

Evolutionary signal enhancement based on Hölder regularity analysis

Jacques LÉVY VÉHEL and Evelyne LUTTON

Projet Fractales — INRIA, B.P. 105, 78153 Le Chesnay cedex, France,
Jacques.Levy_Vehel@inria.fr, Evelyne.Lutton@inria.fr,
<http://www-rocq.inria.fr/fractales>

Abstract. We present an approach for signal enhancement based on the analysis of the local Hölder regularity. The method does not make explicit assumptions on the type of noise or on the global smoothness of the original data, but rather supposes that signal enhancement is equivalent to increasing the Hölder regularity at each point. The problem of finding a signal with prescribed regularity that is as near as possible to the original signal does not admit a closed form solution in general. Attempts have been done previously on an analytical basis for simplified cases [11]. We address here the general problem with the help of an evolutionary algorithm. Our method is well adapted to the case where the signal to be recovered is itself very irregular, e.g. nowhere differentiable with rapidly varying local regularity. In particular, we show an application to SAR image denoising where this technique yields good results compared to other algorithms. The implementation of the evolutionary algorithm has been done using the EASEA (EASy specification of Evolutionary Algorithms) language.

1 Introduction

A large number of techniques have been proposed for signal enhancement. The basic frame is as follows. One observes a signal Y which is some combination $F(X, B)$ of the signal of interest X and a “noise” B . Making various assumptions on the noise, the structure of X and the function F , one then tries to derive a method to obtain an estimate \hat{X} of the original signal which is optimal in some sense. Most commonly, B is assumed to be independent of X , and, in the simplest case, is taken to be white, Gaussian and centered. F usually amounts to convoluting X with a low pass filter and adding the noise. Assumptions on X are almost always related to its regularity, e.g. X is supposed to be piecewise C^n for some $n \geq 1$. Techniques proposed in this setting resort to two domains: functional analysis and statistical theory. In particular, wavelet based approaches, developed in the last ten years, may be considered from both points of view [1, 5].

Our approach in this work is different from previous ones in several respects. First, we do not make explicit assumptions on the type of noise and the coupling between X and B through F . However, if some information of this type

is available, it can readily be used in our method. Second, we do not require that X belong to a given global smoothness class but rather concentrate on its *local* regularity. More precisely, we view enhancement as equivalent to increasing the Hölder function α_Y (see next section for definitions) of the observations. Indeed, it is generally true that the local regularity of the noisy observations is smaller than the one of the original signal, so that in any case, $\alpha_{\hat{X}}$ should be greater than α_Y . We thus define our estimate \hat{X} to be the signal “closest” to the observations which has the desired Hölder function. Note that since the Hölder exponent is a local notion, this procedure is naturally adapted for signals which have sudden changes in regularity, like discontinuities. From a broader perspective, such a scheme is appropriate when one tries to recover signals which are highly irregular and for which it is important that the restoration procedure yields the right regularity structure (i.e. preserves the evolution of α_X along the path). An example of this situation is when denoising is to be followed by image segmentation based on textural information: Suppose we wish to differentiate highly textured zones (appearing for instance in MR or radar imaging) in a noisy image. Applying an enhancement technique which assumes that the original signal is, say, piecewise C^2 , will induce a loss of the information which is precisely the one needed for segmentation, since the denoised image will not contain much texture. The same difficulty occurs in other situations such as change detection from noisy sequences of aerial images or automatic monitoring of the evolution of lung diseases from scintigraphic images: in such cases, the criterion for a change is often based on a variation of the irregularity in certain regions, and one needs to preserve this information.

In addition to the examples given above, other situations where such conditions occur are turbulence data analysis and characterization of non-voiced parts of voice signals. Since the method is highly non linear and quite complex, an analytic solution is not possible. We thus resort to a stochastic optimisation method based on evolutionary algorithms.

The remaining of this paper is organized as follows. Section 2 recalls some basic facts about Hölder regularity analysis, which is the basis of our approach. The denoising method is explained in section 3. The evolutionary implementation based on EASEA is then detailed in section 4, finally numerical results on both 1D and 2D signals are displayed in section 5.

2 Hölder Regularity Analysis

A popular way to measure the local irregularity of signals is to consider Hölder spaces. We will focus in this paper on the Hölder characterizations of regularity. To simplify notations, we assume that our signals are nowhere differentiable. Generalisation to other signals simply requires to introduce polynomials in the definitions [7].

