Damage Recognition for Structural Health Monitoring

Yujie Ying
Dept. of Civil and Env. Eng.
Carnegie Mellon University
Pittsburgh, PA 15213
yying@cmu.edu

Joel Harley
Dept. of Elec. and Comp. Eng.
Carnegie Mellon University
Pittsburgh, PA 15213
jharley@andrew.cmu.edu

Abstract

In the field of structural health monitoring, researchers focus on the design of systems and techniques capable of detecting damage in structures. However, it is difficult to develop robust detection schemes that are invariant to environmental and operational conditions. In this report, we investigate several signal processing and machine learning techniques for developing such robust systems.

From experimental data of a pressurized pipe, we extract 212 different data features and implement three different classification algorithms for detecting and localizing damage: adaptive boosting, support vector machines, and a combination of the two. The third algorithm shows the best overall performance in terms of accuracy, ranging from 81% to 100% in detection tests and 70% to 100% in localization tests. Through feature selection, we also demonstrate the effectiveness of features related to the Mellin transform and curve length.

1 Introduction and background

The fields of nondestructive testing and structural health monitoring (SHM) focus on the development of systems for testing the integrity in solid structures [1]. Such structures include, among other things, pipes, bridges, buildings, airplanes, and ships. These tests can be accomplished in many ways. Systems have been developed which, among other ways, evaluate structures by means of electromagnetic [2], vibration [3], and ultrasonic [4] testing.

We will focus on the use of ultrasonic guided waves produced by piezoelectric transducers. Ultrasonic guided waves travel through the thickness of the structure and are sensitive to structural changes caused by cracks, corrosion, or other forms of damage [5]. However, these waves are also sensitive to benign effects, such as changes in temperature [6,7] or air pressure. As a result, most traditional detection techniques fail under variable environmental conditions. In this report, we apply machine learning classification techniques to distinguish damage from more benign effects, primarily due to changes in pressure, in a steel pipe. We also reveal several informative and robust features for this application, such as Mellin transform descriptors and curve length features.

2 Experimental Setup

For our experiment, we use a synchronized pair of lead zirconate titanate (PZT) ultrasonic sensors to generate guided waves inside of a pressurized, steel pipe. Our pipe specimen is shown in figure 1a. We use a National Instruments PXI data acquisition device to excite a sinc pulse from one PZT and measure the resulting ultrasonic waves from the other PZT.

Data was taken in two collections, thirty-five days apart, denoted as ”Collection 1” and ”Collection 2”. Each collection consists of several sets, each with approximately 20 records of data. Each

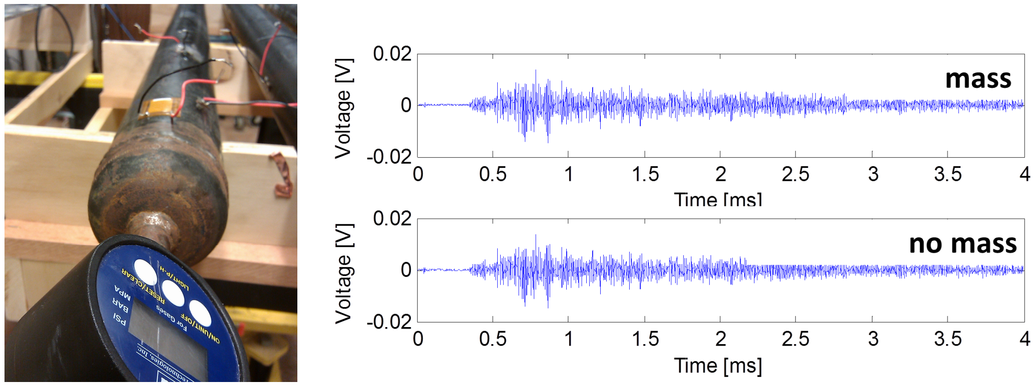

(a) Pressurized pipe specimen

(b) Examples of a complex ultrasonic guided waves with and without a mass on the pipe surface

Figure 1: Experimental setup and raw data.

record of data consists of $10,000$ samples of measurements, sampled at 1 MHz. In total, Collection 1 consists of 240 data records and Collection 2 consists of 360 data records.

