

Random Matrix Theory for Modeling the Eigenvalue Distribution of Images under Upscaling

David Vázquez-Padín, Fernando Pérez-González, and Pedro Comesaña-Alfaro

University of Vigo, Signal Theory and Communications Department, Vigo, Spain
{dvazquez, fperez, pcomesan}@gts.uvigo.es

Abstract. The stochastic representation of digital images through a two-dimensional autoregressive (2D-AR) model offers a proper way to approximate the empirical distribution of the eigenvalues coming from genuine images. By considering this model, we apply random matrix theory to analytically derive the asymptotic eigenvalue distribution of causal 2D-AR random fields that have undergone an upscaling operation with a particular interpolation kernel. This eigenvalue characterization is useful in developing new forensic techniques for image resampling detection since we can use theoretical bounds to drive the decision of detectors based on subspace decomposition. Moreover, experimental results with real images show that the obtained asymptotic limits turn out to be excellent approximations, even when working with images of small size.

Keywords: Image forensics, Marčenko-Pastur law, random matrix theory, resampling detection, two-dimensional autoregressive model.

1 Introduction

In multimedia forensics, the detection of resampling traces in a digital signal is an important step towards determining whether its content has been manipulated. For example, the presence of resampling artifacts in a specific region of a digital image is indicative of the application of a geometric transformation in that part of the captured scene. This is due to the fact that any spatial transformation requires the calculation of new samples through the interpolation of existing ones, which unavoidably leaves characteristic dependencies among adjacent samples.

Since by exposing this characteristic footprint the presence of forgeries can be eventually revealed, the forensic analysis of resampled signals has been largely investigated during last years. Although many different directions have been explored [10, 4, 7, 5], most of the available techniques share a common processing structure: first, a residue signal (where resampling artifacts are observable) is extracted from the image under investigation; then, given the periodic nature of the residue signal in case of resampling, a frequency analysis is carried out to unveil the presence of resampling traces. Very recently, however, we have proposed in [13] and [15] a different way of analyzing resampled images under an

upsampling scenario, i.e., by approaching image resampling detection as a subspace decomposition problem. In these works we have shown that an upscaled image can be seen as a low-dimensional signal surrounded by a high-dimensional noise, which through a Singular Value Decomposition (SVD) as in [13] or a precise eigenvalue characterization as in [15] can be used to expose the implicit signal-plus-noise structure from upscaled images that is not present in genuine ones.

In this paper, we complement the work in [15] by detailing the analytical derivation of the asymptotic eigenvalue distribution of causal 2D-AR random fields (used as an underlying model for genuine images) in an upscaling scenario. To characterize the behavior of the eigenvalues, we resort to Random Matrix Theory (RMT) and by using one of its fundamental results, i.e., the Marčenko-Pastur law [8], we propose theoretical bounds to drive the decision of the resampling detector designed in [15]. Even though the analytically obtained decision threshold is only applicable to images without demosaicing traces, we believe that avoiding the use of a training set to establish an empirical threshold is an advantage over state-of-the-art techniques. Finally, since the robustness of the resampling detector depends on the relation between the variance of the signal and that of the background noise, we analyze the performance limits of the proposed approach from a theoretical perspective using the 2D-AR model.

The structure of the paper is as follows. In Sect. 2, we first introduce the notation used throughout the paper, then we describe the resampling process together with the subspace decomposition approach, and finally we present the 2D-AR model adopted for representing genuine images. The eigenvalue characterization of genuine and upscaled images is carried out in Sect. 3, while the experimental validation with real images is performed in Sect. 4. Conclusions are finally drawn in Sect. 5.

2 Preliminaries

Prior to the formulation of the resampling procedure and the formalization of the model adopted for natural images, we briefly describe the notation used throughout the paper. Matrices will be denoted by capital bold letters. Given an arbitrary square matrix \mathbf{A} of size $N \times N$, we will frequently work with a rectangular submatrix \mathbf{A}_K of size $N \times K$ and aspect ratio $\beta \triangleq \frac{K}{N}$, which is constructed by extracting $K \leq N$ consecutive columns from \mathbf{A} . The sample autocorrelation matrix of \mathbf{A}_K is $K^{-1} \mathbf{A}_K \mathbf{A}_K^T$, but we will mostly work with the renormalized sample autocorrelation $\Sigma_{A_K} \triangleq N^{-1} \mathbf{A}_K \mathbf{A}_K^T$. By convention, we assume that the eigenvalues of a given matrix are sorted in descending order. Table 1 summarizes the notation used for representing different attributes of an arbitrary $N \times N$ matrix \mathbf{A} .

2.1 Resampling Process Description

Let us assume that the input to the resampling operation is a digital image with a single color channel (hereinafter, indistinctly called genuine or natural image) that does not contain traces from any preceding demosaicing process and has negligible quantization noise. Representing the referred input image as a matrix

Table 1. Summary of notation

Symbol	Meaning
$A_{i,j}$	(i,j) -th element of matrix \mathbf{A} , with $i, j \in \{0, \dots, N-1\}$
σ_A^2	Variance of the entries of \mathbf{A}
\mathbf{A}_K	$N \times K$ submatrix with K consecutive columns from \mathbf{A}
β	Aspect ratio of an $N \times K$ matrix, i.e., $\beta = \frac{K}{N}$
Σ_{A_K}	Renormalized sample autocorrelation matrix of \mathbf{A}_K
$\lambda_i(\Sigma_{A_K})$	i -th eigenvalue of Σ_{A_K} , with $i = 1, \dots, N$
$\lambda_-(\Sigma_{A_K})$	Smallest eigenvalue of Σ_{A_K}
$\lambda_+(\Sigma_{A_K})$	Largest eigenvalue of Σ_{A_K}
$f_{\Sigma_{A_K}}(\lambda)$	Probability density function of the eigenvalues of Σ_{A_K}

\mathbf{X} , we consider that the resampling operation scales the two dimensions of \mathbf{X} in a uniform way and by means of a resampling factor $\xi > 1$, which is defined as $\xi \triangleq \frac{L}{M}$ with L and M coprime natural numbers. The pixel values in the new resampled grid are calculated by linearly combining a finite set of neighboring samples from the input matrix \mathbf{X} with a two-dimensional (2D) interpolation kernel that is symmetric and separable (i.e., the kernel can be expressed as the outer product of two vectors). Consequently, each (i, j) -th pixel value $Y_{i,j}$ from an upscaled image \mathbf{Y} can be computed as

$$Y_{i,j} = \sum_{u,v \in \mathbb{Z}} h\left(i\frac{M}{L} + \phi - u\right) h\left(j\frac{M}{L} + \phi - v\right) X_{u,v}, \quad (1)$$

where $X_{u,v}$ denotes the (u, v) -th element of the genuine image \mathbf{X} , ϕ is a shift between the original and the resampled grid,¹ and $h(\cdot)$ is the one-dimensional (1D) impulse response of the interpolation kernel, whose width is $k_w \in 2\mathbb{Z}^+$, i.e., k_w is a positive even number.

