Digital Signal Processing 74 (2018) 14–29

Contents lists available at ScienceDirect

Digital Signal Processing

www.elsevier.com/locate/dsp

A generalised differential sparsity measure for reconstructing compressively sampled signals

Anastasios Maronidis *, Elisavet Chatzilari, Spiros Nikolopoulos, Ioannis Kompatsiaris

Information Technologies Institute, Centre for Research and Technology Hellas, P.O. Box 60361, 57001 Thermi, Thessaloniki, Greece

ARTICLE INFO

Article history:
Available online 6 December 2017

Keywords:
Signal sparsity
Compressive sampling
Differential sparsity
Signal reconstruction

ABSTRACT

Recent advances in signal compression, sampling and analysis have accentuated the importance of sparse representations of signals. A plethora of measures have been presented in the literature for estimating signal sparsity. In this paper, based on the concept that sparsity is encoded in the differences among the signal coefficients, we propose a novel parametric Generalised Differential Sparsity (GDS) measure and we rigorously prove that satisfies a set of objective criteria. Moreover, we prove that GDS interpolates between l_0 norm and Gini Index (GI), both of which prove to be specific instances of GDS, demonstrating the generalisation power of our framework. In showcasing the potential of GDS, we incorporate it in Simultaneous Perturbation Stochastic Approximation (SPSA) method and experimentally investigate its efficacy in recovering compressively sampled sparse signals. In the SPSA context, we prove that GDS, in comparison to GI, loosens the bounds of the assumed sparsity of the original signals and reduces the minimum number of compressive samples, required to guarantee an almost perfect recovery of heavily compressed signals. Finally, through a comparison with various sparse recovery methodologies, we show the superiority of SPSA+GDS in recovering both synthetic and real data.

© 2017 Elsevier Inc. All rights reserved.

1. Introduction

Sparse representation of signals has been introduced as a premise that permits the solution of problems previously unsolvable, paving the way to unprecedented possibilities in fields like signal compression and reconstruction. Roughly speaking, sparsity measures the extent to which the information of a signal is distributed to the coefficients. For highly sparse signals, the information is concentrated to a small portion of the coefficients, while for non-sparse signals the information is uniformly distributed across the coefficients. In this context, sparsity is a desirable property, as it allows for succinct representations of large pieces of information. Recall the Occam's razor, which dictates that among a set of representations, the most compact is always preferred [1].

There are many paradigms stemming from diverse research domains advocating the importance of sparsity. Compressive Sampling (CS) comprises the most vivid example, where the role of sparsity has been demonstrated in the process of compressing and reconstructing a signal [2]. More specifically, through the introduction of the Null Space Property (NSP) [3], and the Restricted Isometry Property (RIP) [4], it has been proven that under the assumption of data sparsity, it is possible to perfectly reconstruct a signal that has been compressed using only few random samples of the original sparse signal. Towards this end, a variety of optimisation algorithms, which incorporate the notion of sparsity, have been proposed. For instance, the Dantzig selector solves an l_1-regularisation problem in an attempt to estimate a ground truth sparse signal from few noisy projections of this signal [5]. In a similar vein, sparsity has also been utilised in the Lasso algorithm for recovering sparse representations of high-dimensional signals [6]. More recently, an interesting parallel implementation of compressive sampling matching pursuit (CoSaMP) offers a very efficient algorithm for sparse signal recovery [7].

Apart from the aforementioned applications, the notion of sparsity has also been exploited by already existing methods in various fields. For instance, it has been incorporated in traditional Bayesian learning methods for the recovery of block-sparse signals [8]. Moreover, in Support Vector Machines (SVM), optimal guarantees on the sparsity of the support vector set encoding the boundary between two classes, have also been investigated [9]. Sparsity appears to play a key role in boosting techniques as well, leading to sparse combinations of a number of weak classifiers [10]. Additionally, it has also found its way in other applications, such as one-bit compressed sensing [11], dictionary learning [12] and sparse principal component analysis (S-PCA) [13].

* Corresponding author.
E-mail addresses: amaronidis@iti.gr (A. Maronidis), ehatzi@iti.gr (E. Chatzilari), nikolopo@iti.gr (S. Nikolopoulos), ikom@iti.gr (I. Kompatsiaris).

https://doi.org/10.1016/j.dsp.2017.11.016
1051-2004/© 2017 Elsevier Inc. All rights reserved.

Given the importance of sparsity, it is essential to find an effective way to measure it. Apparently, the way sparsity is defined and measured is dictated by the specific purpose it is designed to serve. In this paper, in the context of CS, we are particularly concerned with the role of sparsity in the reconstruction of signals which have been heavily compressed using random projections. Signal reconstruction covers a large portion of the problems that concern sparsity. Hence, the conclusions drawn from our analysis are expected to have impact on other case studies as well.

Formally, the core idea of sparsity, as this has originally been introduced in CS, is to count the integer number of non-zero coefficients of a signal, measured by the l_0 norm [2]. In practice though, this proves to be a very strict definition, as rarely in real-world problems signals contain exact zeros. As a consequence, the research community has resorted to new relaxed measures of sparsity whose actual objective is to estimate an approximation of the number of non-zero coefficients, allowing sparsity to take decimal values. Along these lines, the notion of sparsity is usually referred to as signal compressibility. From now on though, in our work, we will consistently use the term sparsity even in the cases where we will actually refer to signal compressibility.

