

Integer Polar Coordinates for Compression

Demba E. Ba and Vivek K Goyal
 Massachusetts Institute of Technology
 {demba,vgoyal}@mit.edu

Abstract— This paper introduces a family of integer-to-integer approximations to the Cartesian-to-polar coordinate transformation and analyzes its application to lossy compression. A high-rate analysis is provided for an encoder that first uniformly scalar quantizes, then transforms to “integer polar coordinates,” and finally separately entropy codes angle and radius. For sources separable in polar coordinates, the performance (at high rate) is shown to match that of entropy-constrained unconstrained polar quantization—where the angular quantization is allowed to depend on the radius. Thus, for sources separable in polar coordinates but not separable in rectangular coordinates—including certain Gaussian scale mixtures—the proposed system performs better than any transform code. Furthermore, unlike unconstrained polar quantization, integer polar coordinates are appropriate for lossless compression of integer-valued vectors. Combination of integer polar coordinates with integer-to-integer transform coding is also discussed.

Index Terms— entropy coding, Gauss circle problem, Gaussian scale mixtures, high-rate quantization theory, integer-to-integer transforms, spherical coordinates, transform coding.

I. INTRODUCTION

TRANSFORM coding has stood the test of time as a method that provides an attractive performance-complexity operating point amongst all lossy compression algorithms. It is central to standards for audio, image, and video coding, though state-of-the-art techniques often combine basic transform coding principles with sophisticated entropy-coding techniques [2].

Transform coding is to use simple (usually scalar) quantization after changing coordinates, and the coordinate change is usually an orthonormal transformation such as a Karhunen-Loève transform (KLT), discrete cosine transform, or discrete wavelet transform. However, polar and spherical coordinates also have a long history of use in source coding theory and practice. For conciseness, this paper discusses only two-dimensional sources and hence only polar coordinates rather than higher-dimensional spherical coordinates.

Polar coordinates are especially attractive when a source distribution is circularly symmetric because in this case the radial (r) and angular (θ) coordinates are independent. However, independence of transform coefficients does not alone make for a good transformation; the shapes and sizes of partition cells are also important, and independent scalar quantization of r and θ gives cell areas and normalized second moments that vary with r . This is illustrated in the middle panel of Fig. 1 and described further in Section II-A. In the entropy-constrained

setting considered in this paper, it is desirable for the partition cells to have equal size. Combined with the restriction to scalar quantization, the best partitioning is a square grid.

This paper develops a way to have the optimal partitioning (squares) while also optimally exploiting circular symmetry in the source distribution. This is achieved with an integer-to-integer (I2I) polar coordinate transformation. For a source with independent radial and angular components, this transformation produces approximately independent components that can subsequently be encoded independently of each other without performance loss. Depending on the source distribution, the performance can thus be better than the best linear transform coder by an arbitrary amount. The analysis is closely related to Gauss’s Circle Problem [3]. Note that although we assume that one starts with a continuous-domain source that is discretized through uniform quantization, one could alternatively directly consider a source supported on \mathbb{Z}^2 that is amenable to polar coordinates. Our integer polar transformations would allow for lossless polar coding of such a source. Prior to this work, this would not have been possible because the centroids of quantization cells in polar quantization schemes do not coincide with \mathbb{Z}^2 (or a scaled version).

The main body of the paper is organized as follows: Section II introduces the relevant background material. In Section III, we develop the I2I polar transform coding scheme and present results of simulations. Concluding remarks are provided in Section IV.

II. BACKGROUND

A. Polar Quantization

Polar quantization is to quantize a source vector (x_1, x_2) via scalar quantization of its polar coordinate representation (r, θ) . Several papers appearing 1978–1986, including [4]–[9], addressed polar quantization of sources with circularly-symmetric pdfs numerically and analytically. These papers mostly focus on the fixed-rate (level-constrained) case and on i.i.d. Gaussian sources. In “strictly” polar quantization (SPQ), r and θ are quantized independently. Wilson [7] introduced “unrestricted” polar quantization (UPQ) in which the number of levels for θ is allowed to depend on r and showed modest improvements over SPQ for i.i.d. Gaussian sources.