Our analysis will be carried out over a block \mathbf{Y} of size $N \times N$ extracted from the upscaled version by ξ of a rectangular $U \times V$ genuine image \mathbf{X} , i.e., with $N \leq \xi \cdot \min(U, V)$. Assuming that both matrices \mathbf{X} and \mathbf{Y} are aligned with respect to their upper-left corner and also that N is a multiple of L , the $N \times N$ block \mathbf{Y} can be expressed using (1) in matrix form:

$$\mathbf{Y} = \mathbf{H}\hat{\mathbf{X}}\mathbf{H}^T, \quad (2)$$

where the entries of \mathbf{H} are given by

$$H_{i,j} = h\left(i\frac{M}{L} + \varphi - j\right), \quad (3)$$

with $\varphi \triangleq \phi + k_w/2 - 1$, such that $\hat{\mathbf{X}}$ only contains the first $R \times R$ pixels from \mathbf{X} that are used in the calculation of \mathbf{Y} with $R \triangleq N\frac{M}{L} + k_w$, i.e., $\hat{X}_{u,v} = X_{u,v}$, for $u, v = 0, \dots, R-1$. Provided that N is larger than R , we will assume that the $N \times R$ interpolation matrix \mathbf{H} is full column rank (i.e., $\text{rank}(\mathbf{H}) = R$), which is true for all kernels considered in this paper. However, the reader is referred to [6] to check the conditions that \mathbf{H} must satisfy to have full column rank.

After the calculation of all the resampled pixels from the upscaled image, a quantization is generally applied as a last step to fit the original precision of the

¹ The value of ϕ is generally taken as $\phi = \frac{1}{2} \left(\frac{M}{L} + 1 \right)$, as in `imresize` from MATLAB.

input image \mathbf{X} . We assume that a uniform scalar quantizer with step size Δ is used. Denoting the resulting quantized image by a matrix \mathbf{Z} , we can model it as

$$\mathbf{Z} = \mathbf{Y} + \mathbf{W}, \quad (4)$$

where we opt for an additive white noise model, such that the entries of the quantization noise matrix \mathbf{W} are i.i.d. $\mathcal{U}[-\frac{\Delta}{2}, \frac{\Delta}{2})$ with zero mean and variance $\sigma_W^2 = \frac{\Delta^2}{12}$. This model is suitable in practice as long as the probability density function (pdf) of the genuine image is smooth and its variance is much larger than the variance of the quantization noise, but it will not be valid for modeling flat regions, which might appear in practice due to undesired artifacts such as saturation. In Sect. 4 we will see how to deal with these special cases.

2.2 Subspace Decomposition to Expose Resampling Traces

As noted above, it was first observed in [13] that an upscaled and later quantized image presents a particular signal-plus-noise structure that does not show up in natural images. This comes from the fact that when a genuine image \mathbf{X} is up-scaled by ξ , the rank of \mathbf{Y} in (2) is at most equal to the rank of the interpolation matrix \mathbf{H} (i.e., $\text{rank}(\mathbf{Y}) \leq \text{rank}(\mathbf{H}) = R$ with $R \approx N/\xi$), which implies that any N -dimensional column/row from \mathbf{Y} will lie in an R -dimensional subspace. On the other hand, the noise \mathbf{W} in (4) is expected to span the full space. Therefore, from the eigendecomposition of the sample autocorrelation matrix of \mathbf{Z} , we count on finding R leading eigenvalues corresponding to the signal subspace, while the remaining ones will correspond to the background noise. This makes possible the detection of resampling traces using any subspace decomposition approach.

In this work, we draw on the most-celebrated result from RMT, i.e., Marčenko-Pastur Law (MPL) [8], whose main characteristic states that for a given random matrix \mathbf{A}_K with i.i.d. entries and aspect ratio β , the eigenvalues of Σ_{A_K} tend to cluster around the variance σ_A^2 as β converges to zero. Respectively, we are interested in checking whether this property also applies when a submatrix \mathbf{Z}_K with small β is extracted from \mathbf{Z} in (4), so that on the one hand the signal eigenvalues cluster around σ_Z^2 , while the noise eigenvalues do so around σ_W^2 . The eigenvalue characterization carried out in Sect. 3 will theoretically support these findings.

2.3 Causal 2D-AR Random Fields for Modeling Genuine Images

Based on our previous analysis in [14], we consider a causal two-dimensional autoregressive model with a single correlation coefficient ρ as the stochastic representation for natural images without traces of demosaicing, so denoting by \mathbf{X} the generated 2D-AR random field of size $N \times N$, it can be expressed as

$$\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{U}^T, \quad (5)$$

where \mathbf{S} is an $(N + Q - 1) \times (N + Q - 1)$ random matrix whose i.i.d. entries follow a zero-mean normal distribution with variance σ_S^2 , and \mathbf{U} is a Toeplitz matrix of size $N \times (N + Q - 1)$, with Q denoting the length of the truncated infinite impulse response of the AR model. Hence, matrix \mathbf{U} is fully described

Table 2. Impulse response and width of several interpolation kernels.

Kernel Type	Impulse Response
Linear ($k_w = 2$)	$h(t) = \begin{cases} 1 - t , & \text{if } t \leq \frac{k_w}{2} \\ 0, & \text{otherwise} \end{cases}$
Catmull-Rom ($k_w = 4$)	$h(t) = \begin{cases} \frac{3}{2} t ^3 - \frac{5}{2} t ^2 + 1, & \text{if } t \leq \frac{k_w}{4} \\ -\frac{1}{2} t ^3 + \frac{5}{2} t ^2 - 4 t + 2, & \text{if } \frac{k_w}{4} < t \leq \frac{k_w}{2} \\ 0, & \text{otherwise} \end{cases}$
B-spline ($k_w = 4$)	$h(t) = \begin{cases} \frac{1}{2} t ^3 - t ^2 + \frac{2}{3}, & \text{if } t \leq \frac{k_w}{4} \\ -\frac{1}{6} t ^3 + t ^2 - 2 t + \frac{4}{3}, & \text{if } \frac{k_w}{4} < t \leq \frac{k_w}{2} \\ 0, & \text{otherwise} \end{cases}$
Lanczos ($k_w = 6$)	$h(t) = \begin{cases} \text{sinc}(t)\text{sinc}\left(\frac{t}{3}\right), & \text{if } t < \frac{k_w}{2} \\ 0, & \text{otherwise} \end{cases}$

as $U_{i,j} = u_Q[j - i]$, where sequence $u_Q[n] \triangleq \rho^{Q-1-n}$, for $n = 0, \dots, Q - 1$, and is zero elsewhere, i.e.,

$$U_{i,j} = \begin{cases} \rho^{Q-1-(j-i)}, & \text{if } (j - i) = 0, \dots, Q - 1 \\ 0, & \text{otherwise} \end{cases},$$

where the value of Q is generally taken as $Q \geq N$ to minimize modeling errors due to truncation. The practical suitability of this model is also reported in [14].

