

Remote Inspection Technologies for Aircraft Skin Inspection

Mel Siegel¹ and Priyan Gunatilake²

¹The Robotics Institute - School of Computer Science

²School of Electrical and Computer Engineering
Carnegie Mellon University
Pittsburgh PA 15213 USA

Phone 412 268 8802 Fax 412 268 5569 Email: mws@cmu.edu WWW: <http://www.cs.cmu.edu/~mws>

Abstract - Aircraft flight pressurization/depressurization cycles cause their skins to inflate and deflate, stressing the skin around the rivets that fasten it to the airframe. The resulting strain, exacerbated by any corrosion that might be present, drives the growth of the initially microscopic cracks that are an unavoidable by-product of rivet installation. To avoid catastrophe, aircraft are periodically inspected for cracks and corrosion. The inspection technology employed is ~90% naked-eye vision. Of the ~10% that is instrumented, ~90% is with hand-held eddy current pencil probes. We have developed and demonstrated robotic deployment of remote eddy current and remote enhanced stereoscopic video instrumentation, augmented by image enhancement and image understanding algorithms, aimed at alleviating human performance problems and workplace hazards in inspection. This paper describes the aircraft skin inspection application, our two robotic deployment systems, the corresponding instrumentation packages we have demonstrated, and our progress toward developing image enhancement and understanding techniques to assist aircraft inspectors in finding cracks and corrosion. A wavelet-based approach is able to select radial features around rivets, and then to discriminate between cracks and scratches.

I. INTRODUCTION

Aircraft are routinely inspected for a variety of actual and potential problems: the airframe and major structures, engines, landing gear, fuel tanks, and skin are all subject to a variety of load and age related problems. Our work focuses on finding skin cracks and the corrosion that accelerates their growth. Skin inspection intervals are conservatively designed to allow a growing crack to be overlooked in two successive inspections and found in a third inspection well before any risk to flight safety is deemed to exist. Nevertheless there have been mishaps, the best known in April 1988 when a Boeing 737 operated by Aloha Airlines lost the crown of its fuselage from floor level to floor level, for a distance of six or seven seats back from the passenger door. This "incident" is generally attributed[1] to an unusual and unfortunate combination of circumstances: an "island hopping" airline with many short low altitude flights, an airplane with an unusually high number of pressurization/depressurization cycles, the corrosion-inducing moist sea air operating environment, a recent inspection that, in retrospect, should have resulted in an order for preventive repairs, and, most interesting from a fundamental perspective, the omission from then-current thinking about failure modes, of the possibility that two short cracks will meet to form one long crack. Since crack growth rate is more-or-less proportional to crack length (crack length increases more-or-less exponentially in time), the result is of this "multiple site damage" (MSD) is unexpectedly rapid growth and catastrophic structural failure. Shortly after this incident Congress mandated that the FAA begin an "Aging Aircraft Research Program". This research program has been highly successful, and its implementation has been rapid, to the point where the risk of an Aloha-type failure is now regarded as neutralized by new procedures and technologies. Nevertheless, there is room for improvement: current human labor-intensive inspection methods are expensive, time consuming, and place the inspectors in some danger to life and limb. The effort which is reported in part in this paper aims at demonstrating that the technologies of robotics, automation, teleoperation, remote data display and imaging, and automated signal and image processing and understanding aiding the inspectors will lead to their doing at least as good a job as they do now, to do it more economically, and to do it with lower risk of human injury.

In Section II we describe in more detail the aircraft inspector's job. In Section III we describe ANDI, a small suction-

cup-based robot designed to deploy eddy current sensors more-or-less anywhere on an airplane's exterior. In Section IV we describe CIMP, a large-wheeled robot designed to deploy a 3D-stereoscopic vision system and other sensors on the crown of medium to large sized aircraft. In Section V we describe recent progress with image enhancement and understanding algorithms for cracks (and corrosion) in CIMP's data stream. In Section VI we draw conclusions and indicate future directions for this work.