Let $\alpha \in (0, 1)$, $\Omega \subset \mathbf{R}$. One says that $f \in C_l^\alpha(\Omega)$ if:

$$\exists C : \forall x, y \in \Omega : \frac{|f(x) - f(y)|}{|x - y|^\alpha} \leq C$$

Let: $\alpha_l(f, x_0, \rho) = \sup \{ \alpha : f \in C_l^\alpha(B(x_0, \rho)) \}$ $\alpha_l(f, x_0, \rho)$ is non increasing as a function of ρ .

We are now in position to give the definition of the local Hölder exponent :

Definition 1 *Let f be a continuous function. The local Hölder exponent of f at x_0 is the real number: $\alpha_l(f, x_0) = \lim_{\rho \rightarrow 0} \alpha_l(f, x_0, \rho)$*

Since α_l is defined at each point, we may associate to f the function $x \rightarrow \alpha_l(x)$ which measures the evolution of its regularity.

This regularity characterization is widely used in fractal analysis because it has direct interpretations both mathematically and in applications. It has been shown for instance that α_l indeed corresponds to the auditive perception of smoothness for voice signals. Similarly, simply computing the Hölder exponent at each point of an image already gives a good idea of its structure, as for instance its edges [10]. More generally, in many applications, it is desirable to model, synthesize or process signals which are highly irregular, and for which the relevant information lies in the singularities more than in the amplitude. In such cases, the study of the Hölder functions is of obvious interest.

In [12], a theoretical approach for signal denoising based on the use of a Hölder exponent and the associated multifractal spectrum was investigated. We develop here another enhancement technique that uses the information brought by the local Hölder function, which is simple from an algorithmic point of view, and yields good results on several kind of data.

3 Signal Enhancement

We adopt in this paper a functional analysis point of view. This means that we do not make any assumption about the noise structure, nor the way it interacts with the data. Rather, we seek a regularized version of the observed data that fulfills some constraints. A statistical approach, classically based on risk minimization, will be presented elsewhere [11].

Let X denote the original signal and Y the degraded observations. We seek a regularized version \hat{X} of Y that meets the following constraints: a) \hat{X} is close to Y in the L^2 sense, b) the (local) Hölder function of \hat{X} is prescribed.

If α_X is known, we choose $\alpha_{\hat{X}} = \alpha_X$. In some situations, α_X is not known but can be estimated from Y . Otherwise, we just set $\alpha_{\hat{X}} = \alpha_Y + \delta$, where δ is a user-defined positive function, so that the regularity of \hat{X} will be everywhere larger than the one of the observations. We must solve two problems in order to obtain \hat{X} . First, we need a procedure that estimates the local Hölder function of a signal from discrete observations. Second, we need to be able to manipulate the data so as to impose a specific regularity. A third difficulty arises from the following analysis: Assume the simplest case of an L^2 signal corrupted by independent white Gaussian noise. It is easy to check that almost surely $\alpha_Y = -\frac{1}{2}$ everywhere, because a) $-\frac{1}{2}$ is the regularity of the noise, b) $\alpha_X \geq 0$ since $X \in L^2$, c) the regularity of the sum of two signals which have everywhere different Hölder exponents is the min of the two regularities. Thus α_Y does not depend on X ,

and one cannot go back from α_Y to α_X . This fact casts doubts on the efficiency on the whole approach, since the information it is based on is degenerate in this case.

All these problems are solved once one realizes that the mathematical notion of Hölder regularity is an abstraction that makes sense only asymptotically. One needs to analyze carefully how it should be adapted to a finite setting, much in the same way as what is done for abstract white noise. In particular, we are interested in a perceptual notion of regularity: If two 2D functions A and B are such that $\alpha_A < \alpha_B$, but an imaging of A and B at a given resolution yields the contrary visual impression that A is smoother than B , then of course our algorithm should go with the perceptual information. In other words, in practical applications we are not interested in the asymptotic behavior, but in the scales which are really in the image, and our estimate of α should reflect this fact. This means precisely one thing: the estimation procedure should yield results in agreement with what is perceived, and not care for the “true” α , which may or may not be accessible from the finite data. To go back to the example above, while it is true that at infinite resolution the sum “signal + white noise” would look much the same as white noise as far as regularity is concerned, this is not the case at finite resolution, where the influence of the signal is still perceptible. In addition, since our procedure is differential, i.e. we wish to impose $\alpha_{\hat{X}} = \alpha_Y + \delta$ or $\alpha_{\hat{X}} - \alpha_X = 0$, for *estimated* α_X and α_Y , we do not care about constant bias.

