PIECEWISE PARAMETERISED MARKOV RANDOM FIELDS FOR SEMI-LOCAL HURST ESTIMATION

J.-B. Regli and J. D. B. Nelson

Department of Statistical Science, University College London

ABSTRACT

Semi-local Hurst estimation is considered by incorporating a Markov random field model to constrain a wavelet-based pointwise Hurst estimator. This results in an estimator which is able to exploit the spatial regularities of a piecewise parametric varying Hurst parameter. The pointwise estimates are jointly inferred along with the parametric form of the underlying Hurst function which characterises how the Hurst parameter varies deterministically over the spatial support of the data. Unlike recent Hurst regularistion methods, the proposed approach is flexible in that arbitrary parametric forms can be considered and is extensible in as much as the associated gradient descent algorithm can accommodate a broad class of distributional assumptions without any significant modifications. The potential benefits of the approach are illustrated with simulations of various first-order polynomial forms.

1. INTRODUCTION

The Hurst parameter determines the spectral decay rate of a process with a power-law spectrum. Since such a simple relationship is ubiquitous in many signal and image processing areas and beyond [1,2] Hurst estimation continues to enjoy many, and disparate, applications including Finance [3], signal/image denoising [4], clutter suppression [5], segmentation [6], the analysis of ECG signals [7, 8], internet traffic flow [1], image texture [9], and turbulence data [10].

The interconnection between wavelets and self-similar processes is a powerful, if not, surprising one. The self-similarity explicitly built in to the wavelet basis functions via the two-scale, or refinement, relations provides a natural representation in which to study processes that exhibit power-law behaviour. However, the localised nature of wavelets also facilitates a localised estimation of the Hurst parameter.

Although there are works, such as those based on the multifractal formalism [11, 12], that describe how regularity varies across an image, less attention has been paid to the case where the main interest is to obtain pointwise estimates of a Hurst parameter that is allowed to vary as a smooth, deterministic function. Such a scenario could, for example present itself in image processing when the texture of an object of interest varies gradually over its spatial support in some assumed manner. In turn this would facilitate tasks such as feature extraction, segmentation, and change detection. Likewise, existing adaptive deniosing methods, which are currently based on a piecewise constant Hurst parameter [13], could also be extended to include more general Hurst functions that vary as piecewise parametric functions.

Since it is reasonable to assume that an image of interest may comprise multiple textures, it is appropriate to consider a piecewise smoothly varying Hurst parameter $H = H(\mathbf{r})$, for $\mathbf{r}$ over some subregion of $\mathbb{R}^2$. Furthermore, we let the way in which this Hurst function varies over space be governed by some parametric form $H = \phi(\mathbf{r}; \boldsymbol{\theta})$ with model parameters $\boldsymbol{\theta}$. We would expect these parameters to be fairly constant over certain subregions of the image domain where the image texture is homogeneous. We allow the spatial support to accommodate multiple textures with a suitable partitioning of disjoint subregions. In each subregion, the $\boldsymbol{\theta}$ are assumed constant (or have very small, smooth variations). However, between subregion boundaries, it is allowed to change arbitrarily. As a consequence the Hurst parameter itself will vary smoothly inside a partition and vary arbitrarily across the respective subregions. We here propose a model and inference scheme that exploits this piecewise parametric outlook. The framework utilises a Markov random field prior to constrain, or penalise, the magnitude of parameter variation over the image.

Spatial regularisation of Hurst estimation has been recently considered as a means to exploit prior knowledge about the spatial smoothness of the Hurst parameter [13]. However, the method was based on the generalised lasso and assumed only a piecewise constant varying Hurst parameter. In contrast our model, and corresponding gradient-descent-like algorithm, are more flexible. The framework can accommodate many different kinds of distributional assumptions and arbitrary models that describe how the Hurst parameter varies deterministically in space. On the other hand, the generalised lasso Hurst estimator simply penalises the ℓ_1-norm of the Hurst parameter spatial derivatives (of some specified order). Therefore, along with a fixed Gaussian assumption on the data, the spatial derivatives of the Hurst parameter are assumed to be Laplacian and it is difficult to incorporate other distributional assumptions without making wholesale changes to the inference scheme. Other

J.-B. Regli is funded by a Dstl/UCL Impact studentship

J. D. B. Nelson is partially supported by grants from the Dstl and Innovate UK/EPSRC

assumptions would necessitate a change in inference strategy (if one existed). Furthermore, unlike the method proposed here, the lasso inference does not obtain any estimate of the underlying parametric form of the Hurst 'function'.

