

LINEAR INTERPOLATION OF BIOMEDICAL IMAGES USING A DATA-ADAPTIVE KERNEL

Hagai Kirshner, Aurélien Bourquard, John Paul Ward, and Michael Unser

Biomedical Imaging Group, EPFL, Lausanne, Switzerland

ABSTRACT

In this work, we propose a continuous-domain stochastic model that can be applied to image data. This model is autoregressive, and accounts for Gaussian-type as well as for non-Gaussian-type innovations. In order to estimate the corresponding parameters from the data, we introduce two possible error criteria; namely, Gaussian maximum-likelihood, and least-squares autocorrelation fit. Exploiting the link between autoregressive models and spline approximation, we use our approach to adapt interpolation parameters to a given image. Our numerical results demonstrate that our adaptive approach yields higher SNR values compared to classical polynomial splines for the task of image scaling. They also indicate that our least-squares-based error criterion nearly achieves the oracle performance for parameter estimation, which provides further support to the practical relevance of our model.

Index Terms— Exponential splines, image interpolation, stochastic modeling.

1. INTRODUCTION

Image interpolation is fundamental to several medical applications such as volume rendering, image rescaling, and image registration [1, 2]. The underlying principle of currently available linear methods is to impose regularity constraints on the continuous-domain solution. In that regard, both theoretical and experimental studies have shown that polynomial B-spline kernels provide a very good trade-off in terms of performance and computational time for general applications [2, 3]. Non-linear methods, on the other hand, tend to be more specific to the type of data in hand. They can rely on local properties of edges, statistical properties of wavelet coefficients, Markov-random-field stochastic models, or variational formulations, to name a few. Compared to the linear case, the main advantage of non-linear methods is their ability to locally adapt the interpolated values to the given data. Their drawback, however, is computational time. Thus, the question that arises is whether one can benefit from both types of approaches, *i.e.*, introduce a linear interpolation method that still adapts to the given data.

This work was funded in part by the ERC Grant ERC-2010-AdG 267439-FUN-SP.

In this work, we propose to interpolate images using a compactly supported generating function that is data-adaptive. We consider a continuous-domain Gaussian autoregressive (AR) model for the image data, assuming that the pixels are sample values of a 2D-separable process. Our approach is motivated by [4] which relies on differential operators corresponding to continuous-domain AR(1) and AR(2) models. In our case, however, we allow for arbitrary model orders and for any admissible set of model parameters. We also distinguish between Gaussian and non-Gaussian innovations, and suggest two possible criteria for adapting the model parameters to the given data. For the Gaussian case, we derive a maximum likelihood cost function and demonstrate its limitations with respect to image-scaling operations. For the non-Gaussian case, we introduce a least-squares fitting criterion and demonstrate its near-optimal properties.

2. THE PROPOSED MODEL

As a starting point, we present the 1D version of the model. We consider the general non-Gaussian continuous-domain AR(p) model with parameters vector θ (Figure 1)

$$v(x) = \int_0^\infty h(x - \xi; \theta) w(\xi) d\xi, \quad (1)$$

where $w(x)$ is a Gaussian or non-Gaussian white-noise process. The equivalent stochastic-integral representation is

$$v(x) = \int_0^\infty h(x - \xi; \theta) dW(\xi), \quad (2)$$

where $W(x)$ is a Brownian motion (in the Gaussian case) or a more general Lévy process (in the non-Gaussian case) [5]. The connection between both representations is

$$W(x) = \int_0^x w(\xi) d\xi. \quad (3)$$

Conversely, $w(x) = \frac{d}{dx}W(x)$, which may be termed white Lévy noise or innovation, is a generalized stochastic process that corresponds to the derivative of $W(x)$ in the sense of distributions [6]. Sparsity is determined by the infinitesimal increments of $W(x)$ (Figure 2). In the case of a Brownian




Fig. 1. A description of a continuous-domain AR process: $w(x)$ is a continuous-domain white noise (Gaussian or non-Gaussian), and $h(x; \theta)$ is an all pole causal filter.




Fig. 2. Examples of sparse and non-sparse stochastic processes: Shown here is $W(x)$ for the Gaussian (non-sparse), the Poisson and the variance-gamma (both sparse) distributions. The probability density function of the increments of $W(x)$ determines the sparsity property of the continuous-domain AR process $v(x)$.

motion, the increments are Gaussian random variables, and $v(x)$ has no sparse innovation. In the Poisson case, however, a large portion of the increments are zero, which results in a sparse innovation [7].

