ITERATIVE MIRROR DECOMPOSITION FOR SIGNAL REPRESENTATION

Fabrizio Guerrini, Alessandro Gnutti, Riccardo Leonardi

Department of Information Engineering, CNIT – University of Brescia, Italy

ABSTRACT

In this paper it is shown how to describe any finite-energy continuous or discrete signal through an ordered set of positions to uniquely represent it. This is obtained by designing an iterative decomposition through a series of mirror operations around those positions. The purpose is to find at any step of the decomposition the location that provides for the maximum decoupling between the even and odd components of the signal with respect to it. The algorithm can then be iterated at infinity determining a sequence of positions. The per location information determines the optimal energy decoupling strategy at each stage providing remarkable sparsity in the representation. Thanks to the sparsity of the resulting representation, experimental simulations demonstrate superior approximation capabilities of this proposed non-linear mirror transform.

Index Terms— Transforms, non-linear approximation, signal decomposition, sparsity, symmetries.

1. INTRODUCTION

Signal decomposition is at the foundation of information representation for analysis tasks such as classification, machine learning, detection, estimation [1, 2, 3], or communication tasks. In 1822, Joseph Fourier was the first to establish that a periodic wave could be represented as a linear combination of harmonic components of its fundamental frequency [4]. More so, the completeness of the representation converges at infinity in the mean square error sense towards the representation of discontinuous waveforms. About a hundred years later in 1909, Alfred Haar showed a converse result that a continuous waveform in $\mathcal{L}^2(\mathcal{R})$ can be represented as an infinite series of discontinuous functions [5]. More generally, it is common to consider how any Hilbert space of infinite dimensionality \mathcal{H} that is separable can be represented using some infinite set of functions representing its basis set. With the concept of frames the decomposition may be overcomplete and the representation becomes non-unique. Consequently, Fourier decomposition and/or multiresolution representations have been at the base of waveform representation.

Lately, more emphasis has been devoted to the study of alternative linear expansions in the representation of any vector x in \mathcal{H} :

$$\underline{x} = \sum_{i \in \mathcal{I}} \alpha_i \underline{\varphi_i} \quad \forall \underline{x} \in \mathcal{H}$$
(1)

learning different sets of dictionaries that would generate sparse representations considering the typical statistical distributions that involve any real phenomenon leading to particular classes of signals (such as natural images, individual ratings, population health, ...), allowing to possibly determine good alternative sparse expansions of waveforms [6], with further advancements for classification, compressive sensing, denoising, super-resolution, data recovery, and so on.

In a broader perspective, there have been attempts to model signal families as lying on manifolds. These approaches try to follow signal geometrical features [7]. This can be framed in the general context of dimensionality reduction, where signals laying in high dimensional spaces are approximated on a set of lower dimensional (sub)spaces. Instead of a signal representation induced by a series of projections on a (possibly overcomplete) basis set or through a sparse expansion on a dictionary, these dimensionality reduction methods do not achieve in general completeness since the signal can only be approximated without any guarantee on the convergence of the representation error. In addition, there is no general solution to the inverse problem of generating a signal given a manifold set, so such problem usually needs a regularization approach (*e.g.*, see [8]).

A final attempt sharing an approximate signal representation paradigm is based on contractive transformations. Iterative function systems (IFS) try to describe a signal through a set of base signals and a pool of contractive transformations so that their iteration converges to an approximate fixed point. Whereas IFS are quite attractive from a signal generation perspective, when the set of transformations is estimated for signal analysis purposes, the solution to the inverse problem is difficult and there exist only bounds on the approximation error to the original signal [9].

In this work, we want to preserve the idea of a converging representation such that any signal \underline{x} may be exactly recovered. However this representation will not come from a projection on the space (or a series of subspaces) spanned by a predefined sets of vectors. Instead, each signal is described through an iterative decomposition into a possibly infinite unique set of orthogonal components, that can be recombined through a chain of summations and mirroring operations. Also differently from IFS and dimensionality reduction methods, our approach is generative in the sense that the signal is exactly represented and the iterative transformation applied at each step is not found by solving an inverse problem but rather it is constructively generated by the characteristics of the signal itself.

At each stage of the proposed iterative transform, to achieve energy compaction the mirroring location is selected so as to provide maximal energy concentration into a single component. The iteration will thus lead to a remarkable energy compaction, so that a very good approximation of the original signal may be constructed from a limited set of its components. By truncating the decomposition at some level and by recombining only its constituent components providing for the highest energy concentration, the reconstructed signal determines the closest reconstruction of the original signal. In other words a much sparser representation can be determined if compared to many linear expansions thus providing a better approximation (in the \mathcal{L}^2 sense) to the original waveform.

