Mechanical Systems and Signal Processing 117 (2019) 383–402

Contents lists available at ScienceDirect

Mechanical Systems and Signal Processing

journal homepage: www.elsevier.com/locate/ymssp

A probabilistic compressive sensing framework with applications to ultrasound signal processing

Ramon Fuentes [a,*], Carmelo Mineo [b], Stephen G. Pierce [b], Keith Worden [a], Elizabeth J. Cross [a]

[a] *Dynamics Research Group, The University of Sheffield, United Kingdom*
[b] *Centre for Ultrasonic Engineering, Strathclyde University, United Kingdom*

ARTICLE INFO

Article history:
Received 20 January 2018
Received in revised form 21 June 2018
Accepted 21 July 2018

Keywords:
NDT
Compressive sensing
Ultrasound
Signal processing
Relevance Vector Machine
Bayesian methods
Sparse representation

ABSTRACT

The field of Compressive Sensing (CS) has provided algorithms to reconstruct signals from a much lower number of measurements than specified by the Nyquist-Shannon theorem. There are two fundamental concepts underpinning the field of CS. The first is the use of random transformations to project high-dimensional measurements onto a much lower-dimensional domain. The second is the use of sparse regression to reconstruct the original signal. This assumes that a sparse representation exists for this signal in some known domain, manifested by a dictionary. The original formulation for CS specifies the use of an l_1 penalised regression method, the Lasso. Whilst this has worked well in literature, it suffers from two main drawbacks. First, the level of sparsity must be specified by the user, or tuned using sub-optimal approaches. Secondly, and most importantly, the Lasso is not probabilistic; it cannot quantify uncertainty in the signal reconstruction. This paper aims to address these two issues; it presents a framework for performing compressive sensing based on sparse Bayesian learning. Specifically, the proposed framework introduces the use of the Relevance Vector Machine (RVM), an established sparse kernel regression method, as the signal reconstruction step within the standard CS methodology. This framework is developed within the context of ultrasound signal processing in mind, and so examples and results of compression and reconstruction of ultrasound pulses are presented. The dictionary learning strategy is key to the successful application of any CS framework and even more so in the probabilistic setting used here. Therefore, a detailed discussion of this step is also included in the paper. The key contributions of this paper are a framework for a Bayesian approach to compressive sensing which is computationally efficient, alongside a discussion of uncertainty quantification in CS and different strategies for dictionary learning. The methods are demonstrated on an example dataset from collected from an aerospace composite panel. Being able to quantify uncertainty on signal reconstruction reveals that this grows as the level of compression increases. This is key when deciding appropriate compression levels, or whether to trust a reconstructed signal in applications of engineering and scientific interest.

1. Introduction

The Nyquist-Shannon theorem is at the centre of most traditional signal processing applications. It states that the frequency resolution obtainable in a signal is given by half of its time resolution, or sample rate. It is at the heart of frequency

* Corresponding author.
E-mail address: ramon.fuentes@sheffield.ac.uk (R. Fuentes).

https://doi.org/10.1016/j.ymssp.2018.07.036
0888-3270/

spectrum analysis methods based on Fourier and wavelet transforms. Compression is an important task in today's scientific computing. From a signal processing perspective, one of the most successful compression schemes is that of shrinkage, or thresholding [1] in a transform domain such as Fourier or wavelet. The problem with this, and any other compression scheme that relies on the application of a transform, is in the inherent wastefulness in the process of first acquiring a full data set in order to then compute the transform and the compression scheme.

Compression usually leverages sparsity: the idea that a signal that is dense in one domain can be sparse in another domain. The field of Compressive Sensing (often also called Compressive Sampling) (CS) was developed following the insight that compression can be achieved whilst also by-passing the usual procedure of first acquiring a full signal, and then transforming it into another "sparse" domain [2,3]. In the particular field of ultrasound signal processing, for example, it has already been shown that using a wavelet transform can achieve as much as a 95% compression ratio [4,5]. However, this requires both the acquisition of the full data set and the computation of a wavelet, or other transform. The main idea in CS is to circumvent the wastefulness of acquiring a large number of samples if one knows that most of them will be discarded anyway.