In this work we limit our attention to entropy-constrained quantization (ECQ). In high-rate ECQ, the optimal point density is uniform. The improvement of UPQ over SPQ can then be understood from its greater uniformity of point density, as illustrated in Fig. 1.

Asymptotic analyses of entropy-constrained SPQ (ECSPQ) and entropy-constrained UPQ (ECUPQ) were given by Vafin

This work was supported in part by the Texas Instruments Leadership University Consortium Program and Analog Devices, Inc. An abstract of preliminary work was published at IEEE Data Compression Conference 2006 (Snowbird, UT, March 28–30) [1].




Fig. 1. As reproduced from [10], partitions of the input space for ECRQ (left), ECSPQ (middle) and ECUPQ (right)

and Kleijn [10]. Let R denote the rate in bits per vector and let

$$D = E[(x - \hat{x})^2 + (y - \hat{y})^2] = E[|re^{j\theta} - \hat{r}e^{j\hat{\theta}}|^2],$$

where \hat{x}_1 , \hat{x}_2 , \hat{r} and $\hat{\theta}$ are the quantized versions of x_1 , x_2 , r and θ respectively.

- 1) Entropy-Constrained Rectangular Quantization (ECRQ): D is minimized subject to a constraint on $R = H(\hat{x}_1) + H(\hat{x}_2)$. This results in the left-most partition of Fig. 1 and performance

$$D_{ECRQ} = \frac{1}{6} 2^{-(R-h(x_1)-h(x_2))}, \quad (1)$$

where $h(x_i)$ is the differential entropy of x_i , $i = 1, 2$.

- 2) ECSPQ: D is minimized subject to a constraint on $R = H(\hat{r}) + H(\hat{\theta})$. This results in the middle partition of Fig. 1 and performance

$$D_{ECSPQ} = \frac{1}{6} \sigma_r^2 2^{-(R-h(r)-h(\theta))}, \quad (2)$$

where $h(r)$ and $h(\theta)$ are the differential entropies of r and θ respectively. Denoting the probability density of r by $f(r)$, $\sigma_r^2 = \int_0^\infty r^2 f(r) dr$.

- 3) ECUPQ: D is minimized subject to a constraint on $R = H(\hat{r}, \hat{\theta})$. This results in the right-most partition of Fig. 1 and performance

$$D_{ECUPQ} = \frac{1}{6} 2^{-(R-h(r)-h(\theta)-b(r))}, \quad (3)$$

where $b(r) = \int_0^\infty p(r) \log_2(r) dr$.

Each of the three quantizer structures can be implemented with two scalar quantizers. Thus they have low computational complexities compared to unstructured vector quantizers.

B. Integer-to-integer (I2I) Transform Coding

In its conventional formulation, transform coding has performance limited by two characteristics: *scalar* quantization and *scalar* entropy coding. With scalar quantization, to minimize the distortion at high rate the best one can hope for is a quantizer partition into hypercubes; and with scalar entropy coding, one desires independent transform coefficients to minimize the rate. Transform coding has a clean analytical theory for Gaussian sources because these two virtues can be achieved simultaneously. It is well-known that, at high-rate, for jointly Gaussian sources and under some mild conditions, a KLT is optimal [2]. For other sources, ideal cell shape and independence of transform coefficients may be competing, and in fact the optimal transform will generally give neither.




Fig. 2. A generic encoder with a continuous-domain transform $T: \mathbb{R}^N \rightarrow \mathbb{R}^N$ and a discrete-domain transform $U: \Delta\mathbb{Z}^N \rightarrow \Delta\mathbb{Z}^N$. $Q(\cdot)$ is an unbounded uniform scalar quantizer with step size Δ .

Various pitfalls in assuming that the optimality of the KLT extends easily beyond the Gaussian case are described in [11].