The shaping filter $h(x; \theta)$ is given in the Fourier domain by

$$H(\omega; \theta) = \frac{1}{\prod_{n=1}^p (i\omega - s_n)}, \quad (4)$$

where

$$\theta = (s_1, s_2, \dots, s_p) \in \mathbb{C}^p. \quad (5)$$

The vector of parameters θ consists of the poles of $H(\omega; \theta)$. The real part of each pole is strictly negative, and complex poles appear in conjugate pairs. Assume that $w(x)$ is white with finite variance σ^2 . Then, for sufficiently large values of x , the autocorrelation function $\varphi(\xi) = \mathbf{E}\{v(x) \cdot v(x + \xi)\}$ is given in the Fourier domain by

$$\Phi(\omega; \theta, \sigma^2) = \sigma^2 |H(\omega; \theta)|^2 = \frac{\sigma^2}{\prod_{n=1}^p (i\omega - s_n)(-i\omega - s_n)}. \quad (6)$$

The sampled version of such processes preserves their autocorrelation structure. That is, the autocorrelation sequence of $\{v(n)\}$ is given by the values of the autocorrelation function $\{\varphi(n; \theta)\}$. It then follows that the discrete-domain model is an ARMA (Auto-Regressive Moving Average) process, having independent discrete-domain innovation in the Gaussian case. In the non-Gaussian cases, the innovation is uncorrelated only. The advantage of the proposed model resides in its ability to describe both sparse and non-sparse signals, while maintaining a linear forward model. The parameterization θ also introduces flexibility in determining

the autocorrelation properties of $v(x)$. This is very useful for image interpolation tasks, as the autocorrelation function provides the optimal MMSE linear estimator for unknown pixel values.

3. EXPONENTIAL-BASED IMAGE SCALING

The polynomial B-spline model is known to provide an excellent trade-off between computational complexity and interpolation performance. One of its key aspects is the use of a compactly supported generating function, namely the B-spline, which makes re-sampling very fast. In order to be computationally efficient, we apply a similar approach and represent the autocorrelation function by means of symmetric exponential B-splines. We use this representation and suggest the following interpolation formula

$$\hat{v}(x, y) = \sum_{m,n} c[m, n] \beta(x - m; \theta_x) \beta(y - n; \theta_y). \quad (7)$$

Here, \hat{v} is an approximation for the unknown continuous-domain image, c is a 2D sequence of representation coefficients and $\beta(x, \theta_x)$ is a symmetric exponential B-spline with a vector of parameters θ_x [8]

$$\mathcal{B}(\omega; \theta) = \prod_{k=1}^p \frac{(1 - e^{j\omega + s_k})(1 - e^{-j\omega + s_k})}{(j\omega - s_k)(-j\omega - s_k)}. \quad (8)$$

Similar to the polynomial B-spline approach, the representation coefficients can be extracted from the pixel values by means of digital filtering.

4. FITTING CRITERIA

In the case where $W(x, y)$ is a Brownian motion, the discrete-domain innovation is independent and one can calculate the likelihood function of θ from the sampled data. In the non-Gaussian case, however, the likelihood function is more involved, and we suggest here to resort to a simpler criterion which fits the sample autocorrelation sequence with a continuous-domain AR model. Although simpler, the autocorrelation-based criterion is potentially more powerful as it does not depend on the type of innovation.

4.1. Likelihood of sampled AR Gaussian images

We assume that the digital image is a sampled version of a continuous-domain AR(p) Gaussian process. We then determine the vector of parameters based on the value of the likelihood function. These parameters can then be used in (7). We rely on results from [9] that describe 2D separable AR models, as given next.

The autocorrelation sequence of the sampled model is given in the z-domain by

$$\Phi_d(z; \theta) = \sigma_d^2 H_d(z; \theta) H_d(z^{-1}; \theta) \quad (9)$$




Fig. 3. Parameter estimation performed on a biomedical image. The *B38701* crop is displayed in (a) as an example. The corresponding autocorrelation sequences along the X and Y coordinates are also shown using black \times marks in the plots (b) and (c). The green squares, the red circles, and the blue dots correspond to the best Gaussian-likelihood, the best l_2 -autocorrelation, and the best oracle-SNR parameter fits for our continuous-domain AR model, respectively.

where

$$\begin{aligned}
 H_d(z; \theta) &= \frac{\prod_{k=1}^{p-1} (1 - \nu_k(\theta)z^{-1})}{\prod_{k=1}^p (1 - \rho_k(\theta)z^{-1})}, \\
 \rho_k(\theta) &= e^{s_k}, \\
 \nu_k(\theta) &= \text{roots of } \mathcal{B}_d(z; \theta) \text{ inside the unit circle,} \\
 \mathcal{B}_d(z; \theta) &= \sum_{n=-p}^p \beta[n; \theta]z^{-n}, \\
 \sigma_d^2(\theta) &= \sigma^2 \frac{\mathcal{B}_d(1; \theta)}{\prod_{k=1}^{p-1} (1 - \nu_k(\theta))^2}.
 \end{aligned} \tag{10}$$