We will use a wavelet based procedure for estimating and controlling the Hölder function. This is made possible by results in [7] and [6] which imply that:

Proposition 1 *Let $\{\psi_{j,k}\}_{j,k}$ be an orthonormal wavelet basis, where as usual j denotes scale and k position, and assumes that ψ is regular enough and has sufficiently many vanishing moments. Then, X has local Hölder exponent α at t if and only if for all (j,k) such that t belongs to the support of $\psi_{j,k}$,*

$$|c_{j,k}| \leq C 2^{-j(\alpha + \frac{1}{2})} \quad (1)$$

where C is a constant and $c_{j,k}$ is the wavelet coefficient of X .

Although (1) is only an inequality, it suggests that one may estimate $\alpha_X(t)$ by linear regression of $\log(|c_{j,k}|)$ w.r.t. to the scale j (\log denotes base 2 logarithm) considering those indices (j,k) such that the support of $\psi_{j,k}$ is centered above t . Of course this will be only approximate, but since (1) is a necessary and sufficient condition, and if enough wavelet coefficients are “large”, we may hope to obtain results sufficient for our purpose.

Two points are essential in this estimation procedure:

- The estimation is obtained through a regression on a finite number of scales, defined as a subset of the scales available on the discrete data. This avoids the pathologies described above concerning the regularity of the sum of two signals. In particular, it is possible to express the Hölder function of the noisy signal $Y = X + \text{Gaussian white noise}$ as a function of α_X , and thus to estimate conversely α_X from α_Y [11].

- The use of (orthonormal) wavelets allows to perform the reconstruction in a simple way: Starting from the coefficient $(d_{j,k})$ of the observations, we shall define a procedure that modifies them to obtain coefficients $(c_{j,k})$ that verify (1) with the desired α , and then reconstruct \hat{X} from the $(c_{j,k})$.

We may now reformulate our program as follows: For a given set of observations $Y = (Y_1, \dots, Y_{2^n})$ and a target Hölder function α , find \hat{X} such that $\|\hat{X} - Y\|_{L^2}$ is minimum and the regression of the logarithm of the wavelet coefficients of \hat{X} above any point i w.r.t. scale is $-(\alpha(i) + \frac{1}{2})$. Note that we must adjust the wavelet coefficients in a global way. Indeed, each coefficient at scale j subsumes information about roughly 2^{n-j} points. Thus we cannot consider each point i sequentially and modify the wavelet coefficients above it to obtain the right regularity, because point $i + 1$, which shares many coefficients with i , requires different modifications. The right way to control the regularity is to write the regression constraints simultaneously for all points. This yields a system which is linear in the logarithm of the coefficients:

$$\Delta L = A$$

where Δ is a $(2^n, 2^{n+1} - 1)$ matrix of rank 2^n , and

$$L = (\log |c_{1,1}|, \log |c_{2,1}|, \log |c_{2,2}|, \dots, \log |c_{n,2^n}|),$$

$$A = -\frac{n(n-1)(n+1)}{12} \left(\alpha(1) + \frac{1}{2}, \dots, \alpha(2^n) + \frac{1}{2} \right)$$

Since we use an orthonormal wavelet basis, the requirements on the $(c_{j,k})$ may finally be written as:

$$\begin{aligned} &\text{minimize } \sum_{j,k} (d_{j,k} - c_{j,k})^2 \text{ subject to: } \forall i = 1, \dots, 2^n, \\ &\sum_{j=1}^n s_j \log(|c_{j,E((i-1)2^{j+1-n})}|) = -M_n \left(\alpha(i) + \frac{1}{2} \right) \end{aligned} \quad (2)$$

where $E(x)$ denotes the integer part of x and the coefficients $s_j = j - \frac{n+1}{2}$, $M_n = \frac{n(n-1)(n+1)}{12}$ and equation (2) are deduced from the requirement that the linear regression of the wavelet coefficients of \hat{X} above position i should equal $-(\alpha(i) + \frac{1}{2})$.

Finding the global solution to the above program is a difficult task; in particular, it is not possible to find a closed form formula for the $c_{j,k}$. In the following, we show how this problem can be addressed with an evolutionary algorithm.

4 Evolutionary signal enhancement with EASEA

An evolutionary technique seems to be appropriate for the optimisation problem described in equation (2): a large number of variables are involved, and the function to be optimised as well as the constraint are non linear. We describe in this section an implementation based on the EASEA [13] language and compiler.

EASEA (EAsy Specification of Evolutionary Algorithms) is a language dedicated to evolutionary algorithms. Its aim is to relieve the programmer of the task of learning how to use evolutionary libraries and object-oriented programming by using the contents of a `.ez` source file written by the user.