2.4 Interpolation Kernels

Even though the upcoming eigenvalue characterization can be applied to any linear and separable interpolation kernel, we will provide examples with the most commonly available kernels, namely: Linear, Catmull-Rom, B-spline, and Lanczos, to highlight the different behavior among them and specific properties of each other. Table 2 collects their 1D impulse response and their width.

3 Eigenvalue Characterization

In this section, we formally characterize the statistical distribution of the sample eigenvalues coming from genuine images (modeled as causal 2D-AR random fields) and their upscaled versions (before and after quantization). As noted above, we base our analysis on well-known results from RMT that allow us to model the asymptotic behavior of the eigenvalues through deterministic functions. Table 3 lists some of the solved problems in RMT that relate to our case.

The first work in Table 3 corresponds to the contribution by Marčenko-Pastur that provides a closed-form expression for the eigenvalue distribution of random matrices with i.i.d. entries. Although this result cannot be applied to filtered white noise processes as the one that emerges from the 2D-AR model, it can actually be used to represent the eigenvalues coming from the matrix \mathbf{W} with i.i.d. quantization noise in (4). On the other hand, the study of the sample eigenvalues of filtered processes has been tackled in the last two works of Table 3, where the considered random matrices have independent columns, but an arbitrary dependence within each column. In both cases, no closed-form

Table 3. Solved RMT problems.

Type of matrix	Eigenvalue characterization from RMT
Random matrix (i.i.d. entries)	Marčenko-Pastur law [8] (closed-form expression)
Random matrix (i.i.d. stationary time series in columns)	Bai and Zhou [1] (through Stieltjes transform)
Random matrix (a linear process in each column)	Pfaffel and Schlemm [9] (through Stieltjes transform)

solution is obtained, but the eigenvalues are characterized indirectly through the Stieltjes transform, whose inversion formula uniquely determines their pdf.

Unfortunately, given the form of our random matrix \mathbf{X} in (5) and its up-scaled version \mathbf{Y} in (2), where both matrices present dependencies among rows and columns, neither of these results can be applied directly to our case. Therefore, we also opt for characterizing the pdf of the sample eigenvalues indirectly through the Stieltjes transform. To that end, we resort to a procedure proposed by Tulino and Verdú in [12, Theorem 2.43] that addresses the calculation of the so-called η -transform [12, Sect. 2.2.2] of an unnormalized sample autocorrelation matrix $\mathbf{B}\mathbf{B}^T$, where \mathbf{B} has the form $\mathbf{B} = \mathbf{C}\mathbf{S}\mathbf{A}$, which coincides with the matrix form of \mathbf{X} and \mathbf{Y} . Similarly, matrix \mathbf{S} has i.i.d. entries, and matrices \mathbf{C} and \mathbf{A} induce the linear dependencies. Once the η -transform is obtained, the relationship in [12, Eq. (2.48)] allows us to retrieve the corresponding Stieltjes transform, whose inversion formula in [12, Eq. (2.45)] finally provides the pdf of the sample eigenvalues.

Algorithm 1 summarizes the above steps to obtain the asymptotic pdf of the eigenvalues of $\boldsymbol{\Sigma}_{B_K} = N^{-1}\mathbf{B}_K\mathbf{B}_K^T$ when considering an arbitrary submatrix \mathbf{B}_K of size $N \times K$, which is made up of $K \leq N$ consecutive columns from an $N \times N$ matrix \mathbf{B} . Accordingly, in the following subsections, we particularize the calculation of the asymptotic eigenvalue distribution for genuine images in Sect. 3.1, and for unquantized upsampled images in Sect. 3.2. Finally, pairing these results with the characterization of the quantization noise through the MPL, we describe in Sect. 3.3 the gap that emerges in images that have been upsampled and later quantized.

For the sake of simplicity, and without loss of generality, in Sects. 3.1 and 3.2 we assume that the entries of \mathbf{S} are i.i.d. with variance $\sigma_S^2 = 1$ as in [12]. More generally, in Sect. 3.3, we assume an arbitrary variance σ_S^2 to further evaluate the effect of the signal-to-quantization-noise ratio. Notice that the eigenvalue distribution for $\sigma_S^2 \neq 1$ is directly obtained by multiplying by σ_S^2 the eigenvalues for $\sigma_S^2 = 1$.

3.1 Eigenvalue Distribution for Genuine Images

Given the 2D-AR causal model for genuine images in (5), here we analyze the eigenvalues of $\boldsymbol{\Sigma}_{X_K}$ when \mathbf{X}_K is constructed by taking K consecutive columns from the $N \times N$ random field \mathbf{X} (without loss of generality, we will assume that we retain the *first* K columns), such that

$$\mathbf{X}_K = \mathbf{U}_N \mathbf{S} \mathbf{U}_K^T, \quad (6)$$

Algorithm 1 Main steps to obtain the asymptotic eigenvalue pdf of Σ_{B_K}

1. Compute the η -transform (see [14, Appendix A]): $\eta_{\Sigma_{B_K}}(\gamma)$
 2. Retrieve the Stieltjes transform through: $\eta_{\Sigma_{B_K}}(\gamma) = \frac{\mathcal{S}_{\Sigma_{B_K}}(-\frac{1}{\gamma})}{\gamma}$
 3. Apply the inversion formula: $f_{\Sigma_{B_K}}(\lambda) = \lim_{\nu \rightarrow 0^+} \frac{1}{\pi} \text{Im} \left[\mathcal{S}_{\Sigma_{B_K}}(\lambda + j\nu) \right]$
-

where \mathbf{U}_K is obtained from \mathbf{U} by keeping the first K rows, and $\mathbf{U}_N = \mathbf{U}$ is so written to stress the fact that it has N rows. In (6) \mathbf{S} is an $(N+Q-1) \times (N+Q-1)$ random matrix with i.i.d. $\mathcal{N}(0, 1)$ entries.