II. THE AIRCRAFT INSPECTOR'S JOB

Some commercial airlines have separate organizations for performing instrumented nondestructive inspection (NDI) and visual inspection (essentially naked-eye), others employ the same staff for both. The military sector generally relies more heavily on nondestructive inspection, and does not always explicitly employ visual inspection, although military NDI, maintenance, and repair crews, *de facto* in the course of everything else they do, perform extensive visual inspection[1]. The following two subsections discuss the tools and procedures employed respectively in nondestructive and visual inspection.

A. Nondestructive Inspection

Nondestructive inspection includes eddy current sensing, ultrasonic sensing, x-ray methods, and a variety of visualization techniques such as shearography, thermography, and D-Sight (a commercial low angle illumination system)[2]. In addition, it includes dye penetrant and magnetic particle methods that fall on the borderline between NDI and visual inspection. Of these methods, at least for skin inspection, eddy current sensors dominate by far. Small pencil probes (with a single miniature coil serving as both transmitter and receiver) and the somewhat larger pitch-catch probes (with separate transmitter and receiver coils) are the mainstay sensors, and cardboard or flexible plastic sheet templates are the mainstay guidance devices. However, with users' increasing expectations for image-like displays, state-of-the-art alternative sensors[3] that display areal images obtained either by mechanical raster scanning of a single probe, linear scanning perpendicular to the length of a linear array of miniature probes, or as a mosaic of areal patches obtained from a two dimensional array of miniature probes, or from an inherently areal magneto-optical probe, are increasingly being seriously considered.

B. Visual Inspection

The visual inspector's tools[8] are his eyes, a high quality flashlight, and a 10x magnifier. There are few standard inspection procedures, and the inspectors do not usually receive any specific training in visual inspection (although they are trained in the applicable regulations and required paperwork). Rather than being "created" by a training process, the visual inspectors are typically drawn from the ranks of mechanics who have worked for ten to fifteen years, and who have exhibited a talent for spotting subtle signs of problems. Their individual talents are sometimes highly specialized: one may be particularly attuned to small cracks, another to rivets that are not set quite right, another to the surface pillowing that accompanies subsurface corrosion. They recognize their own and their co-workers talents, they frequently call each other in for "consultations", and they often organize themselves into teams with complementary strengths.

Nevertheless there is considerable commonality in the way they actually work. Most seem to have a knack for spotting a problem area at a distance. Moving in closer, they first swing the flashlight in an arc to dynamically illuminate the surface at a low angle from a variety of directions. Then, having narrowed the suspicious area to a patch about the size of a coin, they use the 10x magnifier to get very close to the apparent flaw. Again, moving the flashlight to achieve dynamic lighting, they are able to distinguish, e.g., a harmless surface scratch from a structure threatening crack. The scratch/crack distinction is typically based on understanding that light disappears into a crack, but returns specularly from some part of a scratch.

Some specialized equipment, including long-handled mirrors and illumination devices, is used to visually inspect difficult-to-reach interior spaces, and borescopes are routinely used for engine and similar interior inspections, but except for a small pick to remove a paint chip or to poke at a possibly loose rivet head, etc, the flashlight and magnifier are the only tools used in visual skin inspection.

III. THE AUTOMATED NONDESTRUCTIVE INSPECTOR (ANDI)




Fig. 1: Mechanical drawing of ANDI. Only key skeletal components are shown in this figure. Compare with the photograph, inset to Fig. 2

ANDI, the *Automated NonDestructive Inspector*¹[4] was designed in the late 1980s and early 1990s in response to the FAA's interest, in the aftermath of the Aloha incident of 1988, in improving the speed and reliability of eddy current inspection of large numbers of rivet rows that might be subject to multiple site cracking damage. Its mechanical design, known in the robotics literature as a "beam walker", is well suited to moving a pencil probe in a precisely shaped (generally straight line) and precisely oriented path. ANDI is teleoperated via a graphical user interface (GUI) that permits both low level commands to individual actuators and higher level coordinated control of multiple actuators via an on-board microprocessor.