The paper is organized as follows. The mirror transform is defined in Section 2 as the iterative even-odd decomposition modified to pursue maximum energy decoupling in each iteration step. The experiments in Section 3 were carried to show the sparsity properties of the mirror transform, and conclusions are drawn in Section 4.

2. THE MIRROR TRANSFORM

2.1. Optimal Symmetry Decomposition

The even-odd decomposition of an energy signal $x(t) \in \mathcal{L}^2(\mathbb{R})$ states that x(t) can be expressed as the sum of its even and odd parts $x_e(t)$ and $x_o(t)$. However, the classic definition of even-odd decomposition only considers the parity characteristics with respect to the support midpoint. The objective now is to modify the parity decomposition to allow it to track the inherent reflection symmetries possibly present in x(t) as best as possible.

Generalizing the even-odd decomposition around a given time instant t_0 we have:

$$x_e(t;t_0) = \frac{x(t) + x(2t_0 - t)}{2}; \quad x_o(t;t_0) = \frac{x(t) - x(2t_0 - t)}{2}$$
(2)

We need to find the global optimal symmetry point t'_0 , provided we first define what the optimality is [10]. If a signal is decidedly symmetric (or anti-symmetric) around a certain point t_0 , performing the even-odd decomposition using that t_0 as mirroring point would output the even and odd signals whose energies are in great disproportion. Therefore, it is natural to search for the time instant for which there is a maximum energy decoupling between the even and odd parts, so the attention should focus on their energies, E_e and E_o respectively, as t_0 varies.

To find t'_0 , let us concentrate on the extrema of the energy of the even part. Its value as a function of t_0 is:

$$E_{e}(t_{0}) = \int_{-\infty}^{+\infty} |x_{e}(t;t_{0})|^{2} dt = \int_{-\infty}^{+\infty} \left| \frac{x(t) + x(2t_{0} - t)}{2} \right|^{2} dt =$$
$$= \frac{1}{4} \int_{-\infty}^{+\infty} [|x(t)|^{2} + |x(2t_{0} - t)|^{2} + 2x(t)x(2t_{0} - t)] dt$$
(3)

where we can safely extend the integral on the whole real axis without affecting the result for finite-support signals. The first two terms in the last integral give E, the energy of x(t), since reversing the time axis and shifting the origin do not influence the energy value. Hence:

$$E_e(t_0) = \frac{1}{2}E + \frac{1}{2}\int_{-\infty}^{+\infty} x(t)x(2t_0 - t)dt$$
(4)

The usual definition of the linear convolution for energy signals is:

$$(x*y)(t) = \int_{-\infty}^{+\infty} x(t')y(t-t')dt'$$
 (5)

Therefore we can write (repeating the process for E_o as well):

$$E_e(t_0) = \frac{1}{2}E + \frac{1}{2}(x * x)(2t_0); \quad E_o(t_0) = \frac{1}{2}E - \frac{1}{2}(x * x)(2t_0) \quad (6)$$

So the energy of the even part is a function of t_0 dictated by the convolution of the original signal with itself – an "auto-convolution" (which is equivalent to the cross-correlation between the signal and its mirrored version). One could wonder if E_e admits a maximum, which in turn would be implied by the existence of the maximum of the auto-convolution. This can be readily seen by observing that such a function must be limited, as implied by the Cauchy-Schwarz inequality, and on the other hand the energy of the even part cannot be greater than that of the original signal. Furthermore, the auto-convolution is an integral function and therefore must be continue,

so E_e admits at least an extreme point. Of course, since there is no closed-form formula to find the exact position of the extreme points, we need to resort to numerical computation (which takes linear time), and then we need to test all the points to find the absolute maximum and the associated t'_0 .