For the calculation of the η -transform of Σ_{X_K} according to [12, Theorem 2.43], we first need to derive the asymptotic spectra of the matrices that induce the linear dependencies, i.e., $\mathbf{D} = \mathbf{U}_N \mathbf{U}_N^T$ and $\mathbf{T} = \mathbf{U}_K^T \mathbf{U}_K$. It is easy to show that both matrices are full-rank and, for a sufficiently large value of K and N , have identical asymptotic spectra. Therefore, we focus on obtaining the asymptotic spectrum of \mathbf{D} , which can be seen as an unnormalized version of the autocorrelation matrix of \mathbf{U}_N , itself described by $u_Q[n]$ (see Sect. 2.3).

If we assume that $Q \geq N$, we can straightforwardly show that \mathbf{D} is a symmetric Toeplitz matrix, whose (i, j) -th element is given by $(1 - \rho^{2(Q-|i-j|)}) \cdot \rho^{|i-j|} / (1 - \rho^2)$. If we let $Q \rightarrow \infty$, this simplifies to $D_{i,j} = \rho^{|i-j|} / (1 - \rho^2)$. Now, in order to calculate the asymptotic eigenvalue spectrum of \mathbf{D} , we can invoke Szegő's fundamental theorem of eigenvalue distribution [3], which establishes that for a Toeplitz symmetric matrix \mathbf{D} defined by sequence $d[n]$ (i.e., $d[|i-j|] \triangleq D_{i,j}$), where $d[n]$ is absolutely summable,² its asymptotic eigenvalue spectrum tends to the Fourier transform of $d[n]$, i.e., $d(\omega)$. Therefore, when $N \rightarrow \infty$, the eigenvalues of \mathbf{D} and (by extension) \mathbf{T} will asymptotically converge to

$$d(\omega) = \frac{1}{1 + \rho^2 - 2\rho \cos(\omega)}, \quad \omega \in [0, 2\pi). \quad (7)$$

The η -transform of Σ_{X_K} depends on two independent random variables \mathfrak{D} and \mathfrak{T} which are distributed as the asymptotic spectra of \mathbf{D} and \mathbf{T} . Both random variables can be seen as the result of applying the transformation $d(\omega)$ in (7) to a random variable $\Omega \sim \mathcal{U}[0, 2\pi)$, such that $\mathfrak{D} = d(\Omega)$ and $\mathfrak{T} = d(\Omega)$. Therefore, as described in [14, Appendix A], once $d(\omega)$ is available, the η -transform of Σ_{X_K} can be numerically calculated. However, for completing the pdf of the eigenvalues of Σ_{X_K} we need to calculate its asymptotic fraction of zero eigenvalues (in the sequel, AFZE), which is

$$1 - \min \{ \beta P(\mathfrak{T} \neq 0), P(\mathfrak{D} \neq 0) \}. \quad (8)$$

Given that \mathbf{D} and \mathbf{T} have full rank and their asymptotic spectra through (7) satisfy $d(\omega) > 0$, we know that $P(\mathfrak{T} \neq 0) = P(\mathfrak{D} \neq 0) = 1$. Therefore, since $\beta \leq 1$, the AFZE of Σ_{X_K} equals $(1 - \beta)$. So finally, applying steps 2 and 3 from

² There is an additional technical condition that applies in the cases considered in this paper, namely, that the set $\{\omega : d(\omega) = x\}$ has measure zero for all $x \in \mathbb{R}$.




Fig. 1. Partial representation of the analytically derived pdf of the eigenvalues of Σ_{X_K} , i.e., $f_{\Sigma_{X_K}}(\lambda)$, for different values of β and $\rho = 0.97$ in (a). Graphical comparison in (b) of the theoretical result for $\beta = 0.125$ and the histogram corresponding to 50,000 realizations ($N = 1024$). The mass points at $\lambda = 0$ (present for $\beta < 1$) are not shown.

Algorithm 1, we obtain the asymptotic pdf of the eigenvalues of Σ_{X_K} , i.e.,

$$f_{\Sigma_{X_K}}(\lambda) = (1 - \beta)\delta(\lambda) + \lim_{\nu \rightarrow 0^+} \frac{1}{\pi} \text{Im} \left[\mathcal{S}_{\Sigma_{X_K}}(\lambda + j\nu) \right],$$

where $\delta(\cdot)$ denotes the Dirac delta function.

Figure 1(a) depicts the derived pdf $f_{\Sigma_{X_K}}(\lambda)$ for different values of β (fixing $\rho = 0.97$) and quite remarkably, as predicted by MPL for random matrices with i.i.d. entries, we find that the lower the value of β , the more the eigenvalues of Σ_{X_K} get squeezed towards the variance of \mathbf{X} . This generalization of the MPL to stochastic representations of genuine images is key for our purposes. Furthermore, in Figure 1(b), we can see that the conformity between the empirical pdf for finite matrices ($N = 1024$) and its asymptotic version is nearly perfect.

3.2 Eigenvalue Distribution for (Unquantized) Upscaled Images

In this case, we are interested in computing the eigenvalue distribution of Σ_{Y_K} , where \mathbf{Y}_K is a submatrix of size $N \times K$ that is extracted from an upscaled image \mathbf{Y} . Modeling $\tilde{\mathbf{X}}$ in (2) as in (5), we can write

$$\mathbf{Y}_K = \mathbf{H}_N \mathbf{U} \mathbf{S} \mathbf{U}^T \mathbf{H}_K^T,$$

where \mathbf{U} is now a Toeplitz matrix of size $R \times (R + Q - 1)$ described by sequence $u_Q[n]$. In the above equation, matrices \mathbf{H}_N and \mathbf{H}_K are constructed as in (3), with respective sizes $N \times R$ and $K \times R$, containing both shifted copies of the L different polyphase components of $h(iM/L + \varphi)$, $i \in \mathbb{Z}$. Notice that (3) is such that the first row in \mathbf{H}_N (or \mathbf{H}_K) corresponds to the zeroth polyphase component; however, as we will see, this arbitrary assignment has no effect on our analytical derivations.