ANDI carries four cameras for alignment and navigation. Two cameras are mounted fore and aft, looking down at the skin with a field of view that encompasses typically five or six rivets. Shrouds that eliminate ambient light and low angle illumination that accentuates rivets make it possible, with high reliability, to identify rivets and to abstract lines of rivets, and thus to align ANDI fore and aft, with appropriate parallel offset, with the line of rivets under inspection. This image collection and processing is incorporated in a closed loop algorithm that automatically aligns and guides ANDI. The third camera is mounted on the eddy current sensor end effector, giving the operator a view of the sensor and its target. An inspector/operator could use it to obtain visual confirmation of anomalies suggested by the eddy current signal. The fourth camera, mounted above the tail end, has several proprioceptive functions, primarily confirming that the actual configuration of ANDI's body parts corresponds to the commanded configuration, and providing the driver's view for navigation, obstacle avoidance, etc. The cameras are not intended to have sufficient resolution for visual inspection for cracks; however they are potentially useful for more opportunistic kinds of visual inspection, e.g., for spotting the typically centimeter-sized splatter marks of a lightning strike, for noticing dents where a food service cart has bumped the airplane, etc.



Fig. 2: Photograph of ANDI on nose of DC-9 at Aging Aircraft Nondestructive Testing Center.

IV. THE CROWN INSPECTION MOBILE PLATFORM (CIMP)




Fig. 3: Overview photo of CIMP operating on a Boeing 747. See closeup of sensor pod in Fig. 4.

CIMP, the *Crown Inspection Mobile Platform* was designed in the early- and mid-1990s in response to our realization that to move automation technology for aircraft inspection from demonstration to commercialization it is essential to provide an integrated system the heart of whose sensor suite is a high quality remote enhanced visual inspection capability. This capability is provided by a 3D-stereoscopic video system that provides each of the inspector's eyes with its own full color, full NTSC/VGA resolution video stream updated at the full standard 30 frames-per-second². The effect, in addition to the perceptual value of binocular vision and stereopsis, is an increase in perceived spatial resolution because the noise on the two video streams is uncorrelated.

1. Also, for our namesake Andrews, Carnegie and Mellon.

2. This is achieved using special hardware that deinterlaces, buffers, and multiplexes the two camera streams for display on a single 120 Hz VGA monitor viewed through LCD shutter glasses synchronized with the view alternation. However the details need not concern us here.




Fig. 4: Closeup photo of CIMP sensor pod.
Note low angle variable azimuth illumination source at right.

CIMP was tested and demonstrated on a Boeing 747 at Northwest Airlines and on a McDonnell-Douglas DC-9 at US Airways. Several aircraft inspectors operated CIMP and viewed the 3D-stereoscopic images of the aircraft, and of test samples from their lab. Although they were split more-or-less evenly in their estimates of the economic value of remote visual inspection, they unanimously agreed that the perceptual quality of the imagery was more than adequate for the task. Two apparent flaws were located during one of these demonstrations, although subsequent probing (by the inspector who saw it initially in the video) with, in one case, a pick, and in the other an eddy current sensor, proved that one was a paint chip, the other a paint crack.

Despite our initial concern that inspectors would be hostile to any suggestion that the computer could assist them in identifying flaws, or perhaps eventually it could even find flaws autonomously, they actually

made these suggestions and encouraged us to test their feasibility. The following section summarizes our progress in this direction.

V. AUTOMATED IMAGE ENHANCEMENT AND UNDERSTANDING

A. Inspection Console

The inspection console consists of a video monitor that displays CIMP's live high-quality stereoscopic imagery of the aircraft surface, a Silicon Graphics Indy workstation that provides image capture and access to a library of image enhancement and understanding algorithms, and a radio control unit through which CIMP's motion and imaging sensors are controlled remotely. The video monitor provides flicker-free, full spatial and temporal resolution per eye, stereoscopic imagery from either the inspection or proprioceptive cameras. The workstation provides a graphical user interface (GUI), the Intelligent Inspection Window (IIW), that displays live monoscopic or still stereoscopic imagery of the inspection surface. The IIW is also an interface to the library of image enhancement and understanding algorithms. The inspector, by specifying parameters on the IIW, can choose and customize the algorithms for enhancing and detecting defects in the particular image being displayed. Fig. 5 displays the live video station and Fig. 6 the IIW display screen.