2.2. Decomposition Tree: the Mirror Transform and Its Inversion

Let us now apply the just stated optimal symmetry decomposition in an iterative fashion to build a decomposition tree. Such a tree represents what we refer to as the mirror transform of the signal. First, let us assume for simplicity that the original signal x(t) has a [-T, T]finite support, and $t'_0 \in [-T, T]$ optimally decouples the energy of the even and odd part. Both the even and odd parts have total support $2T + 2|t'_0|$. However, both include a $2|t'_0|$ long part which is simply the tail of the original signal divided by 2 (see Fig. 1). Therefore, it is possible to isolate the tail part and treating it as a separate signal. The rationale behind this is that, in practice, the tail part has undergone no modifications and thus should be handled separately. Thus, retaining just half of the even and odd parts (the other halves are obtainable by mirroring), a single decomposition step applied to the root of the tree x(t) outputs three signals: a $T - |t'_0|$ long even signal, a $T - |t'_0|$ long odd signal, and a $2|t'_0|$ long tail signal (whose supports sum is 2T, the support of x(t)).

These three signals are ready to be decomposed again. They constitute the first-level nodes, $x_{\{\mathcal{E}\}}^{(1)}(t)$, $x_{\{\mathcal{O}\}}^{(1)}(t)$ and $x_{\{\mathcal{T}\}}^{(1)}(t)$ respectively of an approximately ternary decomposition tree. The tree is only approximately ternary because a) the tail may not exist in the case that t'_0 is exactly in the middle of the support or b) one of the nodes may turn out to be a zero-energy signal. The latter case could happen because the parent node is perfectly symmetric or anti-symmetric. So, searching for the optimal symmetry point is obviously meaningless for zero-energy signals, and the construction process of the decomposition subtree descending from any zero-energy node should be arrested since it has no meaningful purpose.

In general, a node in the level l of the decomposition tree is referred to as $x_{\{S_1,...,S_l\}}^{(l)}(t)$, where S_i represent labels taking either the value \mathcal{E} , \mathcal{O} or \mathcal{T} to refer to an even, odd or tail component. In principle, such decomposition can go on forever. The number of nodes on a given level l increases exponentially with the level number (ignoring the possibility of zero-energy nodes) and the support of the single nodes also tends to shrink accordingly. Each new level



Fig. 1: The result of the optimal parity decomposition of x(t), using t'_0 as mirroring point. Note the tail part on the extreme right of the [-T, T] support.

shrinks the support of the children nodes, dividing in three parts the support of the parent node, however the total support of all the nodes at a given level is still 2T.

Given a decomposition tree, the inverse transform, which is iterative and starting from the bottom layer of the tree, consists of first mirroring the even and odd signals around their support midpoint (retaining and changing the sign respectively), then adding them and finally concatenating the tail to obtain the parent node on the previous level, until the root is reached. For example, $x_{\{\mathcal{O},\mathcal{E},\mathcal{E}\}}^{(3)}(t), x_{\{\mathcal{O},\mathcal{E},\mathcal{O}\}}^{(3)}(t)$ and $x_{\{\mathcal{O},\mathcal{E},\mathcal{T}\}}^{(3)}(t)$ are combined to reconstruct $x_{\{\mathcal{O},\mathcal{E}\}}^{(2)}(t)$. Correspondingly, the same notation applies to the optimal symmetry point of the root x(t). Then for each successive level l, we define $t_0^{(l)} \{\mathcal{S}_1, \ldots, \mathcal{S}_l\}$ as the globally optimal symmetry point of the corresponding decomposition tree node.

As we mentioned before, the decomposition tree for continuoustime signals can go on forever. We state here an interesting property of such infinite decomposition. A signal x(t) is *uniquely* identified by the infinite succession of symmetry points $t_0^{(0)}$ and $t_0^{(l)}{S_1, \ldots, S_l}$ for $l = 1, \ldots$, besides a possible global scaling factor K that scalar multiplies the entire signal. If two different signals, by absurd, should share the same infinite decomposition, it would mean that all the auto-convolution signals in each level admit the maximum in the same position, which is impossible if the two signals do not differ by anything besides zero-energy signals.

2.3. Discrete-Time Domain

As is imaginable, iterating the decomposition in the discrete time setting is not as straightforward as in the continuous time case. The most important difference is that the symmetry point n_0 cannot be arbitrary but has to correspond to either a sample (*i.e.*, integer) or to a half-sample position. Given the discrete nature of the signal domain, the iteration is bound to stop when single-sample sequences are encountered, that we refer to as "leaves", a situation not found for continuous-time signals. Therefore, decomposition trees for discrete-time finite support signals are always finite, and for a *L*-long sequence exactly *L* leaves will be reached. Using the ternary decomposition process on discrete sequences can build a very unbalanced tree, as its *L* single-valued leaves may be all over the tree levels, depending of the particular n_0 found in each of the nodes.