As in the previous section, in order to compute the η -transform of Σ_{Y_K} , we need to characterize the asymptotic spectra of matrices $\mathbf{D} = \mathbf{C}\mathbf{C}^T = \mathbf{H}_N \mathbf{U} \mathbf{U}^T \mathbf{H}_N^T$ and $\mathbf{T} = \mathbf{A}\mathbf{A}^T = \mathbf{U}^T \mathbf{H}_K^T \mathbf{H}_K \mathbf{U}$. For convenience, we start focusing on the nonzero eigenvalues of $\mathbf{D} = \mathbf{C}\mathbf{C}^T$, which are the same as those of matrix $\mathbf{D}' \triangleq \mathbf{C}^T \mathbf{C}$. Let us consider its inner matrix $\mathbf{R} \triangleq \mathbf{H}_N^T \mathbf{H}_N$, with entries given by

$$R_{i,j} = \sum_{l=0}^{R-1} h\left(\frac{lM}{L} + \varphi - i\right) h\left(\frac{lM}{L} + \varphi - j\right). \quad (9)$$

Although this matrix is not Toeplitz, it can be seen to contain in its rows the different components of a polyphase decomposition of the kernel autocorrelation function. Since these rows are roughly similar, it makes sense to convert \mathbf{R} into Toeplitz by averaging those components. Let $\bar{\mathbf{R}}$ be such matrix, with

$$\begin{aligned}\bar{R}_{i,j} &= \frac{1}{M} \sum_{k=0}^{M-1} \sum_{l=0}^{R-1} h\left(l\frac{M}{L} + \frac{k}{L} + \varphi - i\right) h\left(l\frac{M}{L} + \frac{k}{L} + \varphi - j\right) \\ &= \frac{1}{M} \sum_{k=0}^{R \cdot M - 1} h\left(\frac{k}{L} + \varphi - i\right) h\left(\frac{k}{L} + \varphi - j\right),\end{aligned}\quad (10)$$

which, as it can be readily checked is symmetric Toeplitz, so $\bar{R}_{i,j}$ only depends on $|i - j|$. Then, $\bar{\mathbf{R}}$ is completely characterized by the sequence $r_{hh}[|i - j|] \triangleq \bar{R}_{i,j}$. Now, since \mathbf{U} is also Toeplitz, by expressing products of Toeplitz matrices as convolutions of their corresponding representative sequences, it is possible to see that $\mathbf{D}' = \mathbf{U}^T \mathbf{H}_N^T \mathbf{H}_N \mathbf{U}$ is described by sequence $u_Q[n] * r_{hh}[n] * u_Q[-n]$. This sequence is absolutely summable even for $Q \rightarrow \infty$ and \mathbf{D}' is symmetric Toeplitz; therefore, we can resort again to Szegő's theorem [3] to approximate the asymptotic eigenvalue distribution of \mathbf{D}' when both $Q, N \rightarrow \infty$, as follows

$$d'(\omega) = \left(\frac{1}{1 + \rho^2 - 2\rho \cos(\omega)} \right) \sum_{n=-(k_w-1)}^{k_w-1} r_{hh}[n] \cos(n\omega), \quad (11)$$

where $\omega \in [0, 2\pi)$. This discussion extends to the nonzero eigenvalues of \mathbf{T} , which can also be approximated by $d'(\omega)$.

A crucial difference with respect to the case of genuine images (Sect. 3.1) is that now both \mathbf{D} and \mathbf{T} will have null eigenvalues, given that matrices \mathbf{H}_N and \mathbf{H}_K do not have full rank. Actually, the ranks of both matrices depend on the applied resampling factor ξ , and it is easy to check that $\lim_{N \rightarrow \infty} \text{rank}(\mathbf{H}_N)/N \rightarrow \xi^{-1}$, and $\lim_{K \rightarrow \infty} \text{rank}(\mathbf{H}_K)/K \rightarrow \xi^{-1}$. From the rank properties for real matrices, we have that $\text{rank}(\mathbf{D}) = \text{rank}(\mathbf{C}) = \text{rank}(\mathbf{H}_N)$ because \mathbf{U} has full rank, so $\text{rank}(\mathbf{D})/N \rightarrow \xi^{-1}$ and, by extension, $\text{rank}(\mathbf{T})/K \rightarrow \xi^{-1}$. We can conclude that the AFZE of \mathbf{D} and \mathbf{T} is given by $(1 - \xi^{-1})$.

Figure 2 depicts the eigenvalues of matrix \mathbf{D} for different values of ξ and for the interpolation kernels in Table 2, and the approximation in (11). Although the derived approximation is generally accurate, we observe that for certain pairs of ξ and kernel, it is not able to follow the existing discontinuities. The most evident examples show up with the Linear kernel for $\xi = \frac{4}{3}$ and $\xi = \frac{8}{5}$. This difference comes from approximating \mathbf{R} by $\bar{\mathbf{R}}$. However, since the range of the eigenvalues is well matched and their evolution is tracked up to a good degree, we adopt this approximation.

From the plots in Figure 2, it is interesting to observe that the largest eigenvalues of the four kernels converge to the same value, while the smallest ones differ noticeably. Although for Catmull-Rom and Lanczos kernels the smallest eigenvalues are almost identical, the Linear kernel provides smaller eigenvalues as ξ gets closer to 1, and the B-spline kernel produces the smallest eigenvalues at almost one order of magnitude below.




Fig. 2. Evolution of the nonzero eigenvalues of \mathbf{D} for different values of ξ and the interpolation kernels in Table 2. *Solid lines* represent the ordered eigenvalues $\lambda_i(\mathbf{D})$, while *dashed lines* correspond to the approximation $d'(\omega)$ in (11) sorted in descending order, with $\omega = 2\pi \frac{i-1}{N\xi-1}$, $i \in \{1, \dots, N\xi^{-1}\}$, ($N = 1024$).

The procedure to compute the η -transform of Σ_{Y_K} is exactly the same as the one followed with genuine images by taking $d'(\omega)$ as $d(\omega)$ in [14, Appendix]. However, note that when computing E_1 and E_2 as in [14, Eq. 11] and $\eta_{\Sigma_{Y_K}}$ as in [14, Eq. 12], we must take into account that now the random variable Ω is of mixed type with a probability mass at $\omega = 0$ of size $(1 - \xi^{-1})$, and a continuous pdf $f_\Omega(\omega) = (2\pi\xi)^{-1}$ in $(0, 2\pi)$, i.e., the pdf of a uniform random variable, but scaled by ξ^{-1} so that the total probability adds up to 1.

Once the η -transform has been computed, we have to determine the AFZE of Σ_{Y_K} , which is given by (8). In this case, we know that $P(\mathfrak{F} \neq 0) = P(\mathfrak{D} \neq 0) = \xi^{-1}$ and so the AFZE equals $(1 - \beta\xi^{-1})$ for $\beta \leq 1$. Finally, applying steps 2 and 3 from Algorithm 1, we have

$$f_{\Sigma_{Y_K}}(\lambda) = (1 - \beta\xi^{-1}) \delta(\lambda) + \lim_{\nu \rightarrow 0^+} \frac{1}{\pi} \text{Im} \left[\mathcal{S}_{\Sigma_{Y_K}}(\lambda + j\nu) \right].$$

The upper panels of Figure 3 show the derived pdf $f_{\Sigma_{Y_K}}(\lambda)$ for different interpolation kernels and values of β , whereas the lower ones compare the theoretical pdf against its empirical version under a fixed setting with $N = 1024$. Clearly, the theoretical model fits very well the empirical eigenvalue distribution and, in view of the shape of the pdfs, there is no doubt that the eigenvalues of Σ_{Y_K} are compacted in the same way as those from Σ_{X_K} , which confirms that the squeezing effect of the eigenvalues as β decreases is also valid for upscaled images. This feature will be decisive for distinguishing the signal subspace from the background noise. Moreover, the influence of each particular interpolation kernel on the distribution of the eigenvalues of Σ_{Y_K} for a fixed β can be appreciated in the lower pannels of Figure 3. Consistent with the conclusions drawn from Figure 2, the B-spline kernel concentrates the smallest eigenvalues towards zero. As we will see in Sect. 4, the distinct behavior among these kernels will result in different performance when conducting resampling detection.