Fig. 5: Aircraft inspector observing the live stereoscopic imagery on the video monitor using active eyewear

B. Image Enhancement



Fig. 6: Intelligent Inspection Window (IIW) display screen.

Image enhancement algorithms amplify high spatial frequencies of the imagery, highlighting surface features such as scratches, cracks, lightning strikes and corrosion that are typically rich in high spatial frequencies. Since these algorithms have fast implementations, they allow the inspector to quickly identify high spatial frequency areas of the image that are most likely to contain defects. The enhancement algorithms are also useful to compensate for poor image quality that may arise occasionally due to non-ideal imaging conditions. For example, features of the live remote imagery that cannot be discriminated by naked eye due to poor lighting, motion of the cameras, insufficient depth-of-field, etc, can sometimes be seen clearly after application of an appropriate image enhancement algorithm.

We have developed a library of monoscopic and stereoscopic image enhancement algorithms[7]. Algorithm selection, and the extent of high spatial frequency amplification applied, is controlled by the inspector via buttons and sliders on the IIW control panel. Several options are discussed in the following sections.




Fig. 7: Flow chart for high-frequency emphasis algorithm.

B.1 Monoscopic Image Enhancement

We have developed three monoscopic enhancement algorithms: high frequency emphasis, sharpening, and adaptive histogram equalization. Here we describe only the high frequency emphasis algorithm.




Fig. 8: Live video image (monoscopic view taken from one of the stereoscopic camera pair).

High-frequency emphasis is a two stage process wherein a high-pass spatial filtering operation is performed first, then a fraction of low-frequency content of the original image is added back to the high-pass filtered image. This sequence amplifies high frequency features in the image while attenuating the background low pass information, resulting in the former (suggestive of flaws) being highlighted. The algorithm is illustrated in Fig. 7. This method has proved to be quite useful in highlighting scratches, cracks and corrosion texture, since these defects are typically of high frequency nature. The gain of the high-pass filter is controlled by the sensitivity slider on the IIW. Fig. 8 displays a live image and Fig. 9 an image with a mild degree of high-frequency emphasis. Note the emphasized scratches on this image.

B.2 Stereoscopic Image Enhancement

We have implemented two stereoscopic image enhancement algorithms on the IIW. They are the stereoscopic high-frequency emphasis and the augmented-stereoscopic high-frequency emphasis algorithms. The former algorithm performs a monoscopic high frequency emphasis on the left and right images of the stereoscopic image. The latter algorithm synthetically exaggerates the stereoscopic information content present in the original images. The processed images from these algorithms are displayed stereoscopically on the IIW display screen. The inspector views these images stereoscopically through active eyewear that is synchronized with the display. We have included, in the IIW, controls that adjust the horizontal and vertical disparities between the left and right displayed frames. The inspector, by adjusting these controls to a desired level, can comfortably view the enhanced imagery stereoscopically. We briefly describe the augmented-stereoscopic high-frequency emphasis algorithm.




Fig. 9: High-frequency emphasized image.





Fig. 10: Left and right augmented-stereoscopic high-frequency emphasized images.

An augmented-stereoscopic image is created by synthetically changing the disparity of features of interest in the stereoscopic image. These features then appear, when viewed stereoscopically, to loom above the background. The augmented-stereoscopic high-frequency emphasis algorithm is composed of three steps: (1) the left and right images of the stereoscopic image

are high-frequency emphasized as described in section B.1; (2) the filtered images are dynamically thresholded to identify the features of interest; (3) the identified features are overlaid on the original left and right images with the desired depth offset. The IIW slider controls the depth offset. The result is illustrated in Fig. 10. Note the change in spatial location of edge-like features in the images such as rivet outlines and scratches.