In Section 2 we present the requisite background of wavelet-based Hurst estimation and Li's piecewise (roof-edge) parameterised Markov random field model [14]. We fuse these two concepts in Section 3, propose our parameterised MRF Hurst estimation framework, and describe the inferential machinery. In Section 4 we perform estimation on a selection of simulated imagery where the Hurst parameter is varied according to several first-order polynomial forms. Each one manifests unique roof-like edges in the Hurst parameter and presents different challenges to the estimators. We draw conclusions in Section 5.

2. BACKGROUND

The Hurst parameter controls the spectral slope of a self-similar stochastic process which obeys a power-law relationship. Myriad estimation approaches exist [2]. We here follow the popular wavelet-based framework [1].

2.1. Wavelet-based Hurst estimation

Consider a stochastic field z defined on a subregion of $\mathbb{R}^2$ with weak statistical self-similarity namely $\mathbb{E}z(\lambda\cdot) = \lambda^H \mathbb{E}z$ and $\mathbb{E}z(\lambda\mathbf{r})z(\lambda\cdot) = \lambda^{2H}\mathbb{E}z(\mathbf{r})z(\cdot)$. Then, it is well known (see e.g. [15]), that

$$\mathbb{E}\big|(\mathcal{W}z)(\cdot\,;k,\alpha)\big|^2 \propto 2^{2k(H+1)} \tag{1}$$

where $\mathcal{W}$ is the wavelet transform operator defined by $(\mathcal{W}z)(\mathbf{r};k,\alpha) := 2^{-k}\big\langle z, \psi_\alpha(2^{-k}\cdot - \mathbf{r})\big\rangle$, with wavelet ψ defined over space $\mathbf{r}$, orientation α, and kth finest scale level.

In practice the expectation in Equation (1) is approximated by the sample second moment of the wavelet coefficients magnitudes. When the Hurst parameter varies over space it is still possible to estimate the slope by simply using the squared magnitude of the wavelet coefficients. This pointwise estimate, $E_{k,\alpha}(\cdot) := \big|(\mathcal{W}z)(\cdot\,;k,\alpha)\big|^2$, approximately satisfies the power-law, namely $E_{k,\alpha} \propto 2^{2k(H_\alpha(\mathbf{r})+1)}$. Estimation of H is then performed by taking the log of both sides and regressing the log wavelet magnitude on scale. The Hurst parameter is then easily obtained from the slope of the regression line.

Generally, H can also vary with orientation too. In this case, one can perform separate regressions in each direction as appropriate (cf. [5, 16]). Alternatively, if we assume that the Hurst parameter is isotropic there are two main options. Firstly, one could perform separate regressions over the different orientations and then sum the result. Secondly, one could perform one regression over the orientation -averaged wavelet magnitude. As such, without loss of generality, we can drop any orientation notation and write the log wavelet magnitudes about the spatial location $\mathbf{r}_i$ as $\gamma_k[i]$ where $i \in \mathbb{I}$ simply indexes the spatial locations or 'sites' in Markov random field modelling parlance. This furnishes the set of equations $\boldsymbol{\gamma}[i] = \mathbf{A}\boldsymbol{\beta}[i]$, with

$$\boldsymbol{\gamma}[i] = \begin{bmatrix} \gamma_{k_-}[i] \\ \vdots \\ \gamma_{k_+}[i] \end{bmatrix}, \quad \mathbf{A} = \begin{bmatrix} 1 & k_- \\ \vdots & \vdots \\ 1 & k_+ \end{bmatrix}, \quad \boldsymbol{\beta}[i] = \begin{bmatrix} \beta_1[i] \\ \beta_2[i] \end{bmatrix},$$