For high-rate transform coding of a jointly Gaussian source, it was shown in [12] that the performance of conventional transform coding with transform T can be matched by a system that first quantizes the source and then implements an I2I approximation \hat{T} to T . One can build intuition about I2I transform coding by first considering the per-component mean-squared error (MSE) distortion at high rate of a conventional transform coder:

$$D = \frac{1}{N} \underbrace{\frac{1}{12} \text{trace}(T^{-1}(T^{-1})')}_{\text{shape}} \underbrace{2^{\frac{2}{N} \sum_{i=1}^N h(y_i)}}_{\text{independence}} 2^{-\frac{2}{N}R}. \quad (4)$$

Assuming without loss of generality that $\det T = 1$, the equation above highlights two important facts:

- 1) Orthogonality is good because it minimizes the shape factor over unit-determinant matrices T .
- 2) Independence of transform coefficients is good because it minimizes $\sum_{i=1}^N h(y_i)$.

The idea in I2I transform coding is that one can avoid the shape factor in (4) by first quantizing the source and then implementing an I2I approximation \hat{T} of T obtained through the lifting scheme [12]. Such a strategy achieves independence of transform coefficients with the best possible MSE, while restricting to computationally easy operations, namely *scalar* quantization and *scalar* entropy coding.

Inspired by the I2I framework, this paper takes an expansive view of transform coding in which the encoder performs some signal transformation, uniform scalar quantization, some further signal transformation, and then scalar entropy coding, as illustrated in Fig. 2. Specifically, we go further than in [12] by allowing $U(\cdot)$ to be an integer-to-integer approximation to a (nonlinear) Cartesian-to-polar transformation.

C. Coding of Gaussian Scale Mixtures (GSMs)

A random vector Y is a GSM if it can be expressed as the product of a scalar non-negative random variable z and a zero-mean jointly Gaussian random vector X : $Y = zX$. The scalar random variable z is known as the multiplier, and z and X are independent. GSMs are successfully used to model the statistics of wavelet coefficients of natural images [13].

In [14], the authors derive upper and lower bounds for the rate-distortion functions of GSMs. Ordinary transform coding for GSMs has been previously studied in [15]. In particular, it is shown that a certain class of KLTs is optimal, among orthogonal transforms, for high-rate transform coding. However, by allowing a nonlinear transformation we (asymptotically) obtain the ideal situation of entropy coding independent scalar quantities, rather than simply uncorrelated ones as in [15].

III. I2I POLAR COORDINATES

In this section, we develop an I2I approximation to the Cartesian-to-polar transformation and show its superiority to the conventional approach for transform coding of sources separable in polar coordinates.

A. Designing the Transform

Consider a two-dimensional vector Y with components y_1 and y_2 and identity covariance matrix; these may correspond to the components of a two-dimensional GSM.¹ If Y has independent or nearly independent radial and angular components, we claimed earlier that scalar quantization followed by an I2I approximation to the Cartesian-to-polar transformation (as in Fig. 2) could do better at coding such a source than conventional transform coding. Here, we validate our claim by considering the design of a family of I2I approximations to the Cartesian-to-polar transformation.

Let $(\hat{y}_1, \hat{y}_2) = ([y_1]_\Delta, [y_2]_\Delta)$, where Δ is the step size of the uniform scalar quantizer used in both dimensions and $[\cdot]_\Delta$ denotes rounding to the nearest multiple of Δ . Note that $(\hat{y}_1, \hat{y}_2) \in \Delta\mathbb{Z}^2$, which is the same as saying that $(\hat{y}_1/\Delta, \hat{y}_2/\Delta) \in \mathbb{Z}^2$. The design of our I2I polar transform therefore amounts to a relabeling of the integer coordinate pairs $(\hat{y}_1/\Delta, \hat{y}_2/\Delta)$ into integer pairs (P, A) , where P and A represent our discrete “pseudoradius” and angle variables respectively. In order to avoid loss of information, the mapping from $(\hat{x}_1/\Delta, \hat{x}_2/\Delta)$ to (P, A) should be invertible.