C. Image Understanding

The image enhancement algorithms discussed in section B. above emphasize high spatial frequencies in the image. Since aircraft surface defects such as cracks and surface corrosion are inherently of a high frequency nature, applying enhancement algorithms to imagery results in these defects being highlighted. However, also included in the set of emphasized features are normal features such as scratches, rivet heads, paint cracks and color transitions, discoloration, dirt, etc, which are also of high frequency natures. Hence, it is left to the inspector's judgement to separate the skin defects from the normal features. The goal of image understanding algorithms is to combine the image enhancement capability discussed above with an algorithmic equivalent of the human inspector's interpretation skills. The capability of an algorithm to discriminate defects from normal features of similar appearance depends on correlating features of live imagery with prior or learned knowledge of actual defects. A high correlation of a feature in the live imagery with a defect type will result in the feature being classified as a defect of the correlated type. However actually creating a robust algorithm that detects defects remains an open challenge. Considering economic and human factors issues, one plausible scenario for application of image understanding algorithms in remote visual inspection is screening large volumes of image data. The image understanding algorithm can conservatively label all plausible defects, so that the inspector separates a larger fraction of actual defects from normal features in a smaller volume of raw data. Another scenario is the interactive use of these algorithms by inspectors to obtain a second opinion about a particular suspicious flaw. It is with these scenarios, less demanding than the scenario that asks the algorithm to identify all the flaws and misidentify none of the normal flaw-like features, that we emphasize in the current work described in the following sections. We discuss progress towards a algorithm for surface crack detection. We have made similar progress toward an algorithm for visible surface corrosion detection; however space limitations preclude discussing it in this paper.

C.1 Surface crack detection algorithm

C.1.a Background

The crack detection algorithm that we have developed is modeled closely on the widely practiced test for detection of cracks using directional lighting[8]. The crack detection algorithm restricts its search to the immediate neighborhood of rivets, in a manner similar to the practice of inspectors. This is reasonable because cracks typically grow outward from rivet holes. To take advantage of the differences in specular appearance of cracks and scratches, the crack detection algorithm uses multiple images obtained under systematically varying lighting conditions³. This is analogous to the inspector's use of a manually scanned flashlight to dynamically highlight features that discriminate cracks and scratches. We simulate the directional lighting produced by the inspector's flashlight with a remotely controlled rotatable directional light source inside CIMP's sensor pod. With this control, the inspector can remotely rotate the directional light source around rivet locations of interest, and obtain a set of images similar to those obtained by rotating a flashlight on the inspection surface. The inspector can then either manually, or by the application of the crack detection algorithm, examine the set of images for cracks. Fig. 11 displays a block diagram of the crack detection algorithm.




Fig. 11: Block diagram of the surface crack detection algorithm.

3. At present, our algorithm operates monoscopically on the right or left member of the stereoscopic image pair.

C.1.b Rivet detection and region of interest (ROI) calculation

The first step of our crack detection algorithm is to find rivet locations in the image. Since cracks appear in the neighborhood of rivets, by finding rivet locations the algorithm is able to focus on areas that are most likely to contain cracks. Rivets are identified by detecting the circular arc edges made by their heads. We assume rivets to be the only objects on the inspection surface with circular edges. The rivet detection algorithm consists of the following steps:

1. Find all the edges in the image, discarding edges that are less than a minimum length, e.g., 10 pixels.
2. Fit a circular arc to each edge, calculating the fitting error and discarding all edges with more than a minimum fit error. The remaining edges are taken to be circular arcs belonging to rivet heads.
3. Merge the edges that are close to each other to form groups of circular edges, each belonging to a distinct rivet. The number of distinct groups found is equal to the number of rivets present in the image. Calculate the centroid (i, j) for each group.

We define the region of interest (ROI) as the neighborhood of the rivet that is to be examined for cracks. For each detected rivet, a corresponding ROI is defined as

$$ROI = \{(i, j) \mid i_0 - os \leq i \leq i_0 + os, j_0 - os \leq j \leq j_0 + os\}$$

where (i_0, j_0) is the centroid and os is an offset. The pixel range is bounded by the image size $(0: M-1, 0: N-1)$.

