Improving the Reliability of Automated Non-Destructive Inspection

Nicholas BRIERLEY *, Trevor TIPPETTS *, Peter CAWLEY *
* Department of Mechanical Engineering, Imperial College London, London SW7 2AZ, United Kingdom, +44 (0)207 5947227

Abstract

Automated data acquisition is increasingly common in industrial NDE, driving a recent multiple increase in the availability of data. The data analysis remains a mostly manual task, performed by a skilled operator - a rather painstaking task given that much of data contains no indications. Partial automation, using software to prioritise regions of interest, could simultaneously increase inspection reliability and decrease data analysis time, by optimising the use of the operator's time. The project set out to produce such a software system general enough to fit a wide array of NDE applications, focussing on two specific examples: the ultrasonic inspection of power station rotor bores and the ultrasonic immersion inspection of aerospace titanium turbine discs. The paper presented reviews the progress made, identifies some of the obstacles encountered and demonstrates the developed capabilities using example software outputs.

Improving the Reliability of Automated Non-Destructive Inspection

Imperial College London

RCNDE

Nick Brierley
Trevor Tippetts
Peter Cawley

INTRODUCTION

In industrial NDE it is increasingly common for data acquisition to be automated. This typically gives rise to vast quantities of data, often including multiple interrogations of a given sample volume. The collected data need to be analysed and current processes are predominantly manual. The analysis is a laborious, slow process and does not exploit the available data to the fullest extent possible. There is therefore scope for the inspection reliability to be improved while reducing the time taken through partial automation, with significant potential cost savings.

While the project output is designed to be generic, the development work has considered two specific examples. Here we focus on one of these, provided by Rolls-Royce: an ultrasonic immersion inspection of aerospace engine titanium disks in a mid-manufacture, intermediate shape of rectilinear cross-section.

An aerospace engine titanium disk to be loaded into the immersion tank. The component is approximately 1m in diameter. Note the different surfaces of this intermediate form, which are all individually scanned, often at more than one probe angle.

The disk in the tank, positioned on a turntable below a manipulator arm used to position the ultrasonic probe.
To permit access to the under surfaces of the disk, the component is flipped half-way through the inspection.

The acquired dataset consist of hundreds of thousands of A-scans. The current analysis scheme relies on a global, fixed amplitude threshold for identifying indications. This fails to take into account variations in the achievable signal-to-noise across the component, and makes no attempt to combine signals acquired from different directions but relating to the same component region.

REGISTRATION

Registration, that is the alignment of datasets to a common coordinate system, is critical for any combination or comparison of datasets. Therefore this is the first major processing stage, consisting of the following steps:

- Feature extraction:
  Consistent features within the data channels must be selected for alignment, the disk's surface echoes are used here.
- Featuring pairing:
  Features from different channels are paired, on the basis that one scan's front surface is another scan's back surface.
- Error metric evaluation:
  Metrics quantifying the quality of registration fit are defined and evaluated for each pair of features.
- Optimisation:
  A multi-objective optimisation is run to find a good alignment of the data channels.

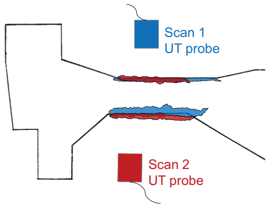

Part of a disk cross-section, showing how surface echo features of scans acquired from different directions may be paired during registration.

The output of the optimisation shows that it is not possible to match all surfaces perfectly, simultaneously. The trade-offs that exist between fitting different surface pairs reflect the inherent uncertainty in the data.

The achievable spatial accuracy of this process is around 2-3mm, consistent across a wide range of possible test displacements. This scale is comparable to the spatial correlation lengths of the data.

DATA FUSION

The data fusion stage that follows registration is designed to fully exploit the available to data to detect possible indications, for sequential examination by a skilled operator, and ultimately improve the Probability of Detection (POD) without incurring a Probability of False Alarm (PFA) penalty. The processing steps are outlined below.

- Data de-correlation:
  Generally, statistics requires independent, identically distributed (i.i.d.) samples. However, the raw samples recorded will be correlated due to e.g. beam-spread, sample microstructure and signal processing effects. To obtain approximately i.i.d. samples the dataset is broken down into basic volume elements of a physical size based on the correlation lengths of the data and the achievable registration accuracy. These elements are termed Resels (Resolution Elements).

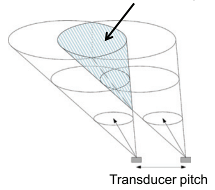

Illustration of how beam-spread introduces correlations into signals.

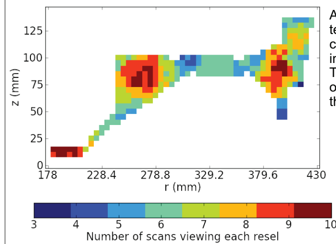

A radial cross-section of the disk, displayed in terms of Resels (each about 5mm across), counting the number of times each is interrogated over the full set of scans acquired. This then provides a map of data fusion opportunities and indicates the extent to which the sensitivity will be a function of position.