where only the k_-th to the k_+th finest wavelet scale levels are used— the coarsest levels will give poor spatial location and the finest levels will typically have low signal-to-noise ratio. Solving in the least-squares sense gives the ordinary least squares (OLS) estimate

$$\hat{\boldsymbol{\beta}}[i] := \operatorname{argmin} \big\|\boldsymbol{\gamma}[i] - \mathbf{A}\boldsymbol{\beta}[i]\big\|_2 = (\mathbf{A}^\top\mathbf{A})^{-1}\mathbf{A}^\top\boldsymbol{\gamma}[i],$$

and then the estimate of the Hurst parameter can be recovered from the second element of the $\boldsymbol{\beta}$ vector, namely $\hat{H}(\mathbf{r}_i) = (\hat{\beta}_2[i]/2 - 1)$.

2.2. Roof edge model

The roof edge model was introduced by Li [14] as a means to recover piecewise planar surfaces from noisy observations. Assuming that the parameters of the underlying true surface are the same, or similar, over contiguous regions of the spatial domain, a Markov random field prior can be introduced to aid inference. This introduces the notion of a Markovian label field $f = \{f_1, \dots f_m\}$ with the property that, conditioned on its neighbours, the field at a site is conditionally independent of all other sites. This allows us to write the full conditional of f as the local conditional: $\mathbb{P}(f_i|f_{-i}) = \mathbb{P}(f_i|f_{\mathbb{I}_i})$.

As a consequence of the Hammersley-Clifford Theorem, the joint prior takes the form $\mathbb{P}(f) \propto \exp(-U(f))$. The prior energy term $U(f)$ therefore determines the manner in which spatially incoherent label configurations are penalised. Given observations d, this is counter-balanced to some extent by the likelihood energy $U(d|f)$. By Bayes rule the posterior $\mathbb{P}(f|d)$ has (posterior) energy $U(f|d) = U(d|f) + U(f)$. Observations are assumed to follow some parametric surface, corrupted by noise $d_i := \phi(\mathbf{r}_i;\boldsymbol{\theta}_i) + \epsilon_i$ but where the underlying labels of the parameters $\boldsymbol{\theta}_i$ satisfy the Markov model. For our problem we exploit this to impose piecewise smooth constraints on the Hurst function model parameters $\boldsymbol{\theta}_i$ and, as a consequence, on the Hurst parameter itself. In Li's basic roof edge model, $\phi(\mathbf{r}_i;\boldsymbol{\theta}_i) := \boldsymbol{\theta}_i^\top\boldsymbol{\rho}_i$, with $\boldsymbol{\theta}_i^\top := (\theta_0[i], \theta_1[i], \theta_2[i])$, and $\boldsymbol{\rho}_i^\top := (1, x_i, y_i)$ but higher-order polynomials can easily be accommodated.

Given data d and the distributional assumptions of ϵ (i.e. the likelihood), and our prior model of the underlying configuration label field (the prior), the goal then is to estimate the maximum a posteriori, namely $f^* = \operatorname{argmin}_f U(f|d)$.

3. PARAMETERISED MRF HURST ESTIMATION

We assume that the Hurst parameter varies as a piecewise parametric function. The parameters which describe how H varies are therefore assumed to change little within a given subregion. However, the parameters may change at the boundaries between subregions. We therefore introduce a Markov random (label) field to assign sites and model parameters to labels.