We use this continuous-to-discrete mapping for approximating the likelihood function of a digital image. In particular,

$$\lim_{\min(M, N) \rightarrow \infty} \mathbf{E}_{\mathbf{x}} \left| l(\theta_x, \theta_y; \mathbf{x}) - \tilde{l}(\theta_x, \theta_y; \mathbf{x}) \right| = 0, \tag{11}$$

where M, N are the horizontal and vertical dimensions of the image, $\mathbf{x} \in \mathbb{R}^{M \times N}$ are the pixel values, \mathbf{E} is the expectation operation, and $l(\theta; \mathbf{x})$ is the likelihood function of θ for a given image \mathbf{x} . The approximated likelihood function is

$$\begin{aligned}
 \tilde{l}(\theta_x, \theta_y; \mathbf{x}) &= M\kappa(\theta_x) + N\kappa(\theta_y) + M^2 \ln \sigma_d^2(\theta_x) + \\
 &\quad + N^2 \ln \sigma_d^2(\theta_y) + \|\mathbf{x} * \mathbf{g}\|_{\ell_2}^2,
 \end{aligned} \tag{12}$$

where $\mathbf{g}[m, n; \theta_x, \theta_y]$ is given in the z-domain by

$$G_d(z_1, z_2; \theta_x, \theta_y) = \frac{1}{\sigma_d(\theta_x) \sigma_d(\theta_y) H_d(z_1; \theta_x) H_d(z_2; \theta_y)}. \tag{13}$$

Also,

$$\kappa(\theta) = \sum_{n=1}^{\infty} n \cdot c[n; \theta]^2, \tag{14}$$

$$c[n; \theta] = \frac{1}{n} \{ \nu_1^n(\theta) + \dots + \nu_{p-1}^n(\theta) - \rho_1^n(\theta) - \dots - \rho_p^n(\theta) \}. \tag{15}$$

The uniform convergence property (11) implies that the global minimum of $\tilde{l}(\theta_x, \theta_y; \mathbf{x})$ provides a good approximation for the maximum-likelihood solution, while introducing low computational complexity.

4.2. Fitting the autocorrelation sequence

In the non-Gaussian case, the discrete-domain innovation is no longer independent. For this reason, we suggest to focus on the interpolation kernel only, rather than deriving the likelihood function on a case-by-case basis. The interpolation kernel depends on the parameters of the horizontal and vertical shaping filters θ_x and θ_y . Our criterion is based on l_2 fitting of the 2D autocorrelation sequence of the digital image with a continuous-domain AR model. It is an intuitive criterion that has a unique solution [10].

5. RESULTS AND DISCUSSION

We examined the proposed criteria through experiments on simulated AR(2) as well as on real images. In order to quantify the interpolation quality, we downsampled each original image by a factor of 3, followed by upscaling by the same factor using the interpolation methods under consideration. The upscaling operation was repeated for many sets of parameters, using different values of θ_x and θ_y within a large enough search space. When using our adaptive method, we first determined which set of values yielded the best results in terms of SNR, and then—besides this parameter-estimation oracle—evaluated its interpolation performance using both practically applicable error measures (*i.e.*, the Gaussian likelihood function (12) and the l_2 fit of autocorrelation values). The numerical results reported in Table 5 also include the non-adaptive cubic polynomial B-spline model for comparison. Note that the used polynomial and exponential interpolation kernels have the same approximation order.

For the synthetic images generated by the Gaussian AR model, our method gave the same results independently from the error criterion, while the polynomial B-spline model gave worse results, as expected theoretically. Interestingly, however, our results on texture and biomedical images still indicated improvements of our method over the polynomial B-splines by 0.5[dB] on average. While the Gaussian prior is