EASEA source files only need to contain the "interesting" parts of an evolutionary language, namely the fitness function, a specification of the crossover, the mutation, the initialisation of a genome plus a set of parameters describing the run. With this information, the EASEA compiler creates a complete C++ source file containing function calls to an evolutionary algorithms library (either the GALIB or EO for EASEA v0.6). Therefore, the minimum requirement necessary to write evolutionary algorithms is the capability of creating non-object-oriented functions, specific to the problem which needs to be solved.

In our case, the evolutionary optimisation involved to enhance a signal (1D or 2D) was implemented using a simple structure on which genetic operators were defined. We used GALib [4] as the underlying evolutionary library.

We describe below the implementation for 1D signals, an implementation for image denoising was also produced based on the same principle, and results for 1D and 2D data are presented in the next section.

The Haar wavelet transform has been used to produce the $d_{j,k}$ associated to the observed signal Y . We also suppose that we know the desired Hölder exponents $\alpha(i)$ (either $\alpha(i) = \alpha_Y(i) + \delta$ where the $\alpha_Y(i)$ are the Hölder exponents of Y and δ is a user defined regularisation factor, or $\alpha(i)$ is set a priori).

The problem is to find some multiplication factor $u_{j,k}$ such that $c_{j,k} = u_{j,k} * d_{j,k}$, $j \in [0..n-1]$, $k \in [0..2^j-1]$. As is usual in wavelet denoising, we let unchanged the first l levels and seek for the remaining $u_{j,k}$ in $[0,1]$. The genome is made of the $u_{j,k}$ coefficients, for $j \in [l..n-1]$ and $k \in [0..2^j-1]$. These coefficients are encoded as a real numbers vector of size `SIZE_MAX = $2^n - 2^l$` , which can be written using EASEA syntax as:

```
GenomeClass { double U[SIZE_MAX]; }
```

The EASEA Standard functions sections contain the specific genetic operators, namely:

1. **The initialisation function:** Each $u_{j,k}$ coefficient is randomly set to a value in $[0,1]$. Two initial solutions are also put in the initial population: $u_{j,k} = 1$. and $u_{j,k} = 2^{-k\delta}$.
2. **The crossover function:** a barycentric crossover has been easily defined as follows : Let `parent1` and `parent2` be the two genomes out of which `child1` and `child2` must be generated, and let `alpha` be a random factor:

```
\GenomeClass::crossover:
  double alpha = (double)random(0.,1.);
  if (&child1) for (int i=0; i<SIZE_MAX; i++)
    child1.U[i] = alpha*parent1.U[i] + (1.-alpha)*parent2.U[i];
  if (&child2) for (int i=0; i<SIZE_MAX; i++)
    child2.U[i] = alpha*parent2.U[i] + (1.-alpha)*parent1.U[i];
\end
```

3. **The mutation function:** Mutation is a random perturbation of radius $SIGMA = 0.01$, applied with probability $PMut$ on each gene.

```

\GenomeClass::mutator: // Must return the number of mutations as an int
  int NbMut=0;
  for (int i=0; i<SIZE_MAX; i++)
  if (tossCoin(PMut)){ NbMut++;
Genome.U[i]+=SIGMA*(double)random(-1.,1.);
Genome.U[i] = MIN(1.,Genome.U[i]); Genome.U[i] = MAX(0.,Genome.U[i]);}
  if (NbMut==0) identicalGenome=true; // saves evaluation time
  return NbMut;
\end

```

4. **The evaluation function:** The fitness function has two aims: minimise $\sum((1 - u_{j,k}) * c_{j,k})^2$, making sure constraint (2) is satisfied, i.e. the Hölder exponents are the ones we want. The constraint is integrated to the fitness function using a high penalisation factor W :


$$Fitness = \sum_{j,k} ((1 - u_{j,k}) * c_{j,k})^2 + W * \sum_i |\alpha_u(i) - \alpha(i)|$$

We use the GALib steady state genetic engine with replacement percentage of 60% and a selection by ranking. Crossover and mutation probabilities are fixed respectively to 0.9 and 0.001. Genome size, Population size, and number of generations are fixed for each experiment, see section 5.


5 Numerical Experiments

We first show results of enhancement on synthetic 1D data. The original signal is a generalised Weierstrass function [2] with $\alpha_X(t) = t$ for all t , (i.e. $X(t) = \sum_{n=0}^{\infty} 2^{nt} \sin(2^{nt})$, $t \in [0.1]$) which has been corrupted by white Gaussian noise. Figure 1 shows the original signal, the noisy one, and the result of the enhancement procedure. For comparison, a denoising using a classical wavelet shrinkage is also displayed. For both procedures, the parameters were set so as to obtain the best fit to the known original signal. It is seen that, for such irregular signals, the Hölder regularity based enhancement yields more satisfactory results. The constraint was to find Hölder exponents that verify $\alpha(t) = t$ for all t . Parameters of the evolutionary algorithm were as follows:


Genome Size	SIZE_MAX = 496
Population Size	25
Number of generations	50000
Computation time	1365.96 seconds for 293506 evaluations



Generalised Weierstrass function (left) + noise (right).