Note that simply quantizing a polar coordinate representation will not work. Let $r = \sqrt{(\hat{y}_1/\Delta)^2 + (\hat{y}_2/\Delta)^2}$ and $\theta = \tan^{-1}[(\hat{y}_2/\Delta)/(\hat{y}_1/\Delta)]$. Further assume $A \in \{0, 1, \dots, M-1\}$ with $M \in \mathbb{Z}^+$. For illustration, one can easily check that

$$(P, A) = \left([r]_1, \left\lfloor \frac{M}{\pi} [\tan^{-1}[(\hat{y}_2/\Delta)/(\hat{y}_1/\Delta)] \bmod \pi] \right\rfloor \right)$$

is not invertible.

Instead of trying to find an approximation to (r, θ) that is invertible on $\Delta\mathbb{Z}^2$, we adopt a sorting approach as follows:


- Divide \mathbb{Z}^2 into M disjoint discrete angle regions by choosing

$$A = \left\lfloor \frac{M}{\pi} [\tan^{-1}(\hat{y}_2/\hat{y}_1) \bmod \pi] \right\rfloor.$$


- Within each angle region (except the zeroth), the $(\hat{y}_1/\Delta, \hat{y}_2/\Delta)$ pair with smallest value of $r = \sqrt{(\hat{y}_1/\Delta)^2 + (\hat{y}_2/\Delta)^2}$ is assigned “pseudoradius” $P = 1$. In the case of the zeroth region, the origin is the point with smallest value of r ; its corresponding “pseudoradius” is $P = 0$. The remaining values of P are assigned by going through each point in increasing order of r and incrementing the value of P at each step; ties are broken by any deterministic rule.

The above construction is illustrated in Fig. 3 for $M = 11$ and 16. By construction, such a mapping is invertible. The main drawback of such an approach is the lack of a simple expression for the mapping from $(\hat{y}_1/\Delta, \hat{y}_2/\Delta)$ to (P, A) .

¹ T in Fig. 2 guarantees that the covariance matrix of Y is a scaled identity matrix.



(a) $M = 11$



(b) $M = 16$

Fig. 3. Visualizing I2I Polar Coordinates. The left panels show cells with equal discrete radius P with the same color. The right panels show regions of constant discrete angle A with constant color.

However, the look-up tables necessary to compute the mapping can easily be generated, both at the transmitter and at the receiver, using only the parameter M . Note in particular that the single parameter M is all that is needed across all rates and sources.

B. High-resolution Analysis

The key feature of the I2I approach is that any invertible transformation of \hat{y} can be used without affecting the distortion. Thus to consider our performance analytically, we only need to approximate the rate, i.e., the entropies $H(P)$ and $H(A)$.

Approximating $H(A)$ is not difficult or interesting. In our examples, A is uniformly distributed on $\{0, 1, \dots, M-1\}$ so $H(A) = \log_2 M$. To evaluate $H(P)$, we relate it to some relevant differential entropy. This is summarized in the following theorem:

Theorem: Let Y be a two-dimensional vector with smooth p.d.f. and components y_1 and y_2 in rectangular coordinates. If the polar coordinate representation (r, θ) has independent components, then

$$H(P) \approx h(r^2) + \log_2(\pi/M) - 2 \log_2 \Delta, \quad (5)$$

with vanishing relative error as $\Delta \rightarrow 0$.²

Proof (sketch): We obtain (5) by relating P to the uniform quantization (with unit step size) of a continuous-valued quantity $\pi r^2 / (M \Delta^2)$.