3.3 Eigenvalue Distribution for Upscaled & Quantized Images

Adopting the $N \times N$ matrix \mathbf{Z} in (4) as the model for upscaled and later quantized images, here we analyze the eigenvalues of Σ_{Z_K} , where \mathbf{Z}_K is an $N \times K$

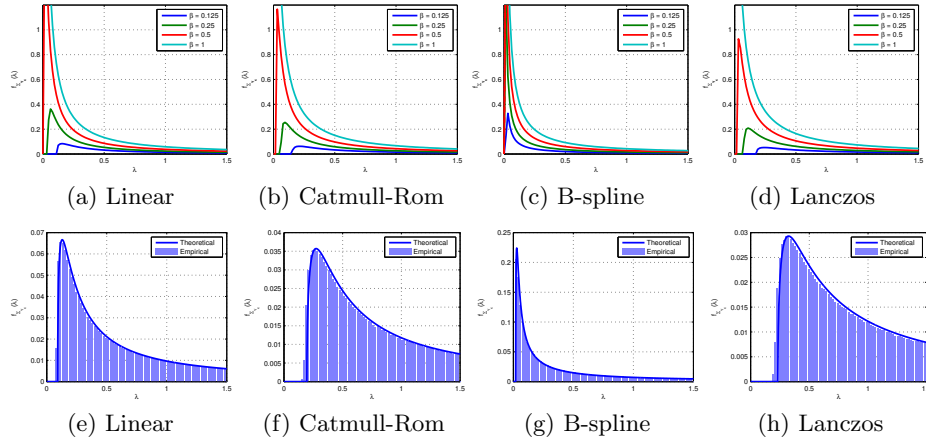


Fig. 3. Partial representation of the analytically derived pdf $f_{\Sigma_{Y_K}}(\lambda)$ for different interpolation kernels and values of β in the *upper panels* ($\xi = 1.5$, $\rho = 0.97$). Graphical comparison in the *lower panels* of the theoretical pdf and the histogram from 50,000 realizations ($N = 1024$, $\beta = 0.125$, $\xi = 2$, $\rho = 0.97$). All mass points are omitted.

submatrix extracted from \mathbf{Z} . In particular, drawing on the previous eigenvalue characterization of Σ_{Y_K} and using the MPL for modeling the eigenvalues coming from the quantization noise matrix \mathbf{W} , we delve into the characteristics of the signal-plus-noise structure of \mathbf{Z} already discussed in Sect. 2.2.

From the preceding analysis in Sect. 3.2 we know that for a sufficiently large value of K , the submatrix \mathbf{Y}_K has rank P strictly smaller than K (i.e., $\text{rank}(\mathbf{Y}_K)/K \rightarrow \xi^{-1}$). Therefore, in Σ_{Z_K} , there are P leading eigenvalues corresponding to the signal subspace and $K - P$ corresponding to the background noise (the remaining $N - K$ are zero). To model this transition between the signal subspace and the noise space, we center our attention on the eigenvalues $\lambda_i(\Sigma_{Z_K})$ located on each side of the boundary. Then, we derive a lower bound for the last eigenvalue coming from the signal subspace, i.e., $\lambda_P(\Sigma_{Z_K})$, and an upper bound for the first eigenvalue belonging to the noise space, i.e., $\lambda_{P+1}(\Sigma_{Z_K})$. As fully reported in [14] and [15], these two bounds can be obtained by resorting to Weyl's inequality applied to the singular values [11, Exercise 1.3.22], so that the following lower bound applies for $\lambda_P(\Sigma_{Z_K})$:

$$\lambda_P(\Sigma_{Z_K}) \geq \lambda_P(\Sigma_{Y_K}) - \lambda_1(\Sigma_{W_K}) \geq \sigma_S^2 \lambda_-(\Sigma_{Y_K}) - \lambda_+(\Sigma_{W_K}), \quad (12)$$

where $\lambda_+(\Sigma_{W_K})$ denotes the largest eigenvalue of Σ_{W_K} given by MPL and $\lambda_-(\Sigma_{Y_K})$ represents the smallest nonzero eigenvalue of Σ_{Y_K} . Remember that $\lambda_-(\Sigma_{Y_K})$ can be obtained through the calculation of $f_{\Sigma_{Y_K}}(\lambda)$ as in Sect. 3.2, albeit now it appears scaled by σ_S^2 because the derived pdf $f_{\Sigma_{Y_K}}(\lambda)$ assumes $\sigma_S^2 = 1$. On the other hand, the upper bound for $\lambda_{P+1}(\Sigma_{Z_K})$ is given by

$$\lambda_{P+1}(\Sigma_{Z_K}) \leq \lambda_1(\Sigma_{W_K}) \leq \lambda_+(\Sigma_{W_K}) \rightarrow \sigma_W^2 (1 + \sqrt{\beta})^2. \quad (13)$$

Finally, combining (12) and (13) we can characterize the asymptotic gap that marks the transition between the signal subspace and the noise as

$$\frac{\lambda_P(\boldsymbol{\Sigma}_{Z_K})}{\lambda_{P+1}(\boldsymbol{\Sigma}_{Z_K})} \geq \frac{\sigma_S^2 \lambda_-(\boldsymbol{\Sigma}_{Y_K}) - \lambda_+(\boldsymbol{\Sigma}_{W_K})}{\lambda_+(\boldsymbol{\Sigma}_{W_K})} \rightarrow \left(\frac{\sigma_S^2}{\sigma_W^2} \right) \frac{\lambda_-(\boldsymbol{\Sigma}_{Y_K})}{(1 + \sqrt{\beta})^2} - 1. \quad (14)$$

From (14), it is clear that the magnitude of the gap will monotonically increase with the signal-to-noise ratio $\left(\frac{\sigma_S^2}{\sigma_W^2} \right)$. In addition, it will also increase as $\beta \rightarrow 0$, since $\lambda_-(\boldsymbol{\Sigma}_{Y_K})$ and $(1 + \sqrt{\beta})^2$ respectively increases and decreases as $\beta \rightarrow 0$. However, as reported in [14], the gap will become smaller as both the resampling factor ξ and the correlation coefficient ρ approach 1.