Denoising using wavelet shrinkage method (left), and using the evolutionary scheme with prescribed $\alpha(t) = t$ (right).




Left: Estimated Hölder exponents of the original function (left), of the function + noise (middle) and of the reconstructed function (right).

Fig. 1. Results on a generalised Weierstrass function $\alpha(t) = t$

Our second example deals with a synthetic aperture radar (SAR) image. A huge literature has been devoted to the difficult problem of enhancing these images, where the noise is non Gaussian, correlated and multiplicative. A fine analysis of the physics of the speckle suggests that it follows a K distribution [9]. Classical techniques specifically designed for SAR image denoising include geometric filtering and Kuan filtering. Wavelet shrinkage methods have also been adapted to this case [3]. Figure 2 show an original SAR image, its denoising with the Hölder method and with soft thresholding. Notice how the river, which appears with a “A” shape in the middle of the image is nicely uncovered by the regularity based enhancement. As no a priori knowledge about Hölder exponents of the signal was available, the constraint was to find Hölder exponents that verify $\alpha_{denoised}(t) = \alpha_{original}(t) + \delta$ for all t . Parameters of the evolutionary algorithm were as follows :

Genome Size	SIZE_MAX = 21845
Population Size	50
Number of generations	100
Computation time	1702.46 seconds for 3051 evaluations
Regularisation factor δ	0.5 and 0.7



Left: Original SAR image. Right: Image denoised using soft thresholding.




Image denoised using the Hölder regularity scheme. Left: $\delta = 0.5$. Right: $\delta = 0.7$

Fig. 2. Experiments on a 256x256 area of a SAR image (courtesy IRD)

6 Conclusion

We have shown in this paper how an evolutionary algorithm can be applied to a signal or image enhancement technique based on a “fractal” analysis. Good results have been obtained as well as on 1D or 2D signal in comparison to other techniques in an affordable computation time.

Quality of the results is also very dependent from the quality of Hölder exponents estimation. Much better estimate of $\alpha_X(t)$ can be obtained by measuring the *oscillations* of X in balls centered at t and of radii ϵ_k , and then regressing the logarithm of these oscillations w.r.t. the logarithm of the ϵ_k . However, this procedure does lead to a more complex inverse problem, i.e. obtaining a signal with prescribed regularity. Future work will be devoted to an evolutionary formulation of this problem.

References

1. R.A. Devore, B. Lucier *Fast wavelet techniques for near optimal image processing*. 1992 IEEE Military Communications Conference, 2–12 (1992).
2. K. Daoudi, J. Lévy Véhel, Y. Meyer, *Construction of functions with prescribed local regularity*, Constructive Approximation, 1989.
3. L. Gagnon, F. Drissi Smaili, *Speckle noise reduction of airborne SAR images with symmetric Daubechies wavelets*, Signal and Data Processing of Small Targets, Proc. SPIE 2759, 1996.
4. M. Wall, *GAlib home page: <http://lancet.mit.edu/ga/>*, MIT.
5. D.L. Donoho, *De-noising by soft-thresholding*. IEEE Trans. Inf. Theory 41, No. 3, 613–627 (1994).
6. B. Guiheneuf, J. Lévy Véhel, *2 micro-local analysis and applications in signal processing*, Int. Conf. on Wavelet, Tangier, 1997.
7. S. Jaffard, *Pointwise smoothness, two-microlocalization and wavelet coefficients*. Publ. Mat. 35, No. 1, 155–168 (1991).
8. J. J. Merelo, *EO home page: <http://geneura.ugr.es/~merelo/EO.html>*, Granada University.
9. C.J. Oliver, *Information from SAR images*, J. Phys. D, **24**, 1493-15144, 1991.
10. J. Lévy Véhel, *Fractal Approaches in Signal Processing*, Fractals, **3** (4), pp 755-775, 1995.
11. J. Lévy Véhel, *Statistical denoising of irregular signals*, INRIA internal report.
12. J. Lévy Véhel, B. Guiheneuf “Multifractal Image Denoising,” *SCIA*, 1997.
13. Pierre Collet, Evelyne Lutton, Marc Schoenauer, Jean Louchet, “Take it EASEA,” Parallel Problem Solving from Nature VI, vol 1917, Springer Verlag pp 891-901, Paris, septembre 2000. *EASEA home page: <http://www-rocq.inria.fr/EASEA/>*