Consider the cell of area Δ^2 centered at (\hat{y}_1, \hat{y}_2) as shown in the left panel of Fig. 4, and denote the “pseudoradius”

²This seems to be a good approximation even if r and θ are not independent; a precise statement would require M to grow as Δ shrinks. Section III-C contains a discussion on the required rate of growth of M as a function of Δ




Fig. 4. For high-resolution analysis of $H(P)$, a square cell with area Δ^2 (left panel) is replaced by a portion of a ring with area Δ^2 (right panel) as shown. The analysis finds the radius and thickness of the ring as a function of P .

label of the cell by $P = k$. We associate it with a section of a ring, also with area Δ^2 , as shown in the right panel. When (a) Δ is small; (b) r has a smooth density; and (c) r and θ are independent, then the probabilities of the two regions are equal.³ We obtained the desired approximation by equating two estimates of the area of the wedge bounded by an arc going through (\hat{y}_1, \hat{y}_2) . First, because of the point density $1/\Delta^2$, the area should be approximately $k\Delta^2$; we have simply counted points that are assigned smaller P values. On the other hand, the area is $\pi r_k^2/M$, where r_k is the distance from the origin to the point of interest. Equating the areas,

$$k \approx \frac{\pi r_k^2}{M\Delta^2}.$$

Thus, the (integral) P values are approximately the same as what would be obtained by rounding $\pi r^2/(M\Delta^2)$ to the nearest integer.

Example: Suppose y_1 and y_2 are i.i.d., zero-mean Gaussian with variance σ^2 . In polar coordinates we have r^2 exponentially distributed with mean $2\sigma^2$. Using standard differential entropy calculations, one can evaluate as follows:

$$\begin{aligned} h(y_1) + h(y_2) &= \frac{1}{2} \log_2(2\pi e\sigma^2) + \frac{1}{2} \log_2(2\pi e\sigma^2) \\ &= [h(r^2) + \log_2(\pi/M)] + H(A) \\ &= [\log_2(2e\sigma^2) + \log_2(\pi/M)] + \log_2 M, \end{aligned}$$

in bits. Our high-resolution analysis thus implies

$$H(\hat{y}_1) + H(\hat{y}_2) \approx H(P) + H(A),$$

as we would expect.

For sources that are independent in polar coordinates but not in Cartesian coordinates, we will obtain $H(P) + H(A) < H(\hat{y}_1) + H(\hat{y}_2)$.

In what follows, we use (5) to compute the distortion-rate performance of the I2I polar scheme. Let $R = H(P) + H(A)$ be the entropy constraint on the encoding and $D = E[(y_1 - \hat{y}_1)^2 + (y_2 - \hat{y}_2)^2] = \Delta^2/6$ be the mean-squared error (MSE) resulting from scalar quantization of y_1 and y_2 (and consequently from the I2I polar scheme since the

transformation from (\hat{y}_1, \hat{y}_2) to (P, A) is lossless). Then for large R (high-rate),

$$\begin{aligned} R &= H(P) + H(A) \\ &= h(r^2) + h(\theta) - 2 \log_2 \Delta - 1 \\ &= h(r) + E[\log_2 r] + h(\theta) - \log_2 \Delta^2. \end{aligned}$$

Rearranging the above equation gives the following expression for D as a function of R :

$$D = \frac{1}{6} 2^{-(R-h(r)-h(\theta)-b(r))}, \quad (6)$$

where $b(r) = \int_0^\infty p(r) \log_2(r) dr$.

Note that the high-rate performance of the I2I polar scheme does not depend on the number of angle regions M , as can be seen from (6). Moreover, **this is the same distortion-rate performance as that achieved by ECUPQ.**

C. Bound for the difference $|k - \frac{\pi r^2}{M\Delta^2}|$

In the preceding high-resolution analysis, we argued that the (integral) P values are approximately the same as what one would obtain by rounding $\pi r^2/M\Delta^2$ to the nearest integer. In actuality, our assertion is only true if the above difference increases *slowly* as a function of r . This would mean (as one would like) that the number of ties to be broken when designing the I2I polar transformation does not grow incommensurately with r .

The problem of bounding the difference $|k - \pi r^2/M\Delta^2|$ is intimately related to that of counting the number of lattice points in a circle of radius r , a classical problem which goes back to Gauss. Let N_r be the number of lattice points in a circle of radius r . According to [3],

$$|N_r - \pi r^2| \leq O(r^{\lambda+\epsilon}), \quad (7)$$

where $\lambda \approx 1/2$ and $\epsilon > 0$. A great deal of work has gone into determining the exact value of λ ; however, for the sake of our argument, Hardy's careful conjecture [16] that $\lambda = 1/2$ will suffice as well as provide better insight than some of the more elaborate results.