C.1.c Multiscale edge detection

An ROI in the live image contains a large number of edges, most of which are caused by rivet edges, scratches, dirt marks and occasionally real cracks. In order to separate cracks from non-cracks, we need an analysis framework with assigns to each edge a feature vector that allows for the discrimination of the small fraction of edges belonging to cracks from the many edges that belong to non-cracks.

A crack is typically very small compared to other objects present on the aircraft surface such as rivets, scratches, etc. Therefore the size of an edge is a important feature for discriminating cracks from non-cracks. It is also of interest to note that the edges corresponding to cracks are typically of the step-edge type due to their small edge width, while edges corresponding to non-cracks are mostly of the ramp-edge type.

These observations motivated us to select a multiscale (multiresolution) edge detection framework for the detection and analysis of edges in the ROIs. Multiscale edge detection is defined as the detection of edges at multiple scales or equivalently multiple resolutions[9]. Scale implies the size of the neighborhood in which the intensity changes are detected for edge determination. In multiscale edge detection, edges belonging to small objects appear only at low scales or high resolutions while edges of large objects appear at higher scales or low resolutions. Hence, the multiscale edge detection process allows us to characterize the size of each edge by assigning a relative size corresponding to the object that created the edge. In addition, Mallat and Zhong[10] have shown that the local shape of an edge can be characterized within the multiscale framework, by analyzing the evolution of wavelet⁴ local maxima across scales[9]. This is of importance for crack characterization and detection, since cracks are typically of the step-edge type. In short, we have chosen the multiscale edge detection framework because it allows us to estimate the size of the object that created the edge and also describe the type of the edge. Both are important features in discriminating cracks from non-cracks.

Multiscale edge detection is a two step process wherein the ROI is first decomposed into different resolutions by successive smoothing, followed by edge detection at each resolution. We have selected wavelet based filters for the projection of ROI to different resolutions and the estimation of intensity variations for edge detection. The wavelet based filters were chosen for two reasons. First, it is easy to define filters corresponding to a certain scale using wavelets. Secondly, as mentioned above, the set of wavelet local maxima of an edge in different scales define the type of edge. Therefore the wavelet maxima values can be used to indirectly characterize edge types. We now show the use of wavelet filters in multiscale edge detection.

We have chosen the cubic spline and its derivative, the quadratic spline, described by Mallat[10], as our scaling and wavelet functions. Let Φ_s and Ψ_s be the scaling and wavelet functions at scale s . Assume, without loss of

4. Wavelets are time and frequency localized basis functions. See [10] and [11] for more details.

generality, that $f(x)$ is a 1D continuous and differentiable function. The wavelet transform of $f(x)$ is given by

$$W_s f(x) = f(x) \otimes \psi_s(x)$$

where the operation is a convolution. Now, the scaling and wavelet functions are related by

$$\psi_s(x) = \frac{d}{dx} \phi_s(x)$$

Substituting

$$W_s f(x) = \frac{d}{dx} (f(x) \otimes \phi_s(x))$$

Therefore the wavelet transform $W_s f(x)$ is the first derivative of the signal smoothed at scale s . This procedure is very similar to the classical edge detection process. The local extrema of the wavelet transform at scale s corresponds to the edges of the image at that scale.

Listed below are the steps taken to generate the multiscale edges within each ROI in our crack detection algorithm using the wavelet filters.

1. Filter each ROI with the filter bank shown in Fig. 12. This gives a dyadic scale (scale = 1, 2, 4) decomposition of the ROI. Filters h_0 and g_0 denote the filters corresponding to the scaling and wavelet functions at scale 1. Filters h_n and g_n used at dyadic scale $2n$ ($n > 0$) are derived




Fig. 12: Multiresolution decomposition of the live image.