Regarding the impact of the interpolation kernel, since the magnitude of the gap is directly proportional to $\lambda_-(\boldsymbol{\Sigma}_{Y_K})$, we can conclude that the smallest gap will take place with the B-spline kernel since it contributes with the smallest magnitude of $\lambda_-(\boldsymbol{\Sigma}_{Y_K})$ (see Figure 3(e-h)). Conversely, the Lanczos kernel will produce the largest gap together with the Catmull-Rom, whereas the Linear kernel will be halfway between Catmull-Rom and B-spline.

4 Experimental Results

Exploiting the above eigenvalue characterization, we have proposed in [15] a resampling detector whose test statistic κ (cf. [15, Eq. 11]) is used to determine whether an observed matrix \mathbf{Z} comes from a genuine image (i.e., hypothesis \mathcal{H}_0), or from an upscaled and quantized image (i.e., hypothesis \mathcal{H}_1). The experimental validation of the proposed detector proved its practical applicability, but here we are more interested in evaluating its performance from a theoretical perspective.

As an example, given that we know through (14) that the magnitude of the gap between the signal subspace and the noise depends on the signal-to-noise ratio $\left(\frac{\sigma_S^2}{\sigma_W^2} \right)$, we can evaluate the detector performance as a function of this quantity using synthetic 2D-AR random fields. The obtained results of Area Under the Curve (AUC) corresponding to the Receiver Operating Characteristic (ROC) of the detector are shown in Figure 4(a) for different values of $\left(\frac{\sigma_S^2}{\sigma_W^2} \right)$ and ξ (using the Linear kernel and performing 1000 realizations). As expected, the larger the signal-to-noise-ratio, the better the detection performance. Notice also that the convergence to perfect detection is faster for larger resampling factors.

On the other hand, it is interesting to note that we can determine a theoretical threshold to drive the decision of our detector. This comes from the analysis performed in Sect. 3: first, we know that under \mathcal{H}_0 the last nonzero eigenvalue of $\boldsymbol{\Sigma}_{Z_K}$, i.e., $\lambda_K(\boldsymbol{\Sigma}_{Z_K})$, is lower bounded by $\sigma_S^2 \lambda_-(\boldsymbol{\Sigma}_{X_K})$; then, from (13) we know that the upper bound for $\lambda_K(\boldsymbol{\Sigma}_{Z_K})$ under \mathcal{H}_1 is given by MPL at $\sigma_W^2 (1 + \sqrt{\beta})^2$; so, provided that $\sigma_S^2 \gg \sigma_W^2$ (which is typically the case for real images), the condition $\sigma_S^2 \lambda_-(\boldsymbol{\Sigma}_{X_K}) > \sigma_W^2 (1 + \sqrt{\beta})^2$ will be generally satisfied (recall that $\lambda_-(\boldsymbol{\Sigma}_{X_K})$ increases as $\beta \rightarrow 0$). As a conclusion, assuming that this condition is commonly satisfied in most practical cases, the proposed detector can operate with a fixed threshold at $\sigma_W^2 (1 + \sqrt{\beta})^2$.

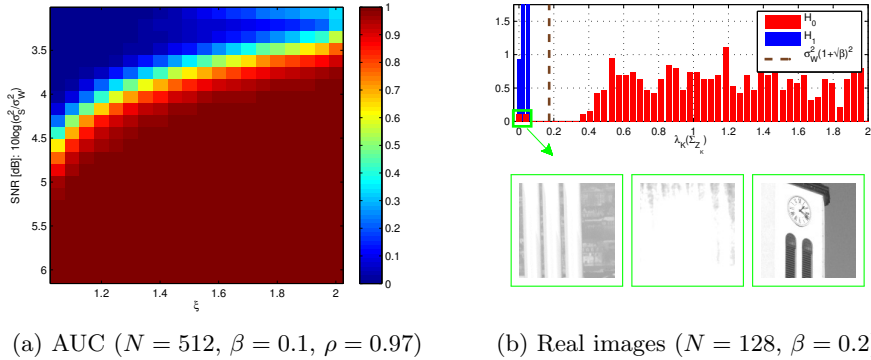


Fig. 4. Evaluation of the detector performance in terms of AUC for different values of $\left(\frac{\sigma_Z^2}{\sigma_W^2}\right)$ and ξ in (a). In (b), histogram of the values of $\lambda_K(\Sigma_{Z_K})$ when considering real images (*top*) and examples of genuine blocks with linear dependencies (*bottom*).

The soundness of this limit with real images is verified by testing a total of 1317 uncompressed images from the Dresden Image Database [2] (only raw images from different Nikon cameras are selected). To avoid the presence of demosaicing traces, the genuine images are constructed by getting access to the output of the camera sensor using the tool `dcraw` and picking always the same green pixel position from each 2×2 Bayer pattern. For performing each full-frame resampling operation we employ the tool `convert` from ImageMagick’s software. As interpolation kernels, we use those described in Table 2. We constrain the set of resampling factors to the interval $[1.05, 2]$ uniformly sampled with step 0.05.

Figure 4(b) shows part of the histogram of $\lambda_K(\Sigma_{Z_K})$ under each hypothesis after processing the central 128×128 block from each genuine and upscaled image by $\xi \in [1.05, 2]$ with the Linear interpolation kernel. As can be observed, the discussed threshold almost perfectly separates both hypotheses, excepting the highlighted cases that correspond to genuine image blocks which present linear dependencies in their content (e.g., flat regions). This is the reason why the test statistic κ defined in [15, Eq. 11] must consider different blocks and identify rank-deficient matrices to circumvent these special cases and also their upscaled versions (which do not follow the model in Sect. 2.1). Hence, when using the test statistic κ with non-demosaiced images, we can configure the detector to work with the above threshold so that it labels an observed image block \mathbf{Z} as upscaled whenever $\kappa < \sigma_W^2(1 + \sqrt{\beta})^2$, and genuine otherwise.