Using (7) and substituting $1/2$ for λ , one can obtain a bound for $|k - \pi r^2/M\Delta^2|$ as a function of Δ and M . This is achieved by realizing that, at high rate, the number of squares of area Δ^2 in a circle of radius r is the same as the number of lattice points in a circle of radius r/Δ . Therefore,

$$\left| kM - \frac{\pi r^2}{\Delta^2} \right| \leq O\left(\left(\frac{r}{\Delta}\right)^{1/2+\epsilon}\right), \quad (8)$$

which leads us to the bound

$$\left| k - \frac{\pi r^2}{M\Delta^2} \right| \leq O\left(\frac{r^{1/2+\epsilon}}{M\Delta^{1/2+\epsilon}}\right). \quad (9)$$

Equation 9 states that, for bounded r , the number of ties to be broken when designing the I2I polar transformation only grows as $r^{1/2}$. One also requires that the number M of discrete angle regions grow as $1/\Delta^{1/2}$.

³Requirements (b) and (c) could be replaced with smoothness of the joint density of r and θ if the angular extent were small.

TABLE I

PERFORMANCE COMPARISON (RATIO OF DISTORTIONS AT HIGH RATE)
FOR VARIOUS ENTROPY CODING TECHNIQUES APPLIED TO TWO GSMs.

Prior on z	RQ/I2I	UPQ/I2I	SPQ/I2I	J-RQ/I2I
$z \sim E(\cdot)$	1.048	1	1.79	1
$z \sim R(\cdot)$	1.23	1	3.36	1

D. Results

Table 1 compares the performance of entropy-coded I2I (ECI2I) polar transform coding to that of the various entropy-coded quantization techniques introduced in Section II-A, as well as J-RQ, which refers to rectangular quantization (RQ) followed by *joint* entropy coding. The criterion used is the ratio of distortions at high rate. The comparison is based on two examples of GSMs, one where the multiplier z is exponentially-distributed (denoted $E(\cdot)$) and another where it has a Rayleigh distribution (denoted $R(\cdot)$).

The first column of Table I shows that ECI2I outperforms ECRQ for both of the sources considered. While the improvement is relatively small, it should be noted that it is for a source for which no improvement was thought possible with transform coding [15]. The fourth column of the table shows that, in both cases, ECI2I is asymptotically optimal: its performance matches that obtained by scalar quantization followed by joint entropy coding. The second column of Table I reflects the fact, at high rate, the performance of ECI2I and that of ECUPQ are the same. The third column of the table shows that both ECI2I and ECUPQ significantly outperform ECSPQ for both of the GSMs considered. Intuition behind such a result can be built by taking a closer look at Fig. 1. Indeed, both ECI2I and ECUPQ aim at maintaining quantization partition cells that are as close as possible to squares (in ECI2I, the quantization partition cells are exact squares), while the cells that result from ECSPQ become further away from squares as one gets further away from the origin. Therefore, the outermost cells resulting from ECSPQ have higher normalized second moments than those induced by ECI2I and ECUPQ, which translates into higher distortion (at the same rate). The gain is larger with the Rayleigh prior on z (3.36 vs. 1.79) because the p.d.f. of the radius decays slower, which means that outermost ECSPQ quantization cells have more probability mass associated with them than in the case of the exponential prior on z .

IV. CONCLUSION

In conventional transform coding, the desire to preserve the shape of quantization partition cells as well as have the luxury of encoding independent quantities are competing goals. Integer-to-integer (I2I) transform coding [12] is a method that aims at simultaneously achieving both of these goals by allowing quantization to occur first and then implementing

an invertible integer-to-integer approximation to a possibly nonlinear transformation that will yield independent components. Inspired by the I2I framework, this paper introduces the generic encoder of Fig. 2 for nonlinear transform coding of polar-separable sources. Specifically, we design an I2I approximation to the Cartesian-to-polar transformation and analytically show its superiority to the conventional approach for high-rate transform coding of polar-separable sources. Our analysis is supported by simulation results on GSMs, a class of sources for which no improvement was previously thought possible under the conventional transform coding paradigm [15].