2.1 Methodology

Before measuring each collection of data, the air in the pipe was pressurized to 100 PSI. Over each sets's 20 records, the pressure in the pipe was continuously discharged. This added uncertainty to the measurements so that no two records should be the same. Temperature may have also varied slightly between measurements. The first set of data is named ”U1” (undamaged).

After measuring the U1 data, a grease-coupled mass is placed onto the pipe to simulate damage. The second set of data is taken with the mass on the pipe and recorded as ”M1” (mass). The mass is then removed and another set of undamaged data is taken. This is done to accommodate for any changes due to the grease coupling used. This process of placing and removing the mass is repeated for six locations in Collection 1 and nine locations in Collection 2. All of the data sets may be clustered into three groups according to the location of the mass. Those locations are is near the PZT transmitter (Zone 1), at the center region of the pipe (Zone 2), or close to the PZT receiver (Zone 3). The data sets and the associated zones are shown in figure 2.

Figure 1b shows examples of two received signals, one from an undamaged data record in Collection 1, and one from a damaged record in Collection 1. From the figure, we can conclude at least two challenges in detecting damage in a pipe: first, the damage (mass) only causes very subtle changes in the signals and is difficult to distinguish; second, harmless environmental conditions may produce more changes over time and may masquerade as damage in the signal characteristics.

3 Feature extraction

For our analysis, we extracted 212 different features from each of our sets of data. These features were extracted using the concepts and tools such as the Mellin transform, the Fourier transform, correlation, principal component analysis, and the analysis of local maxima. In this section, we will briefly detail the features we extracted. In section 4.1, we will discuss which of these features are informative with respect to our problem.

3.1 Advanced signal processing technique: Mellin transform

For several of the informative features, we use a signal processing tool known as the Mellin transform. The Mellin transform is a tool that may be used for the analysis of scale behavior in a signal. The transform is often defined by

$$\mathcal{M}\left\{x\left(t\right)\right\} \quad = \quad X\left(c\right) = \int_0^\infty x\left(t\right) t^{-jc-1} dt\,, \tag{1}$$

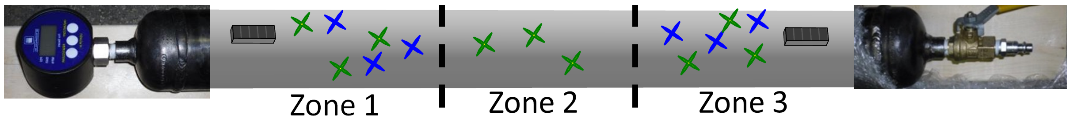

(a) Locations of masses on the pipe (blue crosses for Collection 1 and green crosses for Collection 2)

Figure 2: Mass locations and zone designations for localization experiments.

where j is the imaginary number and c is the dependent variable, sometimes referred to as the scale parameter [8].

The Mellin transform treats scaling behavior just as the Fourier transform treats shifting behavior. Due to this, we can effectively measure scale effects and generate features invariant to scale operations. This can be useful in the study of waves because uniform changes in wave velocity can be equated to a change scale of the received signal. In order to compute the Mellin transform of a discrete time signal, we use the Fast Mellin Transform algorithm. Details about implementing the Fast Mellin Transform can by found from reference [9].

3.2 Baseline-free features

Baseline-free detection of damage in structures is a subject of particular interest to the SHM community. A baseline is any measurement, or collection of measurements, used to remove features unrelated to the presence of damage. A baseline is usually useful due to the complexity of SHM signals. We consider 46 different baseline-free features.

3.2.1 Peak amplitude and location features

The peaks of a complex wave measurement are often important features of the signal. In the time domain, peaks often indicate the arrival of a new wave mode (a wave traveling at an independent group velocity and oscillating orthogonally to all other wave modes) or of a wave reflection. Therefore, we would expect certain peaks to be affected differently from other peaks when damage is introduced.

To construct features, we extract local maxima from the time domain, envelope of the time domain, and the magnitude of the frequency domain. Our features include the number of local maxima, the means and variances of the locations and amplitudes of the maxima, and the locations and amplitudes of the first 3 maxima. This is computed for all three domains.