Table 1. Results for image scaling¹

Image	Fitting criterion ²			Poly. B-spline	
	Best SNR	Best autocorrelation fit	Best Gaussian likelihood		
Gaussian AR	SNR	2.96	2.96	2.96	2.22
	x	$-0.8-i, -0.8+i$	$-0.8-i, -0.8+i$	$-0.8-i, -0.8+i$	
	y	$-0.2-0.6i, -0.2+0.6i$	$-0.2-0.6i, -0.2+0.6i$	$-0.2-0.6i, -0.2+0.6i$	
D2 [640 × 640]	SNR	7.87	7.79	3.97	7.51
	x	$-0.1, -4$	$-0.1, -4$	$-0.6-1.3i, -0.6+1.3i$	
	y	$-0.1, -4$	$-0.1, -1$	$-0.1-0.1i, -0.1+0.1i$	
D79 [640 × 640]	SNR	2.64	2.52	2.12	1.95
	x	$-0.2-i, -0.2+i$	$-0.5-i, -0.5-i$	$-0.1, -1.2$	
	y	$-0.5, -5$	$-0.2, -1.8$	$-0.1, -1$	
B37251 [256 × 256] ³	SNR	6.03	5.74	2.35	5.53
	x	$-0.2, -4$	$-0.2, -0.9$	$-0.1, -0.2$	
	y	$-0.2, -4$	$-0.2, -1$	$-0.9-1.5i, -0.9+1.5i$	
B38701 [256 × 256] ³	SNR	3.17	3.06	1.22	2.44
	x	$-0.5, -4$	$-0.2, -4$	$-1-1.7i, -1+1.7i$	
	y	$-0.4, -4$	$-0.2, -4$	$-0.1, -0.2$	

¹ The textures *D2* and *D79* originate from the Brodatz database (<http://www.ux.uis.no/~tranden/brodatz.html>) while the biomedical images *B37251* and *B38701* originate from the Cell Image Library (www.cellimagelibrary.org). The image data are used without the DC components to allow for proper parameter estimation; this preserves the differences between the SNR values for a given image.

² The real and imaginary parts of the poles lie in the search spaces $[-5, -0.1]$ and $[0.1, 3]$. The corresponding step is 0.1 for each variable.

³ These data consist of cropped zones that are visually stationary.

mostly inadequate in that case (it yields SNR values that are consistently lower than all other methods), the l_2 autocorrelation fit was found to be efficient, including when the image resolution was low. The corresponding SNR values were nearly optimal compared to oracle performance within the considered search space (they were off by no more than 0.29 dB). This observation corroborates the theory: fitting the autocorrelation function does not depend on the probabilistic properties of the continuous-domain innovation process, and is thus more robust. As an illustration, Figure 3 shows one image of Table 5 with its X and Y autocorrelation values fitted using our estimation methods. We observe that the best-SNR and l_2 -based criteria are able to fit the autocorrelation satisfactorily, unlike the maximum-likelihood estimation approach.

6. CONCLUSIONS

We proposed an estimation method that quantifies some statistical properties of an image based on a continuous-domain autoregressive model. We exploited the link between this model and exponential splines to perform resolution enhancement. Using adaptive exponential-B-spline kernels, we implemented our interpolation algorithm through digital filtering, the only additional computational cost consisting in the estimation of the model parameters. Our experiments with texture images and two-dimensional cell images demonstrated better results compared to the piecewise-polynomial-B-spline model of the same order.

7. REFERENCES

- [1] J.V. Hajnal, D.L.G. Hill, and D.J. Hawkes, *Medical Image Registration*, CRC Press, 2001.
- [2] P. Thévenaz, T. Blu, and M. Unser, “Interpolation revisited,” *IEEE Trans. Med. Imag.*, vol. 19, pp. 739–758, July 2000.
- [3] E. H. W. Meijering, W. J. Niessen, and M. A. Viergever, “Quantitative evaluation of convolution-based methods for medical image interpolation,” *Med. Image Anal.*, vol. 5, no. 2, pp. 111–126, 2001.
- [4] T. Chen and R. deFigueiredo, “Two-dimensional interpolation by generalized spline filters based on partial differential equation image models,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 33, no. 3, pp. 631–642, Jun 1985.
- [5] P. Protter, *Stochastic Integration and Differential Equations*, Springer, New York, 2004.
- [6] I. Gelfand and N. Y. Vilenkin, *Generalized Functions*, vol. 4 of *Applications of Harmonic Analysis*, Academic, New York, 1964.
- [7] M. Unser and P. D. Tafti, “Stochastic models for sparse and piecewise-smooth signals,” *IEEE Trans. on Signal Processing*, vol. 59, no. 3, pp. 989–1006, February 2011.
- [8] H. Kirshner, S. Maggio, and M. Unser, “A sampling theory approach to continuous ARMA identification,” *IEEE Trans. Sig. Proc.*, vol. 59, no. 10, pp. 4620–4634, October 2011.
- [9] A. Bourquard, H. Kirshner, and M. Unser, “Resolution-invariant separable ARMA modeling of images,” in *ICIP’11*, 2011, pp. 1873–1876.
- [10] J. P. Ward, H. Kirshner, and M. Unser, “Is uniqueness lost for under-sampled continuous-time auto-regressive processes?,” *IEEE Signal Processing Letters*, vol. 19, no. 4, pp. 183–186, April 2012.