3.1. Markov random field model

The ordinary least squares estimate $\hat{\boldsymbol{\beta}}[i] = (\mathbf{A}^\top\mathbf{A})^{-1}\mathbf{A}^\top\boldsymbol{\gamma}[i]$ gives rise to a 'noisy' version of the true value of $\boldsymbol{\beta}$, namely $\hat{\boldsymbol{\beta}}[i] = \boldsymbol{\beta}[i] + \boldsymbol{\epsilon}[i]$. For notational convenience, and without generality, denote the observed spectral log-slope (i.e. $\hat{\beta}_2[i]$) as $\hat{\beta}[i]$. Assume that the true spectral slope follows some parametric model: $\beta[i] = \phi(\mathbf{r}_i;\boldsymbol{\theta}_i)$, where $\mathbf{r}_i = (x_i, y_i)$ determines pixel location and where $\boldsymbol{\theta}_i$ denotes the parameters of β. Then, assuming that the noise is iid Gaussian[1] $\epsilon_i \overset{iid}{\sim} \mathcal{N}(0,\sigma^2)$ we have the likelihood energy

$$U(\hat{\beta}|f) = \lambda \sum_{i\in\mathbb{I}} \big(\hat{\beta}[i] - \phi(\mathbf{r}_i;\boldsymbol{\theta}_i)\big)^2 \tag{2}$$

Exploiting the Markov structure of the label field, we use a prior energy function of the same form as Li [14, 17]:

$$U(f) = \sum_{i\in\mathbb{I}} \sum_{i'\in\mathbb{I}_i} g\big(\big\|\mathbf{W}(\boldsymbol{\theta}_i - \boldsymbol{\theta}_{i'})\big\|_2\big), \tag{3}$$

where $\mathbb{I}_i$ is the neighbourhood of site i and the diagonal weight matrix $\mathbf{W}$ provides the option to penalise the lack of smoothness of each parameter to different degrees. Again, following Li [14], we choose $g(z) = \ln(1+z^2)$. The choice of ϕ determines the complexity with with the underlying Hurst parameter is assumed to vary. In contrast to the work of Nafornita et al [13], who considered a piecewise constant Hurst, we here consider a Hurst parameter which varies as a piecewise order-1 polynomial. However, it should be noted that higher-order terms can easily be accommodated by recalling that $\phi(\mathbf{r}_i;\boldsymbol{\theta}_i) = \boldsymbol{\theta}_i^\top\boldsymbol{\rho}_i$ and noting that the vectors $\boldsymbol{\rho}_i$ and $\boldsymbol{\theta}_i$ can be extended accordingly. For example higher order products $(x^{p_0}y^{p_1})_{p_0,p_1}$ can be concatenated on to the vector $\boldsymbol{\rho}_i$ for suitable ranges of p_0 and p_1.

3.2. Inference

Given the least-squares estimate of the Hurst parameter and the Markov random field roof-edge piecewise parametric model, we find the maximum a posteriori solution to the problem, namely

$$U(f|\hat{\beta}) := U(\hat{\beta}|f) + U(f).$$

This is an unconstrained optimization problem and can be solved using a gradient-descent-like algorithm. The derivatives with respect to the model parameters can be expressed analytically as

$$\frac{1}{2}\frac{\partial U(f|\hat{\beta})}{\partial\boldsymbol{\theta}_i} = \lambda\left(\hat{\beta}[i] - \phi(\mathbf{r}_i;\boldsymbol{\theta}_i)\right)\boldsymbol{\rho}_i + \sum_{i'\in\mathbb{I}_i} g'\big(\big\|\mathbf{W}(\boldsymbol{\theta}_i - \boldsymbol{\theta}_{i'})\big\|_2\big)\,\mathbf{W}(\boldsymbol{\theta}_i - \boldsymbol{\theta}_{i'}). \tag{4}$$