from h_0 and g_0 . W_y^n and W_x^n denote the wavelet transform images at scale $2n$ in the y (row) and x (column) directions of the ROIs while r and c denote row and column filtering⁵. Calculate the magnitude M_i and angle A_i images for each scale of the wavelet transform images W_x and W_y using

$$M_i = \sqrt{W_{xi}^2 + W_{yi}^2} \quad \text{for } i=0,1,2$$

$$A_i = \text{atan}(W_{yi}/W_{xi}) \quad \text{for } i=0,1,2$$

2. Threshold each magnitude image $M_{i=0,1,2}$ using a dynamic threshold calculated using its histogram. Pixels above the threshold are as marked edge points.
3. Link edge points based on 8-neighbors if their corresponding angles differ less than a maximum angle. This produces edges that are smoothly varying in direction which are characteristic of natural edges such as cracks. The above process produces a set of edges in each ROI at dyadic scales 1, 2 and 4.




Fig. 13: Coarse-to-fine edge linking of edges A and B

C.1.d Coarse-to-fine edge linking

Multiscale edge detection described in the last section generates edges at a set of scales for each ROI in the image. We model an edge belonging to an object appearing at multiple scales as the propagation of the edges along scale space. We define the propagation depth as the number of scales in which the edge appears. We then assign a propagation depth value to each edge at scale 1. The propagation depth captures the size information of the

5. h_n (g_n) is formed by including 2^n-1 zeros between the coefficients of h_0 (g_0); $h_0 = \{0.125, 0.375, 0.375, 0.125\}$, $g_0 = \{-2, 2\}$.

object represented by multiscale edges. For example, edges of objects that are small will have a lower propagation depth than edges of objects that are large. The propagation depth of an edge is found using a coarse-to-fine edge linking process. The coarse-to-fine edge linking process attempts to trace an edge from a coarse resolution (large scale) to a finer resolution (small scale). Fig. 13 illustrates the coarse-to-fine edge linking process. Note that the edges A and B have propagation depths 1 and 2 respectively.

C.1.e Feature vector calculation

We are now in a position to assign a feature vector to each edge in scale 1. We have selected the following attributes of an edge found through the multiscale edge detection process.

1. Average wavelet magnitude of active pixels.
2. Propagation depth number.
3. Average wavelet magnitudes of any linked edges in scale 2 and scale 4.
4. Signs of $\sum(W_x)$ and $\sum(W_y)$ where W_x , W_y are the wavelet coefficients in the x and y directions of an active pixel at scale 1.
5. Number of active pixels.

Active pixels are defined as pixels belonging to an edge.

C.1.f Feature Classification




Fig. 15: Image containing real cracks.

The feature vectors are classified into one of two classes, cracks and non-cracks. We use a neural net with six inputs, one hidden layer with four elements, and one output neuron. The network is trained using backpropagation with momentum. We generated fourteen feature vectors of simulated cracks (fine black fibers of approximately 10 micrometer diameter) and thirty feature vectors of non-cracks corresponding to actual rivet edges and scratches. Seven of the simulated cracks and fifteen of the non-cracks were used to train the network. After



Fig. 14: Surface containing several simulated cracks.

1000 training cycles, the network was approximately 71.5% accurate in predicting cracks, with a 27% false alarm rate for the test set edges. The simulated cracks and the discrimination of cracks from crack-like non-cracks are illustrated in Fig. 14 and Fig. 16 respectively.




Fig. 17: Detected cracks (dark) and non-cracks (gray) in image containing real cracks.

Fig. 15 shows a section of an aircraft surface containing two natural cracks and several scratches appearing in the neighborhood of a rivet hole. The first crack emanating from the rivet hole is ~0.5 inch long, while the second crack (which is partially masked by a scratch beside) it is ~0.3 inch long. Fig. 17 shows the result of applying the algorithm described. In Fig. 16 and Fig. 17, edges that are marked in black indicate suspected cracks, edges marked in gray indicate edges that the algorithm detects and classifies as non-cracks. The algorithm detects the known cracks in both images. It also correctly classifies the edges of the rivet hole and the scratches as non-cracks. The other edges that are marked in black are false alarms.




Fig. 16: Detected cracks (dark) and non-cracks (gray) in image containing simulated cracks.