- Local data model:
  Variations in coherent noise level mean that global processing is unsuitable. Instead the component symmetry is exploited to examine each Resel relative to the distribution of Resels in the same hoop of the data. That distribution is then used to normalise samples by converting amplitudes to p-values, providing a measure of the likelihood that the signal is noise.

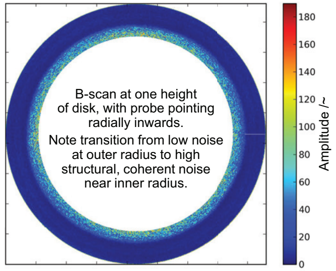

B-scan at one height of disk, with probe pointing radially inwards.
Note transition from low noise at outer radius to high structural, coherent noise near inner radius.

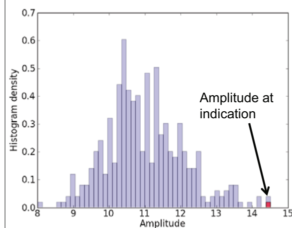

The local amplitude distributions for two scans viewing a particular known indication.
The data model applied is able to reflect the fact that the indication signal in the right distribution is much more clearly separated from the noise than in the left distribution.

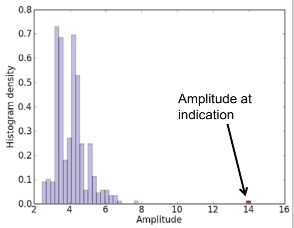

- Fusion by consensus test:
  Using a multivariate extension of simple hypothesis testing, p-values relating to the same spatial location may be fused to a single p-value by a consensus test, a tool widely applied in biological meta-studies.

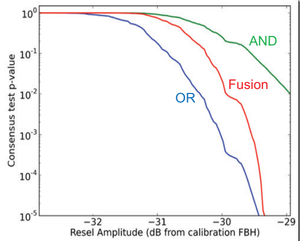

The output of the consensus test fusion as a function of indication signal amplitude, compared to the bounding combination possibilities provided by logical AND & OR operations. This shows that the fusion test transitions from behaving like an AND combination to act as an OR combination as the signal level increases. This is qualitatively ideal behaviour.

RESULTS

The performance of the developed program may be evaluated using the Relative Operating Condition (ROC) plot that relates POD and PFA.

The results shown were obtained for a dataset acquired from a test disk made of contaminated billet, containing multiple known indications. The first two plots relate to the most prominent known indication, the third to one that is much harder to detect.

First we consider only two input scans that view the indication of interest.

Both channels individually already outperform the fixed threshold approach by using locally derived thresholds, based on the local normalisation of amplitudes, significantly reducing PFA.

Fusing these two inputs provides a further improvement, demonstrating the benefit of the data fusion.

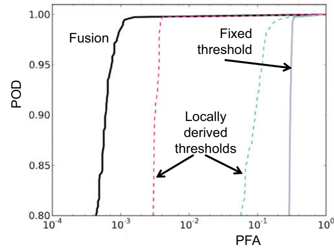

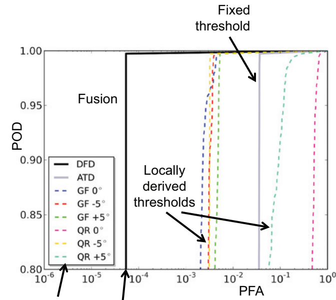

Legend details scan surfaces and probe angles

Effectively, PFA = 0 for these data

Then we consider all six available scans for the indication of interest.

Two of the channels are individually worse, even using the locally derived threshold detector, than the global, fixed amplitude threshold detection system.

However, fusing all the channels gives an ROC curve that effectively represents a perfect detector: a step-function at a PFA that is as low as can be numerically estimated from the available data.

Finally, we consider the five scans that view a different indication that is far more difficult to detect than the last.

Applying the locally derived threshold detection here, three channels are individually worse than the global, fixed amplitude threshold detector.

Fusing all the available channels still yields an ROC curve that substantially outperforms any of the input scans or the traditional detection system.

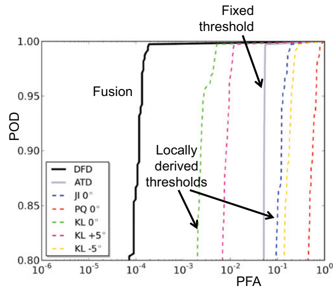

CONCLUSIONS

- **The developed analysis scheme, based on registration and data fusion, convincingly out-performs traditional approaches:**
  - Higher POD and / or lower PFA readily achievable
  - Analysis partially automated: more reliable, faster, lower cost
  - Process requires almost no changes to the data collection procedure
- **System needs to advance from a lab environment to the shop floor**