3.2.2 Fourier and Mellin distribution features

For these features, we treat the magnitudes of the Fourier and Mellin transforms as probability distribution functions and extract the mean and variance values for those distributions. In the Fourier domain, any shift or scale in frequency may change the distribution of energy across frequencies. In the Mellin domain, the peak is dependent on the distribution of energy across the time domain.

3.2.3 Curve length

The curve length of a signal is useful for describing the signal complexity. A variation in curve length may be caused by changes in waves' modal amplitudes or locations. The curve length is also robust to time-scale changes since the signal's shape remains the same. Lu and Michaels [7] showed that the differential curve length between two signals is an excellent feature for damage detection.

3.3 Baseline-dependent features

It is often advantageous to use a known baseline signal, where we know there is no damage in the structure under test. We generate several baseline signals from the mean of 20 undamaged measurement records and from the principal components of those measurements. Principal component analysis is used to uncover certain properties of the signal that may better define the presence of damage. The following section will define several features that take advantage of these baselines.

3.3.1 Standard and invariant correlation coefficients

A very common metric for the similarity of two signals (a discrete-time baseline signal $x_b[n]$ and another discrete-time measured signal $x[n]$) is the correlation coefficient. Although the correlation coefficient is excellent at detecting differences in signals, it is generally not ideal for our applications. Benign effects often decrease the correlation coefficient by a greater degree than damage.

However, we may be able to modify the correlation coefficient formula to be more robust. In particular, we can generate correlation coefficients which are invariant to shifting and invariant to scaling. As alluded to previously, the magnitude of the Fourier transform is invariant to shifting and the magnitude of the Mellin transform is invariant to scaling. By correlating the magnitudes of the transforms, we can generate both shift and scale-invariant correlation coefficients.

3.3.2 Estimate of shifting or scaling

Scale-invariant or shift-invariant computations may be useful to gain robust against certain environmental effects. Therefore, estimating the degree to which a signal is actually a shifted or scaled replica of some baseline signal may also be informative. The estimate of a shift between two signals is defined by the parameter which maximizes the cross-correlation. Similarly, the estimate of a scale between two signals is defined by the parameter which maximizes the Mellin correlation. More information about Mellin correlation can be found in references [10] [8].

3.3.3 Differential curve length

Lu and Michaels [7] showed that the differential curve length was a good feature for damage detection. The feature is computed similar to the curve length measurement shown previously, but instead with the residual signal.

4 Experimental Results

In this section, we show the results of our damage detection and damage localization by using three classification approaches: adaptive boosting (AdaBoost), support vector machines (SVM), and a combined method of the two (AdaSVM). AdaBoost is an ensemble classification approach that linearly superposes a number of weighted ”weak” binary classifiers to generate a final ”strong” classifier [11]. The weak learners are usually simple and moderately inaccurate. Each weak classifier focuses on the misclassified instances of the previous classifier [12, 13]. SVM is a linear maximum margin classifier. It can provide a mapping of data into a higher dimensional space by applying kernel tricks so that a linear hyperplane or set of hyperplanes can be determined to classify data in that high dimensional feature space [14, 15].

One limitation of AdaBoost is that it can only linearly combine weak classifiers. So the final classifier may not necessarily be optimal [16, 17]. In contrast, SVM can effectively incorporate features through kernel functions. However, applying SVM to 212 features is not computationally efficient and some features may generate adverse effects in classification by adding noise. As a result, we demonstrate a combination method that uses Adaboost to select principle features, followed by SVM for classification. We define the principle features as the selected features used when the AdaBoost classifier reaches its lowest error rate after T iterations.

4.1 Feature Selection by AdaBoost

We apply AdaBoost to automatically select principle features. In figure 3, we demonstrate two examples of feature spaces containing the first two principle features produced by AdaBoost. In the first example (figure 3a) we use all of Collection 1 as the training data. AdaBoost ranks the top three features to be the curve length, shift invariant correlation coefficient of the second principle component, and the differential curve length with the 18th principle component as the baseline. In the figure, we see that the two classes, undamaged and damaged, are well separated. This demonstrates that the selected features are sensitive to the damage while robust to the environmental conditions. Overall, when the training set is extracted from Collection 1, our results show that the curve length and differential curve length features often rank as one of the top three principal features in detection tests.

