so far.

5.1.1 Performance assessment

Figure 13 shows the reference data for the western abutment of the Winnersh bridge. The reference data has been grouped into three categories for ease of display – red shows defects, blue shows features, and green indicates 'other'. This data can be compared visually against the results of the segmentation process to see how well the segmentation is performed.

The data is split into three parts for ease of display. The top part of Figure 13 shows the first 40m of images starting from the southern end of the abutment and working north. This part of the bridge shows the presence of some defects, particularly at the southern end. A lighting cable runs along the top part of this portion of the structure, and can be seen in the reference data. The middle part is the next 40m of images, and includes a large area which was deemed damp by the inspector. The final part is the remainder of the surface, to the northern end of the abutment. The northern end of the abutment exhibits a number of defects, as did the southern end. A series of construction joints are seen throughout the reference data.

Note that the physical length of the abutment was approximately 35m, but the large amounts of overlap in each image have resulted in almost 100m of images.




Figure 13: Reference data – western abutment of Winnersh bridge.




Figure 14: Areas of structure segmented by image entropy-based segmentation method – west internal abutment of Winnersh bridge.

Figure 14 shows the areas of the structure segmented using the method based on image entropy. This correctly picks up the defects at the southern end of the surface (left end of top part of Figure 14), and some of the other defects and features, but misses a lot of the damp surface reported in the middle of the structure.



Figure 15: Areas of structure segmented by wavelet-based segmentation method – west internal abutment of Winnersh bridge.

Figure 15 shows the parts of the structure which the wavelet based method has segmented as containing a feature of potential interest to the engineer. This has reported far more features than the entropy method, and also far more than the reference data shows. However, despite these false positives, the general pattern of the bridge is still discernible, with a lot of defects at either end of the bridge, and the

regularly spaced construction joints being visible. A lot of the damp area in the mid section of the abutment is still not reported.




Figure 16: Areas of structure segmented by entropy (blue), wavelet (green) and both (red) segmentation methods – south end of west internal abutment of Winnersh bridge.

Figure 16 shows a section of the abutment where the objects segmented by either or both of the methods are shown. Those objects reported by the entropy method only are shown in blue, objects reported by only the wavelet method are shown in green, and red shows objects which have been segmented and reported by both methods. By combining the methods in this way and only investigating those reported by both methods the incidence of false positives is reduced, and the confidence levels in those that are reported is increased. Figure 17 shows the areas of the western abutment which have been segmented by both methods. The areas of defects at either end of the bridge are clearly visible, as are a number of the construction joints, the lighting cable, and some of the damp patch.



Figure 17: Areas of structure segmented by both segmentation methods – west internal abutment of Winnersh bridge.

Figure 18 shows how well the segmented data presented in Figure 17 compares with the reference data shown in Figure 13. Data from outwith the surface of concern (soffit or pavement have been excluded from this figure). In Figure 18 green parts show correctly segmented features, amber parts show false positives, and red areas show false negatives. These false negatives are of concern as these are areas which may need attention, but which the segmentation is missing. The majority of the false negatives, which are not related to the edge of the structure are from the damp area in the middle of the bridge. A number of defects have been reported here, but the extent of the dampness is very under reported.




Figure 18: Areas of structure segmented by both segmentation methods – west internal abutment of Winnersh bridge – red = false positive, amber = false negative, green = hit.

A crude numerical assessment of the performance suggests that the segmentation was correct 81.4% of the time, with 4% of the data being false positives, and 14.6% being false negatives.

5.1.2 Classification

The reference data shown in Figure 19 (the same as that in Figure 13, but reproduced here for convenience) showed the features present in the images as one of three categories – defects, features or 'other'. The segmented objects as reported in Section 5.1.1 have been characterised, and these characteristics used to classify the segmented objects.

Figure 20 shows the results of the classification on the segmented objects shown in Figure 18. In Figure 20 defects are shown in red, features are shown in green, and 'other' are shown in blue.




Figure 19: Reference data – west internal abutment of Winnersh bridge




Figure 20: Areas of structure segmented by both segmentation methods and classified – west internal abutment of Winnersh bridge – red = defect, green = feature , blue = 'other'.

It can be seen in Figure 20 that the characterisation and classification has been very successful at identifying construction joints and other features which are supposed to be present, and has successfully identified areas of defects at both the north and south ends of the abutment.

6 Discussion and Conclusions

The enhanced multiple position collection set up which was used appears to be the most suitable approach at this stage. This is a reasonably simple and straightforward method, requiring minimal training for the operators. Currently the method requires manual intervention to orientate the camera correctly and record distance measurements, but these tasks could be automated, as could the actual taking of the picture. If the original images are of high enough quality and resolution then the camera orientation data and the distance information can be used to reproject the images and remove the effect of parallax. This reprojection produces images which are far easier to interpret manually and are still suitable for use with the various segmentation and classification algorithms.

The segmentation methods developed previously which were based on the use of localised image entropy information have been enhanced by the addition of a new method based on the use of wavelets to detect features in the images. Combining these methods seems to improve the segmentation results by reducing the number of false positives reported when using any of the methods on their own.

The characterisation work, although in a very early stage, has shown very encouraging results, particularly at distinguishing between linear features (such as cabling or construction joints) and defects. It is expected that this aspect of the research could be considerably developed and improved in future phases.

Acknowledgements

The work described in this report was carried out in the Technology Development Group of the Transport Research Laboratory. The authors are grateful to Richard Woodward who carried out the technical review and auditing of this report.

References

Abdel-Qader I, Abudayyeh O and Kelly M, (2003). Analysis of edge-detection techniques for crack identification in bridges (v17(n4) p255-263), *Journal of Computing in Civil Engineering*. Reston, VA: American Society of Civil Engineers.

Appendix A

General inspection of King Street Lane Bridge

HA Structure no:	M4/59.2
SMIS Key No:	020
O/S grid no:	477840 170120
County/Borough:	Wokingham
Maintenance Area:	3
Maintaining Region:	OD South East
Structure Agent:	Mott McDonald with HCC
Structure owner:	HA
Date of construction:	1971
Designer:	Sir Alexander Gibbs and Partner
Design load:	HA + 45 units of HB
Bridge carries:	M4 over B3030 (King Street Lane)
Structure type:	Beam and slab bridge: precast post-tensioned beams (PSC), insitu concrete slab, reinforced concrete cantilever abutments, spread footings.
Span:	12.9m
Skew:	38°
Width:	33.8m
Headroom:	5.05m
Bearings:	Laminated rubber
Joints:	Flexcell
Parapets:	N2
Waterproofing:	HD Bituthene with Bitu-Dek
Carriageway:	2 x 3 lane + hard shoulder

Appendix B

Table 2: Defects (red), Features (blue) and 'Other' (green) looked for and recorded in reference data.

Seen on image
Cracking
Steel reinforcement
Spalling
Rust staining
Wet surface
Other Defect
Joint
Cable/drainpipe
Drilled (filled?) hole
Other Feature
Graffiti
Other 'Other'
Edge of structure / surface / image