The contribution of this paper is a formulation of a probabilistic CS scheme. This is brought about through the use of the Relevance Vector Machine (RVM), a sparse Bayesian learning technique developed by Tipping [6]. The approach taken here is to replace the Lasso with an RVM in the sparse coding step of CS. This is a simple idea with profound implications. The result is a signal reconstruction that is fully probabilistic: it involves a mean and variance around the prediction, so that a confidence interval can be established regarding the quality of the signal estimate. Being a Bayesian method, it also naturally solves the problem of the appropriate level of sparsity with little user intervention. The estimation of uncertainty in the predictions is useful; it adds a layer to the understanding of prediction quality, which is paramount if compressing signals of scientific interest, or in safety critical applications.

2. Ultrasound-based NDT

The motivation for developing this CS framework is as an aid in the analysis of ultrasound waveforms, commonly used for Non-Destructive Testing (NDT). Ultrasound-based NDT has long been used in the structural integrity assessment of engineering components. The method, akin to the echo-location principle used by bats, relies on estimating the distance to an object by emitting a sound wave and listening to the response. The time it takes to receive the sound wave back can be used as a proxy to the location of the object, given some knowledge of the speed of sound properties of the medium. In order to use this principle to detect flaws in materials, one has to assume that a measurable amount of energy will be reflected back at the boundary between the medium and the flaw. Back at the source of the sound, this is measured as an echo. By mapping the time-of-flight (TOF) of these echoes across the surface of a material, it is possible to create a "depth map" otherwise known as a C-scan.

One of the characteristic features of ultrasound-based NDT is the acquisition of large quantities of high frequency data in the form of sound waveforms. Due to the high frequencies involved, often in the range between one and ten MHz, high sample rates are required to capture these waveforms, and this results in large quantities of data.

There is, however, a large disparity between the information content in a given ultrasound waveform, and the number of data points recorded in the time history of a waveform. So far, industry has solved this problem by extracting two key features from the echoes of ultrasound pulses: their attenuation and TOF difference. These features can yield useful information about material specimens and engineering structures if they are related to the material properties. The simplest and most widely used being the connection between TOF and material thickness, given the speed of sound of a material [7].

TOF estimation has therefore attracted significant attention from the NDT community. The methods developed over the years can be split into two categories: 1) those methods that use thresholds and changes in signal phase in order to separate the main pulse from the resulting echoes and compute the time differences between these two [8], and 2) those that use physical insight and attempt to solve a deconvolution problem to recover the impulse response function of the material being scanned [9,10]. The idea of estimating the full impulse response function of the material under question can be more attractive than characterising it with a few summarising features (such as a TOF) as this would capture all the information contained through the depth of a material. One interesting thing, from the point of view of this paper, is that the blind deconvolution problem is equivalent to the sparse coding step in compressive sensing, under an appropriate dictionary.

2.1. Features of ultrasound pulses

The Bayesian CS framework being presented will be demonstrated on ultrasound C-scan data from a carbon fibre composite wing panel. Although the results for this will be presented in Section 6, some key features of ultrasound signal processing will be introduced here, mostly as a motivation for the development of the method.

A typical ultrasound pulse is shown in Fig. 1a, with two time indices marked as t_a and t_b. These times correspond to reflections from the front and back wall of the composite panel respectively. A pulse of this kind effectively constitutes an A-scan. The information extracted from this is the time difference $t_f = t_b - t_a$, and this is often referred to as the ultrasonic TOF. This can be related to the thickness of the plate, if the propagation speed of bulk waves for the material is known. Another feature of interest is the ratio $x(t_a)/x(t_b)$ (where $x(t)$ is the measured amplitude of the ultrasound pulse), as this contains information