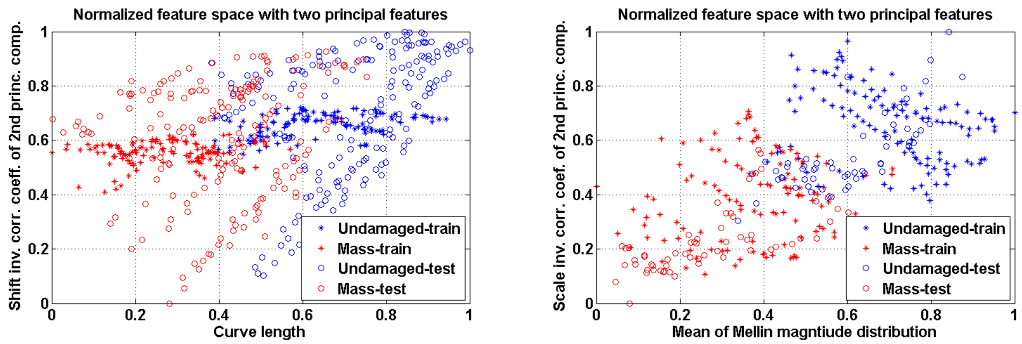

(a) Features, selected by AdaBoost, from Collection 1 (training set) and Collection 2 (testing set).

(b) Features, selected by AdaBoost, from Collection 2. Unique sets of data from Collection 2 were used as the training and testing sets.

Figure 3: The normalized feature spaces of the principle features, as selected by AdaBoost.

In the second example (figure 3b), the testing set consists of data sets $U1_2$, $U4_2$, $U7_2$, $M1_2$, $M4_2$ and $M7_2$ (120 data points), and the remaining 240 data points in Collection 2 are used for training. Again, the undamaged and damaged classes are very distinguishable under the first two principle features selected by AdaBoost: scale-invariant correlation of the signal with the second principle component and the mean of the Mellin magnitude distribution. Similar to the previous example, we have observed these that Mellin related two features are very effective in most of the detection and localization tests when the training sets originate from Collection 2. While most pairs of our extracted features do not provide any separation between the undamaged and mass data, our feature selection results support both Mellin transform and curve length features to be robust to environmental and operation effects on guided waves.

4.2 Results of damage detection

In order to detect and localize the mass, we employ three binary classification approaches, AdaBoost, SVM, and AdaSVM. For cross-validation purposes, we have conducted several tests with different training and testing sets. All of the trials are categorized into three groups: 1) the training and testing data originate from Collection 1, 2) the training and testing data originate from Collection 2, and 3) all of Collection 1 is used as the training data while all of Collection 2 is used as the testing data. Since we have the same number of undamaged instances as that of the damaged instances in our datasets, we convert the confusion matrices to the representation of accuracy, false-positive rate (FPR), and false-negative rate (FNR).

The Collection 1 and Collection 2 results (figures 4a, 4b) show that the three classification algorithms all obtain more than 85% accuracy, with AdaSVM providing 100% accuracy in Collection 1. When Collection 1 and Collection 2 are used together (figure 4c), SVM shows better performance than Adaboost and AdaSVM, with an accuracy around 86.7%. Figures 4d, 4e, 4f show these rates for each classification algorithm. It is evident that AdaSVM has an overall higher accuracy, lower FPR, and lower PNR. On average, the AdaSVM method improves the detection performance of AdaBoost by 4.3% and SVM by 2.1%.

4.3 Results of damage localization

In the damage localization tests, we apply the classification algorithms only to the damaged instances. Here we present the results of a binary-class localization problem. For multiple damage zones, we can utilize the same methods but for a hierarchical classifier structure. Similar to the detection trials, the localization tests are also divided into three groups, based on the training and testing data used. For all three types of trials, AdaSVM shows better performance than AdaBoost and SVM (figures 5a, 5b, 5c). The accuracy, FPR, and FNR over every trials with AdaBoost, SVM, AdaSVM are compared in figures 5d, 5e, 5f . On average, the AdaSVM method improves the localization performance of AdaBoost by 2.5% and SVM by 37.8%.