VI. FUTURE WORK AND CONCLUSION

Based on our experience with the surface crack detection algorithm, we are convinced that the multiscale edge analysis framework on which the algorithm is based is suitable for extraction and analysis of surface cracks. We are encouraged by its performance in detecting real and simulated cracks despite the small sample of available cracks for training and testing. We have the following plans for further development of this algorithm:

1. Inclusion of additional features in the feature vectors that describe edges in a multiscale framework.
2. Providing multiple images of the same surface obtained under different lighting conditions to the algorithm to enrich the data set.
3. Obtaining a larger library of natural surface cracks for algorithm training and testing.

Another algorithm (that space limitations preclude our discussing here) has been similarly successful in detecting visible surface corrosion in test images; these results will be discussed in a future paper.

ACKNOWLEDGMENTS

We thank William M. Kaufman, Christopher Alberts, Christopher Carroll, Court Wolfe, John Hudak, John Harn, Caroline Perlee, and many other colleagues now or formerly at Carnegie Mellon Research Institute (CMRI), and Alan Guisewite, Gregg Podnar, Ian Davis, and many other colleagues now or formerly at the Robotics Institute, for their contributions to the research and development of ANDI, CIMP, and their instrumentation. We thank the FAA for its support of the ANDI project, and the Ben Franklin Technology Center of Western PA for its support of the CIMP project and Aircraft Diagnostics Corporation's efforts to commercialize it. And with immeasurable gratitude we thank Russell Jones at US Airways, Jeff Register at Northwest Airlines, their open minded supervisors, and the many dedicated members of their staffs who always cooperated cheerfully during our many visits to their hangars.

REFERENCES

- [1] Facts not explicitly referenced in this section are derived from "personal communications" with management, supervisory, and line personnel in the commercial and military aircraft inspection communities.
- [2] Ansley, G., S. Bakanas, M. Castronuovo, T. Grant, and F. Vichi, Current Nondestructive Inspection Methods For Aging Aircraft, FAA Technical Report DOT/FAA/CT-91/5, June 1992.
- [3] Beattie, A., L. Dahlke, J. Gieske, B. Hansche, G. Phipps, D. Roach, R. Shagam, and K. Thompson, Emerging Nondestructive Inspection Methods For Aging Aircraft, Sandia Report SAND92-2732. January 1994.
- [4] Siegel, M. W., W. M. Kaufman, and C. J. Alberts, Mobile Robots for Difficult Measurements in Difficult Environments: Application to Aging Aircraft Inspection, Robotics and Autonomous Systems, Elsevier, Vol. 11, pp 187-194, 1993.
- [5] Alberts, C. J., W. M. Kaufman, M. W. Siegel, and C. L. Wolfe, Robot with Cruciform Geometry, U. S. Patent issued 95-Jul-04, Patent No 5,429,009.
- [6] Seher, C., M. W. Siegel, and W. M. Kaufman, Automation Tools for NonDestructive Inspection of Aircraft: Promise of Technology Transfer from the Civilian to the Military Sector, Fourth Annual IEEE Dual-Use Technologies and Applications Conference, May 1994.
- [7] Gunatilake P., M. W. Siegel, A. G. Jordan, and G. W. Podnar, Image Enhancement and Understanding for Remote Visual Inspection of Aircraft Surface, Proceedings of the SPIE Conf. on Nondestructive Evaluation of Aging Aircraft, Airports and Aerospace Hardware, Vol 2945, pp 416-427, December 1996.
- [8] Bobo, S. N, Visual Nondestructive Inspection Advisory Circular, 1993.
- [9] Akansu, A. and R. Haddad, Multiresolution Signal Decomposition - Transforms, Subbands and Wavelets, Academic Press, Boston. 1992.
- [10] Mallat, S. and S. Zhong, Characterization of signals from multiscale edges, IEEE Trans. on Patt. Anal. and Mach. Intell., vol 12, pp 629 - 639, July 1990.
- [11] Chui, C. K., An Introduction to Wavelets, Academic Press, Boston 1992.