In our implementation we use the unconstrained version of the BFGS algorithm proposed by Yuan [18] instead of a simple gradient descent. It is a variation of second order newton's method where the Hessian matrix is estimated rather than computed at every steps. The optimization procedure is detailed in Algorithm 1. Therein, for a given step ℓ, we define $\boldsymbol{\theta}^{(\ell)} := \big(\boldsymbol{\theta}_i^{(\ell)}\big)_{i=1}^m \in \mathbb{R}^{3\times M}$, $\boldsymbol{\rho}^{(\ell)} := \big(\boldsymbol{\rho}_i^{(\ell)}\big)_{i=1}^m \in \mathbb{R}^{3\times M}$, $\mathbf{B}_{(\ell)} := \big(\mathbf{B}_{(\ell)}[i]\big)_{i=1}^m \in \mathbb{R}^{3\times 3\times M}$ and where the products between the elements are defined pixel-wise, namely: $\boldsymbol{\theta}_{(\ell)}^\top\boldsymbol{\rho}^{(\ell)} = \big(\boldsymbol{\theta}^{(\ell)\top}[i]\boldsymbol{\rho}_i^{(\ell)}\big)_{i=1}^m$.

The meta-parameter λ in Equation (2) is used to control the importance of the likelihood over the prior. The weights in the diagonal matrix allows variable emphasis to be placed on each of the model parameters $\boldsymbol{\theta}_i$.

Meta-parameters:
λ, $\mathbf{w}$
Initialization:
$\ell = 0$ and $\mathbf{B}_{(0)}[i] = \mathbf{I}_3\ \forall i \in [\![1, m]\!]$
while *convergence* **do**
- *Descent direction:*
$\mathbf{p}^{(\ell)} = -\mathbf{B}_{(\ell)}^{-1}\nabla U(f|\hat{\beta}^{(\ell)})$
- *Optimal step in the direction* $\mathbf{p}^{(\ell)}$*:*
$\mu^{(\ell)} = \operatorname{argmin}_{\mu\in\mathbb{R}}\left[U(f|(\boldsymbol{\theta}_{(\ell)} + \mu\mathbf{p}^{(\ell)})^\top\boldsymbol{\rho})\right]$
$\boldsymbol{\theta}^{(\ell+1)} = \boldsymbol{\theta}^{(\ell)} + \mu^{(\ell)}\mathbf{p}^{(\ell)}$
$\hat{\beta}^{(\ell+1)} = \boldsymbol{\theta}^{(\ell)^\top}\boldsymbol{\rho}$
- *Hessian matrix estimate:*
$\boldsymbol{\eta}^{(\ell)} = \nabla U(f|\hat{\boldsymbol{\beta}}^{(\ell+1)}) - \nabla U(f|\hat{\boldsymbol{\beta}}^{(\ell)})$
$\mathbf{B}_{(\ell+1)} = \mathbf{B}_{(\ell)} + \dfrac{\boldsymbol{\eta}^{(\ell)}\boldsymbol{\eta}^{(\ell)\top}}{\mu_{(\ell)}\boldsymbol{\eta}^{(\ell)\top}\mathbf{p}^{(\ell)}} - \dfrac{\mathbf{B}_{(\ell)}\mathbf{p}^{(\ell)}\mathbf{p}^{(\ell)\top}\mathbf{B}_{(\ell)}}{\mathbf{p}^{(\ell)\top}\mathbf{B}_{(\ell)}\mathbf{p}^{(\ell)}}$
$l = l + 1$
end

Algorithm 1: Minimization of the posterior energy

4. EXPERIMENTS

Experiments were carried out to test the utility of the proposed method for scenarios where the Hurst parameter varied as a first-order polynomial surface. In particular, the behaviour of the estimator was investigated when H varied as a selection of different roof-edge-like functions. These might model the way in which a texture becomes gradually smoother or rougher in space. The second column of Figure 1 illustrates the different roof-edge shapes. For simplicity, we let H vary as a function of it's ℓ_∞-norm distance from the centre of the image, namely $H(\mathbf{r}) = h(\|\mathbf{r}\|)$. The function h is a projection of the Hurst values onto the ℓ_∞-ball; we shall refer to it as the *Hurst signature*. The signatures of the four different roof-types are plotted in the first column of Figure 1.