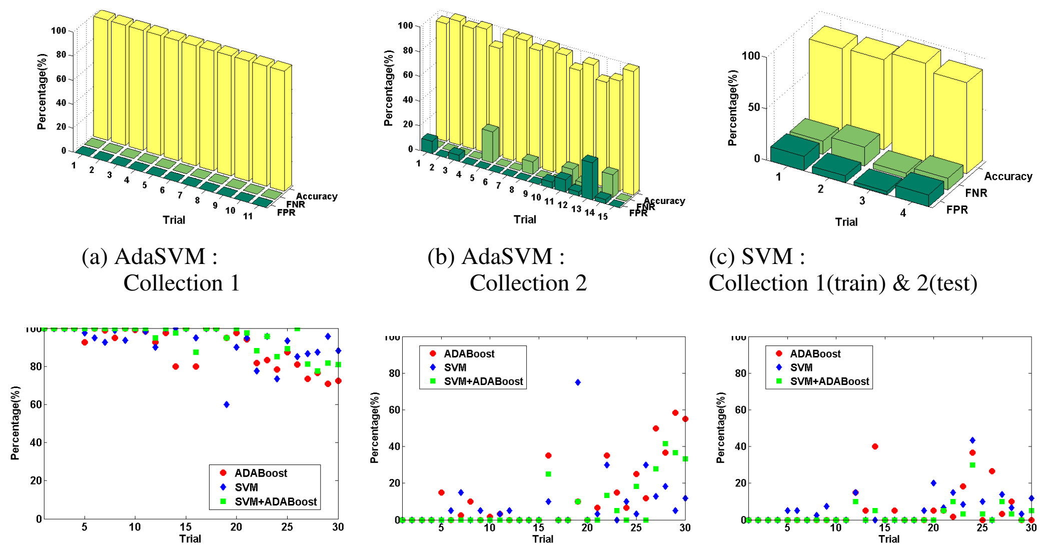

(a) AdaSVM : Collection 1

(b) AdaSVM : Collection 2

(c) SVM : Collection 1(train) & 2(test)

(d) Accuracy rates

(e) False negative rates

(f) False positive rates

Figure 4: Results for detecting damage

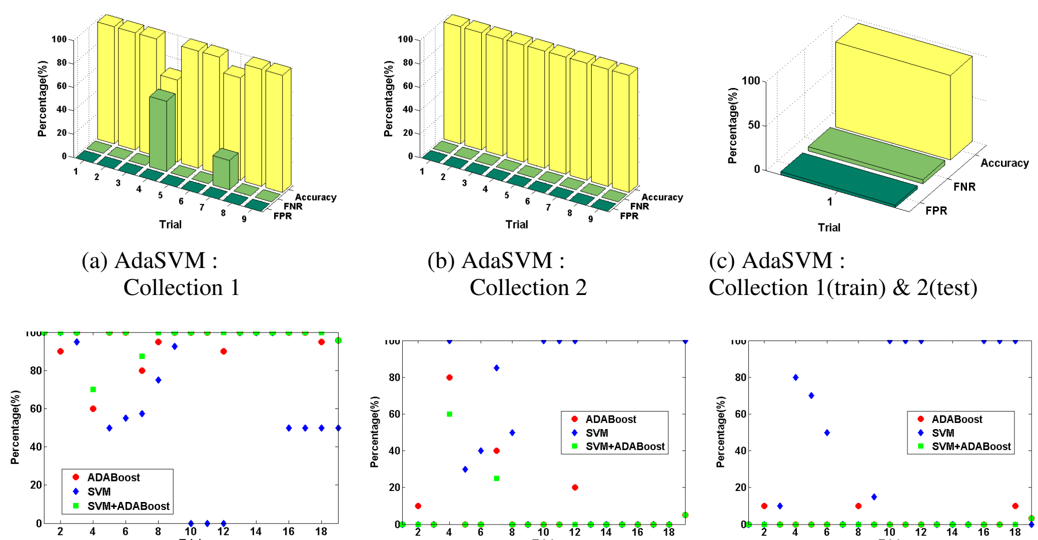

(a) AdaSVM : Collection 1

(b) AdaSVM : Collection 2

(c) AdaSVM : Collection 1(train) & 2(test)