ACKNOWLEDGMENT

The authors would like to thank the reviewers for making them aware of the related work in [5], [8], [9] and for comments that improved the focus and organization of the paper.

REFERENCES

- [1] D. E. Ba and V. K Goyal, "Nonlinear transform coding: Polar coordinates revisited," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 2006, p. 438, IEEE Computer Society.
- [2] V. K Goyal, "Theoretical foundations of transform coding," *IEEE Sig. Proc. Mag.*, vol. 18, no. 5, pp. 9–21, Sept. 2001.
- [3] P. M. Bleher, Z. Cheng, F. J. Dyson, and J. L. Lebowitz, "Distribution of the error term for the number of lattice points inside a shifted circle," *Commun. Math. Phys.*, vol. 154, no. 3, pp. 433–469, 1993.
- [4] N. C. Gallagher, Jr., "Quantizing schemes for the discrete Fourier transform of a random time-series," *IEEE Trans. Inform. Theory*, vol. IT-24, no. 2, pp. 156–162, Mar. 1978.
- [5] W. A. Pearlman, "Polar quantization of a complex Gaussian random variable," *IEEE Trans. Comm.*, vol. COM-27, no. 6, pp. 892–899, June 1979.
- [6] J. A. Bucklew and N. C. Gallagher, Jr., "Quantization schemes for bivariate Gaussian random variables," *IEEE Trans. Inform. Theory*, vol. IT-25, no. 5, pp. 537–543, Sept. 1979.
- [7] S. G. Wilson, "Magnitude/phase quantization of independent Gaussian variates," *IEEE Trans. Comm.*, vol. COM-28, no. 11, pp. 1924–1929, Nov. 1980.
- [8] P. F. Swaszek and J. B. Thomas, "Optimal circularly symmetric quantizers," *J. Franklin Inst.*, vol. 313, no. 6, pp. 373–384, June 1982.
- [9] P. F. Swaszek and T. W. Ku, "Asymptotic performance of unrestricted polar quantizers," *IEEE Trans. Inform. Theory*, vol. IT-32, no. 2, pp. 330–331, Mar. 1986.
- [10] R. Vafin and W. B. Kleijn, "Entropy-constrained polar quantization and its application to audio coding," *IEEE Trans. Speech Audio Proc.*, vol. 13, no. 2, pp. 220–232, Mar. 2005.
- [11] M. Effros, H. Feng, and K. Zeger, "Suboptimality of the Karhunen-Loève transform for transform coding," *IEEE Trans. Inform. Theory*, vol. 50, no. 8, pp. 1605–1619, Aug. 2004.
- [12] V. K Goyal, "Transform coding with integer-to-integer transforms," *IEEE Trans. Inform. Theory*, vol. 46, no. 2, pp. 465–473, Mar. 2000.
- [13] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, "Image denoising using scale mixtures of Gaussians in the wavelet domain," *IEEE Trans. Image Proc.*, vol. 12, no. 11, pp. 1338–1351, Nov. 2003.
- [14] H. M. Leung and S. Cambanis, "On the rate distortion function of spherically invariant vectors and sequences," *IEEE Trans. Inform. Theory*, vol. 24, no. 3, pp. 367–373, May 1978.
- [15] S. Jana and P. Moulin, "Optimality of KLT for high-rate transform coding of Gaussian vector-scale mixtures: Application to reconstruction, estimation and classification," *IEEE Trans. Inform. Theory*, vol. 52, no. 9, pp. 4049–4067, Sept. 2006.
- [16] G. H. Hardy, "The average order of the arithmetical functions $p(x)$ and $\delta(x)$," *Proc. London Math. Soc.*, vol. 15, no. 2, pp. 192–213, 1916.