4.1. Simulation

The data was synthesised by adapting the incremental Fourier synthesis approach of Kaplan and Kuo [19], as implemented in the Fraclab toolbox [20]. We partition the spatial domain into disjoint ℓ_∞ tori: $\mathbb{I}[j] := \{i \in \mathbb{I}\colon \|\mathbf{r}_i\|_\infty \in [j\Delta r, (j+1)\Delta r)\}$. Then, fractional Brownian surfaces are simulated which have a Hurst parameter of $h(j\Delta r)$ on the region $\mathbb{I}[j]$ and

[1] Strictly speaking there exists a small bias term due to non-linearities introduced by the log function [1] but we neglect them here and leave such considerations as further work

Table 1. Mean absolute error (and standard deviation) of the OLS and MRF Hurst estimators

Type	OLS	MRF
1	0.2145 (0.1502)	**0.1335** (0.0955)
2	0.1906 (0.1355)	**0.1330** (0.0983)
3	0.1874 (0.1376)	**0.1286** (0.1038)
4	0.1495 (0.1174)	**0.1117** (0.0928)

which take zero values elsewhere (and which all have the same global white noise driving process). Finally, the surfaces are simply summed over all j. The result is a fractional Brownian surface with a piecewise, order-one polynomial, varying Hurst parameter.

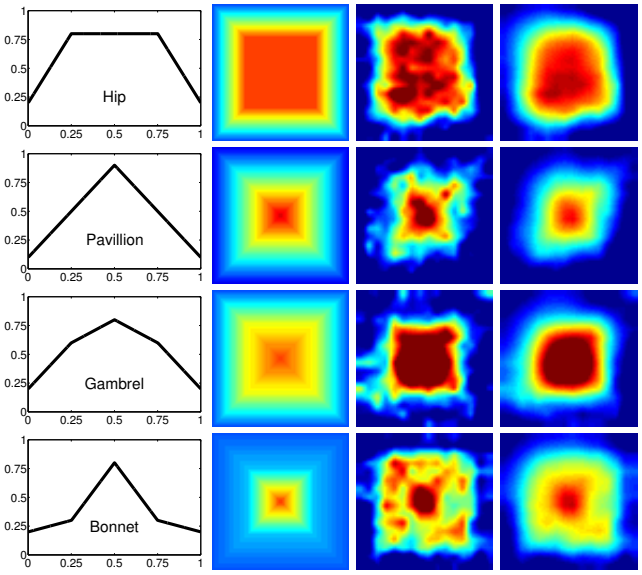

Fig. 1. Indicative Hurst estimates of four fractional Brownian surfaces. 1st column: true Hurst projected onto ℓ_∞ ball; 2nd: spatial map of true Hurst; 3rd: OLS; 4th: MRF.

4.2. Hurst estimation

Hurst estimation was performed on the four image types 'Hip', 'Pavillion', 'Gambrel', and 'Bonnet' illustrated in Fig 1. Ordinary least-squares estimates were used as a baseline for our proposed MRF-based approach although we note that a direct comparison is not necessarily fair as we were free to select an optimal value of λ in our approach to balance the effects of the prior and likelihood functions. Nevertheless, the comparison does offer some intuition as to some of the advantages that one might buy from the addition of an extra parameter. For example, the third and fourth columns of Figure 1 depict the Hurst parameter estimates from the OLS and MRF methods, respectively. The spatial regularisation, or smoothing, effect of the MRF method can be clearly seen for all edge types.

Experiments were performed over 100 instances of each of the edge types. The value of λ was chosen by testing over a smaller subset of data as 0.001 in all cases. The mean absolute errors are listed in Table 1 and the error histograms are plotted in Figure 3. The advantage of exploiting the spatial smoothness