The number of decomposition levels in the discrete-time domain is minimized by symmetry points found nearby the node sequences midpoints because the maximum length of the children even/odd sequences is limited by approximately half that of the parent sequence and the tail node is very short. On the other hand, symmetry points found near the end points of the original sequence generate very short even/odd sequences and a very long tail, almost as long as the parent sequence. If this latter case happens frequently in the decomposition process, the number of decomposition levels is bound to increase. In the extreme cases, if the global symmetry is always found in the midpoint of any given node the number of decomposition levels is $\lceil \log_2 L \rceil$, while if it is always put on the first or last sample the number of levels tends to L.

3. EXPERIMENTAL RESULTS

In this section some experiments conducted to verify the sparsity properties of the mirror transform are reported. To that end, the scope of the tests is of course limited for practical reasons to finitesupport signals in the discrete domain. In those cases, the decomposition tree has a finite number of levels.

Specifically, we have conducted our experiments on a variety of 1-D signal types. First, we generated some random white noise sequences using the **randn** routine in Matlab[®]. Next, we selected around 30 electrocardiogram (ECG) signals, obtained with a sampling frequency between 250Hz and 350Hz, taken from the MIT-BIH Arrhythmia Database [11]. We also prepared some audio clips of a few seconds each, from two sources: an audio signal from a speaking source and a live rendition of "The Star-Spangled Banner", both sampled at CD quality. Furthermore, we have included in our tests some seismic signals taken from the IRIS database [12]. Last, the 1-D signals database is completed by the image rows taken from classic images, coming in either 512×512 and 1024×1024 format, and from a collection of texture images taken from [13].

To measure the sparsity of the information provided by the decomposition tree leaves, we handled the decomposition process in the same way as one would with a standard processing transform. This way, we can compare our method with two other basic discrete signal processing transforms: the Discrete Cosine Transform (DCT) and the Discrete Wavelet Transform (DWT), implemented using both the 'db1' and 'db4' wavelets with 5 decomposition levels. The comparison has been done as follows: first, compute the whole decomposition tree (which is at least $\lceil \log_2 L \rceil$ levels deep) and then sort the L leaves according to their decreasing magnitude. Then, discard a given percentage P of the least significant leaves (in terms of magnitude) and compute the ratio of the energy of the error signal in the reconstructed sequence and that of the original signal. The same is built for the other transforms, where instead of the leaves we discarded the least significant (frequency or time-frequency) coefficients in terms of their magnitude.

In Fig. 2, the average results on 10^6 realizations of white Gaussian noise sequences with $L = 10^3$ (Fig. 2a) are reported. In addition, the standard deviation of the mirror transform performance in terms of SNR is also depicted in Fig. 2b, showing very tight intervals around the average result reported in Fig. 2a. The results are very similar for higher L (up to 10^6) as well. Noise sequences are notoriously hard to represent with frequency related transforms given their spread-spectrum frequency content. Instead, a transform such as the mirror transform effectively follows the locations where to decompose the signal so that in the end a better reconstruction of the noisy input is obtained. This happens since the mirror transform maps the waveform complexity into the coupling of few significant coefficients and of the locations $n_0^{(l)} \{S_1, \ldots, S_l\}$ combined with extreme flexibility by the tree decomposition. This does not in any way mean that the mirror transform is able to "compress" noise sequences, but that it can systematically find a small subset of



Fig. 2: Mean reconstruction SNR for noise signals as P varies on the x-axis.

those $n_0^{(l)}{S_1, \ldots, S_l}$ that allows a good reconstruction of the input noise, achieving a superior degree of sparsity. Of course, due to the random nature of the noise, even for the proposed transform the SNR is still lower with respect to those obtained for the other types of signals, which we analyze in the following.

Next, we show in Fig. 3 the results obtained for the rows of four images: *Lena*, *Baboon*, the satellite image of Richmond, VA (1024×1024) and a texture image (brick wall pattern). For the former classic images, the mirror transform performance is comparable with those given by the 'db4' wavelets, with actual values depending on the image characteristics (in particular, how many rows can be well approximated by polynomial functions). For images with strong presence of local symmetries like the satellite and texture images, the mirror transform performance is clearly superior.