(d) Accuracy rates

(e) False negative rates

(f) False positive rates

Figure 5: Results for localizing damage

5 Conclusions

Two sets of physical experiments were conducted on a pipe under varying operational and environmental conditions. This includes changes in the internal pressure of the pipe, changes in the ambient temperature, and several unknown changes and that may have occurred in the 35 days between data collections. Various signal processing techniques have been applied to extract 212 features. When evaluated by the AdaBoost algorithm, the curve length and Mellin domain features were largely selected as principle features for fine separation between classes. Three classification approaches, adaptive boosting, support vector machines and AdaSVM, have been investigated in order to detect and localize damage in the pipe. On average, the AdaSVM algorithm shows superior performance with average accuracy of 95.1% for damage detection and 97.5% for damage localization.

Acknowledgments

The authors would like to thank Jun Shi for assisting with the experimental measurements.

References

[1] K. Worden, C. R. Farrar, G. Manson, and G. Park, "The fundamental axioms of structural health monitoring," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 463, no. 2082, pp. 1639–1664, Jun. 2007.

[2] K. P. Chong, N. J. Carino, and G. Washer, "Health monitoring of civil infrastructures," *Smart Materials and Structures*, vol. 12, no. 3, pp. 483–493, Jun. 2003.

[3] H. Sohn, "Effects of environmental and operational variability on structural health monitoring." *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, vol. 365, no. 1851, pp. 539–60, Feb. 2007.

[4] D. N. Alleyne and P. Cawley, "The interaction of Lamb waves with defects." *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 39, no. 3, pp. 381–97, Jan. 1992.

[5] P. Wilcox, M. Lowe, and P. Cawley, "Mode and Transducer Selection for Long Range Lamb Wave Inspection," *Journal of Intelligent Material Systems and Structures*, vol. 12, no. 8, pp. 553–565, 2001.

[6] A. J. Croxford, J. Moll, P. D. Wilcox, and J. E. Michaels, "Efficient temperature compensation strategies for guided wave structural health monitoring," *Ultrasonics*, vol. 50, no. 4-5, pp. 517–28, Apr. 2010.

[7] Y. Lu, J. E. Michaels, and S. Member, "Feature Extraction and Sensor Fusion for Ultrasonic Structural Health Monitoring Under Changing Environmental Conditions," *IEEE Sensors Journal*, vol. 9, no. 11, pp. 1462–1471, Sep. 2009.

[8] L. Cohen, "The scale representation," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3275–3292, Dec. 1993.

[9] A. De Sena and D. Rocchesso, "A Fast Mellin and Scale Transform," *EURASIP Journal on Advances in Signal Processing*, vol. 2007, pp. 1–10, 2007.

[10] D. Casasent and D. Psaltis, "Scale invariant optical correlation using Mellin transforms," *Optics Communications*, vol. 17, no. 1, pp. 59–63, Apr. 1976.

[11] P. Viola and M. Jones, "Robust Real-time Object Detection," *International Journal of Computer Vision*, pp. 1–25, 2001.

[12] R. E. Schapire, P. Avenue, and R. A, *The Boosting Approach to Machine Learning An Overview*. New York: Springer, 2002.

[13] Y. Freund, "A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting,," *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119–139, Aug. 1997.

[14] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, Sep. 1995.

[15] V. N. Vapnik, *Statistical Learning Theory*. New York: Wiley-Interscience, 1998.

[16] L. Shen, L. Bai, D. Bardsley, and Y. Wang, "Gabor Feature Selection for Face Recognition using Improved AdaBoost Learning," *Advances in Biometric Person Authentication*, vol. 3781, pp. 39–49, 2005.

[17] J. H. Morra, Z. Tu, L. G. Apostolova, A. E. Green, A. W. Toga, and P. M. Thompson, "Comparison of AdaBoost and support vector machines for detecting Alzheimer's disease through automated hippocampal segmentation." *IEEE transactions on medical imaging*, vol. 29, no. 1, pp. 30–43, Jan. 2010.