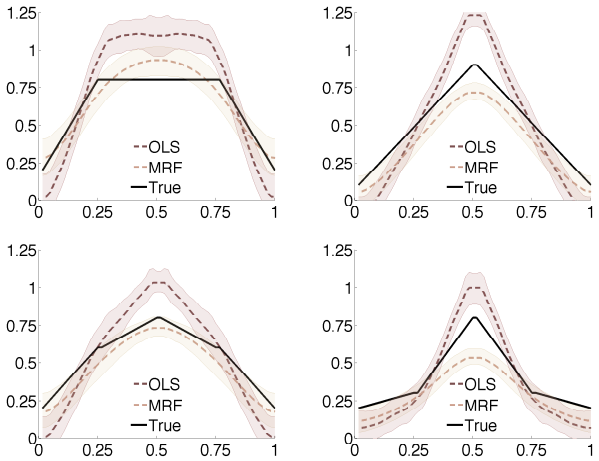

Fig. 2. Mean Hurst estimates of the four fractional Brownian surfaces projected onto the ℓ_∞ ball. Top left: Hip; top right: Pavillion; bottom left: Gambrel; bottom right: Bonnet. The shaded error bars indicate the upper- and lower-quantiles over all experiments and pixel estimates.

of the Hurst parameter is evident. However, this advantage is not as marked in the 'Bonnet'. The reason for this can be seen by inspecting the error as a function of the Hurst signature—i.e. the distance it away from the centre as measured by the ℓ_∞-norm. We see, in Fig. 2, that the MRF method's tendency to smooth the edge features somewhat is more pronounced when the edge is sharp or concave. Nevertheless, MRF still holds an advantage here because the OLS method overshoots the edge point. For convex ridge shapes, the advantage becomes significant.

5. CONCLUSION

A piecewise parameterised Markov random field was introduced to jointly estimate a spatially regularised pointwise Hurst parameter and the model parameters which govern how it varies over the spatial support. The model is flexible in that the model can easily accommodate other likelihood or prior assumptions without any significant changes in the gradient-descent-like inferential machinery. Experiments confirm that the introduction of the Markov random field prior successfully furnishes spatially regularised Hurst estimates with more accuracy than ordinary least squares although this advantage is tempered somewhat when the Hurst function displays concave ridge shapes.

6. REFERENCES

[1] P. Abry, R. Baraniuk, P. Flandrin, R. Riedi, and D. Veitch, "Multiscale nature of network traffic," *IEEE Transactions on Signal Processing Magazine*, vol. 19, pp. 28–46, 2002.

[2] B. Pesquet-Popescu and J. Lévy Véhel, "Stochastic fractal models for image processing," *IEEE Signal Processing Magazine*, vol. 19, no. 5, pp. 48–62, 2002.

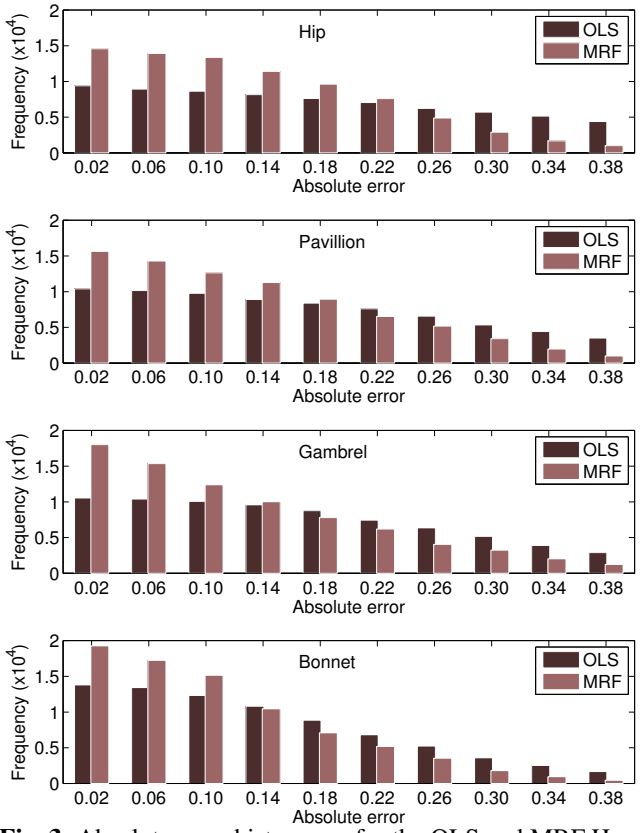

Fig. 3. Absolute error histograms for the OLS and MRF Hurst estimates over 100 instances of each roof-edge type.