We also show in Fig. 4 the visual appearance of a magnified detail of the *Baboon* image when P = 0.95 for the proposed transform and 'db4' DWT. In this case, we transformed separately each row of the images and then we have retained the most significant transform coefficient across the whole image. In accordance, the PSNR has been computed on the whole image instead of separately for each row. The obtained PSNR values are 29.99 dB for the mirror transform, compared to 26.63, 28.58 and 26.76 respectively for the DWT 'db1', DWT 'db4' and DCT. This time too, the DCT and the DWT performance are worse than those given by the proposed trans-



Fig. 3: Mean reconstruction SNR for image rows as *P* varies on the *x*-axis.





(a) Detail on mirror transform reconstruction.

Fig. 4: Visual results obtained with P = 0.95.

reconstruction.

form, given that the best looking result is clearly the one obtained using the proposed even-odd ternary decomposition tree, reinforcing the impression given by the PSNR values. The mirror transform is particularly well suited to represent local symmetries, in particular edges and gradients, thus giving a better visual result. Of course, this does not mean that the DCT and DWT performance are worse when used for natural images compression, since it is well known that for example when the 2-D DCT is employed, its strong decorrelation properties allow for a very compact representation of visual data.

Last, the results for the other considered 1-D data types are given in Fig. 5. For these data types as well, the comments already offered remain valid: the signals having the greatest amount of smooth, local symmetries have the best chances for being sparsified by the proposed mirror transform, which is especially true for ECG signals.

4. CONCLUSIONS

We presented the mirror transform, that is a new signal transform based on the iterative application of the even/odd decomposition, valid for both continuous and discrete time domains. In the former case, the ternary decomposition tree constituting the transform is in principle infinitely deep even for finite support signals, while for finite-length sequences the decomposition tree always has a finite number of levels. Furthermore, we stated how using the proposed transform a continuous-time signal can be equivalently represented by an exponentially increasing number of symmetry positions, highlighting an interesting connection between sparsity and symmetry.

We employed the proposed mirror transform in a variety of experiments, aimed at analyzing its sparsity properties, limiting our scope to finite trees obtained from 1-D signals of various origin. Of course, the compactness of the alternative representation is hampered by the amount of information concerning the symmetry point employed in each decomposition step. Nevertheless, the over-completeness of the signals set, implicit in the aforementioned ability to translate a signal into a set of positions, allows for sparse representations for signals with certain characteristics.



Fig. 5: Mean reconstruction SNR for the other 1-D data types.

5. REFERENCES

- V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," ACM Computing Surveys (CSUR), vol. 41, no. 3, pp. 15, 2009.
- [2] P. Over, G. Awad, M. Michel, J. Fiscus, G. Sanders, W. Kraaij, A. Smeaton, and G. Quenot, "Trecvid 2014 – an overview of the goals, tasks, data, evaluation mechanisms and metrics," in *Proc. of TRECVID 2014*. NIST, USA, 2014.
- [3] E. Hjelmås and B. K. Low, "Face detection: A survey," Computer vision and image understanding, vol. 83, no. 3, pp. 236– 274, 2001.
- [4] J. Fourier, *Theorie analytique de la chaleur, par M. Fourier*, Chez Firmin Didot, père et fils, 1822.
- [5] A. Haar, "Zur theorie der orthogonalen funktionensysteme," Mathematische Annalen, vol. 69, no. 3, pp. 331–371, 1910.
- [6] S. Mallat, "A theory for multiresolution signal decomposition: the wavelet representation.," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, no. 7, pp. 674–693, 1989.
- [7] M. Belkin and P. Niyogi, "Laplacian eigenmaps for dimensionality reduction and data representation," *Neural computation*, vol. 15, no. 6, pp. 1373–1396, 2003.
- [8] G. Peyré, "Manifold models for signals and images," Computer Vision and Image Understanding, vol. 113, no. 2, pp. 249–260, 2009.
- [9] A. Jacquin, "Image coding based on a fractal theory of iterated contractive image transformations," *IEEE Transactions* on *Image Processing*, vol. 1, no. 1, pp. 18–30, 1992.
- [10] A. Gnutti, F. Guerrini, and R. Leonardi, "Representation of signals by local symmetry decomposition," in *Proc. of the 23rd European Signal Processing Conference (EUSIPCO)*, 2015, pp. 983–987.
- [11] PhysioToolkit PhysioBank, "Physionet: components of a new research resource for complex physiologic signals," *Circulation.* v101 i23. e215-e220, 2000.
- [12] "Incorporated research institutions for seismology data," https://www.iris.edu/hq/resource/ bb-processing-matlab, Accessed: 2017-04-27.
- [13] "The usc-sipi image database," http://sipi.usc.edu/ database/, Accessed: 2017-10-01.