To confirm the good behavior of our detector using this predefined threshold, we conduct the same experiments over all the interpolation kernels, but now processing smaller blocks of size 32×32 , so as to test our detector in realistic conditions where the tampered regions might be small. The performance of the proposed detector is measured in terms of AUC and detection rate at a fixed False Alarm Rate (FAR). For comparison, we consider our previous work in [13] based on subspace decomposition, namely the “SVD-based” detector, and also the state-of-the-art “LP-based” detector proposed in [5] (where LP stands for

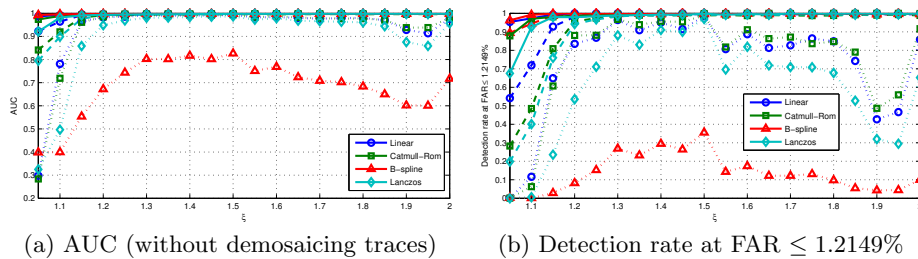


Fig. 5. Evaluation of our detector (*solid lines*) against the SVD-based [13] (*dashed lines*) and the LP-based [5] (*dotted lines*) for image blocks of size 32×32 .

Linear Predictor). We configure our detector to work with submatrices of small aspect ratio $\beta = 0.2812$ (i.e., $K = 9$, $N = 32$), because it makes the separation between the smallest eigenvalue under each hypothesis more evident. For [13] we take $\xi_{\min} = 1.05$ and for [5] we fix a neighborhood of 3 rows/columns.

The AUC for each combination of ξ and interpolation kernel is obtained by applying the three detectors on both the genuine and the correspondingly upscaled images. Given that our detector works with a predefined threshold, we can avoid the use of a training set for computing the detection rates. However, for deriving the empirical thresholds for the other two methods, the database is randomly split in two disjoint sets, where 1/3 of the images are used for training and the remaining ones for testing. The FAR obtained with the application of our method is fixed as a reference for comparing the detection rates of the three detectors. This means that for the detectors in [13] and [5], the training set is used to empirically determine the thresholds that give the same FAR. Finally, using the obtained thresholds, these detectors are applied on the test set of upscaled images to compute the detection rates.

Figure 5(a) shows the obtained AUC for each detector. As expected from the analysis in Figure 4(a), the detection performance of our method improves as the resampling factor increases and, in coherence with the statistical distribution of the eigenvalues shown in Figure 3, the best results are achieved for the interpolation kernels B-spline and Linear since their smallest eigenvalues converge faster to zero than for Catmull-Rom and Lanczos. This improves the separability between upscaled and genuine images and, interestingly, the performance of the SVD-based detector, which also exploits the idea behind the subspace decomposition, resembles that of our method, still showing worse results. The LP-based detector is outperformed by the other two because the spectral density estimator used in [5] becomes inaccurate when dealing with small block sizes.

Figure 5(b) depicts the detection rates at $\text{FAR} \leq 1.2149\%$. Remember that this value of FAR, which turns out to be appealing in practice, results from the application of our detector using the predefined threshold over the genuine images. In all cases, our method always achieves the best detection performance. Furthermore, the results obtained in the range $1.05 \leq \xi \leq 1.2$ are clearly superior to those of the other methods, which is also of interest in practical scenarios, since credible forgeries are typically performed using slight transformations.

5 Conclusions

The eigenvalue characterization of genuine and resampled images has been addressed in this paper using a causal 2D-AR model and drawing on valuable results from RMT. The remarkable agreement between the theoretical distributions and their empirical versions made possible the translation of theoretical bounds to practical solutions for resampling detection. Further research should focus on introducing a suitable model for the demosaicing process, so as to generalize all these theoretical findings to more realistic scenarios.

Acknowledgments. This work is funded by the Agencia Estatal de Investigación (Spain) and the European Regional Development Fund (ERDF) under project WINTER (TEC2016-76409-C2-2-R), and by the Xunta de Galicia and the ERDF under projects Agrupación Estratégica Consolidada de Galicia accreditation 2016-2019 and Red Temática RedTEIC 2017-2018.

References

1. Bai, Z., Zhou, W.: Large sample covariance matrices without independence structures in columns. *Statistica Sinica* 18(2), 425–442 (Apr 2008)
2. Gloe, T., Böhme, R.: The Dresden Image Database for benchmarking digital image forensics. In: ACM SAC. pp. 1584–1590 (Mar 2010)
3. Grenander, U., Szegö, G.: Toeplitz Forms and Their Applications. University of California Press (1958)
4. Kirchner, M.: Fast and reliable resampling detection by spectral analysis of fixed linear predictor residue. In: ACM MM&Sec. pp. 11–20 (Sep 2008)
5. Kirchner, M.: Linear row and column predictors for the analysis of resized images. In: ACM MM&Sec. pp. 13–18 (Sep 2010)
6. Li, Y., Ding, Z.: Blind channel identification based on second order cyclostationary statistics. In: IEEE ICASSP. pp. 81–84 (Apr 1993)
7. Mahdian, B., Saic, S.: Blind authentication using periodic properties of interpolation. *IEEE Trans. on Information Forensics and Security* 3(3), 529–538 (Sep 2008)
8. Marčenko, V.A., Pastur, L.A.: Distribution of eigenvalues for some sets of random matrices. *Mathematics of the USSR-Sbornik* 1(4), 457–483 (Apr 1967)
9. Pfaffel, O., Schlemm, E.: Eigenvalue distribution of large sample covariance matrices of linear processes. *Probability and Mathematical Stat.* 31(2), 313–329 (2011)
10. Popescu, A.C., Farid, H.: Exposing digital forgeries by detecting traces of resampling. *IEEE Transactions on Signal Processing* 53(2), 758–767 (Feb 2005)
11. Tao, T.: *Topics in Random Matrix Theory*. American Math. Soc. (2012)
12. Tulino, A.M., Verdú, S.: Random matrix theory and wireless communications. *Foundations and Trends in Com. and Information Theory* 1(1), 1–182 (Jun 2004)
13. Vázquez-Padín, D., Comesaña, P., Pérez-González, F.: An SVD approach to forensic image resampling detection. In: EUSIPCO. pp. 2067–2071 (Sep 2015)
14. Vázquez-Padín, D., Pérez-González, F., Comesaña Alfaro, P.: Derivation of the asymptotic eigenvalue distribution for causal 2D-AR models under upscaling. [arXiv:1704.05773](https://arxiv.org/abs/1704.05773) [cs.CR] (Apr 2017)
15. Vázquez-Padín, D., Pérez-González, F., Comesaña Alfaro, P.: A random matrix approach to the forensic analysis of upscaled images. *IEEE Transactions on Information Forensics and Security* 12(9), 2115–2130 (Sept 2017)