[3] L. Kristoufek and M. Vosvrda, "Measuring capital market efficiency: Global and local correlations structure," *Physica A: Statistical Mechanics and its Applications*, vol. 392, pp. 184–193, 2013.

[4] A. Echelard and J. Lévy Véhel, "Wavelet denoising based on local regularity information," *Proceedings of the European Signal Processing Program*, 2008.

[5] J. D. B. Nelson and N. G. Kingsbury, "Dual-tree wavelets for estimation of locally varying and anisotropic fractal dimension," *IEEE International Conference on Image Processing*, pp. 341–344, 2010.

[6] M. Chen and Josef Strobi, "Multispectral textured image segmentation using a multi-resolution fuzzy markov random field model on variable scales in the wavelet domain," *International journal of remote sensing*, vol. 34, pp. 4550–4569, 2013.

[7] V. Chudáček, J. Andén, S. Mallat, P. Abry, and M. Doret, "Scattering transform for intrapartum fetal heart rate variability fractal analysis: A case-control study," *IEEE Transactions on Biomedical Engineering*, pp. 1100–1108, 2014.

[8] M. Julian, R. Alcaraz, and J. Rieta, "Study on the optimal use of generalized hurst exponents for noninvasive estimation of atrial fibrillation organization," *Computing in Cardiology Conference*, pp. 1039–1042, 2013.

[9] S. Deguy, C. Debain, and A. Benassi, "Classification of texture images using multi-scale statistical estimators of fractal parameters," *British Machine Vision Conference*, 2000.

[10] P. Kestener and A. Arneodo, "A multifractal formalism for vector-valued random fields based on wavelet analysis: application to turbulent velocity and vorticity 3d numerical data," *Stochastic Environmental Research and Risk Assessment*, vol. 22, no. 3, pp. 421–435, 2008.

[11] P. Abry, P. Gonçalvès, and J. Lévy Véhel, *Scaling Fractals and wavelets*, Wiley, 2009.

[12] S. Jaffard, B. Lashermes, and P. Abry, "Wavelet leaders in multifractal analysis," in *Wavelet Analysis and Applications*, M. I. Vai T. Qian and X. Yuesheng, Eds., Applied and Numerical Harmonic Analysis, pp. 201–246. Birkhäuser, 2007.

[13] C. Nafornita, A. Isar, and J. D. B. Nelson, "Regularised, semi-local hurst estimation via generalised lasso and dual-tree complex wavelets," *IEEE International Conference on Image Processing*, pp. 2689–2693, 2014.

[14] S. Z. Li, "Roof-edge preserving image smoothing based on mrfs," *IEEE Transactions on Image Processing*, vol. 9, pp. 1134–1138, 2002.

[15] J. D. B. Nelson and N. G. Kingsbury, "Fractal dimension based sand ripple suppression for mine hunting with sidescan sonar," *International Conference on Synthetic Aperture Sonar and Synthetic Aperture Radar*, 2010.

[16] C. Nafornita and A. Isar, "Estimating directional smoothness of images with the aid of the hyperanalytic wavelet packet transform," *International Symposium on Signals, Circuits, and Systems*, 2013.

[17] S. Z. Li, "On discontinuity-adaptive smoothness priors in computer vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, no. 6, pp. 576–586, 1995.

[18] Y. Yuan, "A modified bfgs algorithm for unconstrained optimization," *IMA Journal of Numerical Analysis*, vol. 11, no. 3, pp. 325–332, 1991.

[19] L. M. Kaplan and C. C. J. Kuo, "An improved method for 2-d self-similar image synthesis," *IEEE Transactions on Image Processing*, vol. 5, no. 5, pp. 754–761, 1996.

[20] Inria Saclay/Ecole Centrale de Paris Regularity Team, "Fraclab Toolbox," `http://fraclab.saclay.inria.fr`, 2014.