For assessing the credibility of any sparsity measure, a number of objective criteria have been proposed in the literature, enabling the comparison between different measures [14,15]. The origin of these criteria stems from the financial science, where the notion of sparsity is analogous to the inequity of wealth distribution in a human society [16]. So far, to the best of our knowledge, the only measures that satisfy these criteria are the Gini Index (GI) [14], and the S^* [15]. In particular, in connection with our work, the GI has led to impressive results in reconstructing compressively sampled signals [17].

In this paper, we propose a novel Generalised Differential Sparsity (GDS) measure, which is based on the differences among the signal coefficients. Due to an adjustable parameter, which from now on we call *the order*, GDS can offer different measure-instances. We rigorously prove that these GDS measure-instances satisfy all the objective criteria for sparsity measures [14,15]. In addition, although the computation of GDS using its original formula is tractable even for large values of its order, it proves to be cumbersome for high-dimensional data. For dealing with the above shortcoming, we provide an equivalent formula of GDS, which allows for its efficient calculation when the number of dimensions is high. The drawback though of the latter formula is that in contrast to the original one, it is costly for big values of the order of GDS. Consequently, both formulas prove to be useful and can be used interchangeably according to the given circumstances.

As part of our analysis, we prove that the order of GDS determines the tendency of the corresponding measure-instance to qualify an arbitrary signal as sparse. Moreover, interestingly, we prove that both GI and l_0 norm comprise measure-instances of GDS for the two extreme values of the order parameter, i.e. 1 and $+\infty$, respectively. This finding highlights the generalisation power of GDS in unifying and extending already existing measures, but most importantly shows that GDS constitutes an interpolation between GI and l_0. This proves to be a great advantage, offering GDS the flexibility to adjust to data whose coefficients are generated using several fundamental distributions, i.e., Binomial, Uniform, Normal and Exponential, where either GI or l_p norms may fail.

In verifying the above claim, we have used GDS to reconstruct sparse signals which have been heavily compressed via random projections. For this purpose we have employed the reconstruction approach presented in [17], which combined with GI has returned excellent results. The reconstruction is performed by incorporating a sparsity measure into Simultaneous Perturbation Stochastic Approximation (SPSA) method that solves a dedicated sparsity maximisation problem. More specifically, given a compressed signal and based on the prior assumption that the original signal before compression was sparse, the idea is to find in the original space, the signal with highest sparsity that gives the smallest reconstruction error.

In contrast to l_p norms which use the energy of a signal to estimate its sparsity, GDS could be characterised as an “entropy-based” sparsity measure, since through the use of the differences among the magnitudes of the coefficients, it indirectly uses the distribution of the energy to the coefficients. Conceptually, the above viewpoint of GDS may prove to be beneficial in the above-mentioned sparse signal reconstruction problem. More specifically, put in the above maximisation scheme as an objective function, GDS has the advantage to increase the sparsity of the compressed signal by re-distributing at each iteration the energy to the coefficients, thus changing its form, and not necessarily reducing the total energy of the signal. Based on this feature, GDS has the potential to extend the applicability of sparsity maximisation approaches to signals with all coefficients far from zero and non-uniformly distributed. On the contrary, l_p norms may miss the correct form of the original signal as they basically aim to reduce the energy of the compressed signal and the sparsity maximisation method relies strongly on the problem constraints to prevent signal-estimates arbitrarily close to zero.

The above potential to outperform l_p norms in sparse signal reconstruction has already been proven in [17] by GI, which builds on the same rationale as GDS. In this paper, we experimentally prove that incorporating GDS to the previous reconstruction approach, in comparison with GI, loosens the assumptions of both the underlying sparsity of the original signal and the minimum number of compressive samples, required to perfectly reconstruct an original signal from its compressed version. In other words, GDS offers further compression capacity to lowly sparse signals and simultaneously allows for using a smaller number of compressive samples without increasing the reconstruction error. Along the same lines, it is proven that the optimal order of GDS is strongly dependent on the type and sparsity of the original data as well as the desired compression level. This finding justifies the rationale behind using different values for the order of GDS and provides a useful rule of thumb in deciding what order of GDS is the appropriate for certain problem parameters.

Extending our analysis, through a comparison with numerous state-of-the-art sparsity measures, we prove the superiority of GDS in the context of SPSA and we propose SPSA+GDS as an effective method for reconstructing sparse signals. Towards this end, we compare SPSA+GDS in both synthetic and real data with three sparse recovery methodologies, namely Basis Pursuit (BP), Orthogonal Matching Pursuit (OMP) and Compressive Sampling Matching Pursuit (CoSaMP). The results illustrate the ability of SPSA+GDS to reconstruct signals from their compressive samples more accurately and effectively than conventional approaches, showing the strong potential of our proposed methodology.

2. Related work

2.1. Sparsity measures

As already mentioned, the most straightforward way to measure sparsity is through the l_0 norm [18]. However, in the context of CS, l_0 was very early replaced by l_p norms of higher order, surpassing some of the shortcomings accompanying the former [18]. Towards this direction, it has been proven that the classical error correcting problem can be translated into an l_1-optimisation problem, which can be trivially solved with linear programming methods such as the Homotopy method [19,20]. In [21], the authors propose a methodology for sparse signal recovery that often outperforms the l_1-minimisation problem by reducing the number

Download English Version:

https://daneshyari.com/en/article/6951850

Download Persian Version:

https://daneshyari.com/article/6951850

Daneshyari.com