RATE-DISTORTION TRADE-OFFS IN ACQUISITION OF SIGNAL PARAMETERS

Miguel R. D. Rodrigues[†], *Nikos Deligiannis*[‡], *Lifeng Lai*^{*}, and Yonina C. Eldar^{**}

[†] Dept. Electronic and Electrical Engineering, University College London, United Kingdom
 [‡]Dept. Electronics and Informatics, Vrije Universiteit Brussel, Belgium and iMinds vzw, Ghent, Belgium
 ^{*}Dept. Electrical and Computer Engineering, University of California, Davis, CA, USA
 ^{**}Dept. Electrical Engineering, Technion – Israel Institute of Technology, Israel

ABSTRACT

We consider problems where one wishes to represent a parameter associated with a signal source – subject to a certain rate and distortion – based on the observation of a number of realizations of the source signal. By reducing these indirect vector quantization problems to a standard vector quantization one, we provide a bound to the fundamental interplay between the rate and distortion in the large-rate setting. We specialize this characterization to two particular quantization scenarios: *i*) the representation of the mean of a multivariate Gaussian source; and *ii*) the representation of the eigen-spectrum of a multivariate Gaussian source. Numerical results compare our quantization approach to an approach where one recovers the parameters from the representation of the source signals itself: in addition to revealing that the characterization is sharp in the large-rate setting, the results also show that our approach offers considerable gains.

Index Terms— Signal Acquisition, Signal Parameters Acquisition, Rate-Distortion, Vector Quantization

1. INTRODUCTION

Modern signal processing systems – which manipulate digital versions of real-world analog signals – rely on analog-to-digital and digital-to-analog converters to convert the signal from the analog to the digital domain and vice-versa. The fundamental principles associated with these signal acquisition and reconstruction operations are laid out by the two seminal theories of sampling and rate-distortion. The Shannon-Nyquist sampling paradigm asserts that it is possible to reconstruct a continuous-time signal from its discrete-time version provided that the signal sampling rate is at least twice the signal bandwidth [1], whereas rate-distortion theory quantifies the trade-off between rate and distortion associated with the digital representation of the analog signal samples [2]. The interplay between sampling frequency, rate and distortion has been studied in recent works [3].

Recent years have also witnessed the introduction of other signal acquisition and reconstruction paradigms. Of particular relevance, compressive sensing and its variants [4], [5], [6], [7], [8], [9] demonstrate that it is possible to recover a signal from a sub-Nyquist sampled version, provided that the signal admits a sparse representation in an orthonormal dictionary (e.g. wavelet basis) or overcomplete dictionary. See also [10].

There are however various practical applications where one does not wish to capture the signal itself but rather a certain signal parameter. For example, one is often interested in the second-order statistics of an underlying stochastic process in applications such as communications, audio and speech processing, passive sonar and radar, radioastronomy or seismology. This has led to a renaissance of interest in various sampling approaches, such as covariance sensing methods, that exploit structural information in the statistical domain to capture very efficiently the relevant signal parameters [11] (see also [12], [13], [14]). In particular, it has been shown that covariance sensing methods can lead to substantial sampling frequency gains in relation to schemes where the signal parameter is derived from a conventionally sampled signal. However, these approaches assume that the signal samples are captured perfectly, so that quantization is not taken into account. Here, we wish to understand the fundamental trade-off between quantization rate and distortion associated with these emerging approaches.

This paper introduces a framework to study rate-distortion tradeoffs associated with such problems, where we are only interested in certain parameters of the signal distribution given signal samples. The contributions include a bound to the distortion-rate function (Theorem 2) and its specialization to two particular problems (Theorems 3 and 4). The parameter-tailored rate-distortion problem is formulated in Section 2. Section 3 introduces the bound to the distortion-rate function (applicable to the large-rate setting) that relies on the reduction of the original indirect rate-distortion (quantization) problem onto a surrogate direct one. Section 4 introduces specializations of the bound to two problems: representation of the mean and representation of the eigen-spectrum of a multivariate Gaussian signal given a number of signal realizations. Section 5 presents a number of numerical results that showcase that parametertailored signal acquisition can lead to substantial gains over other approaches. Finally, concluding remarks are drawn in Section 6.

2. RATE-DISTORTION PROBLEM

Figure 1 illustrates the setting, where one wishes to represent a parameter associated with a signal source subject to a certain rate and distortion, based on the observation of a number of realizations of the source signal. This (indirect) rate-distortion setting differs from the traditional one because the quantizer does not have direct access to the parameter of interest itself but rather to a manifestation of the parameter.

In particular, the encoder

$$f(\cdot): \mathbb{R}^m \times \dots \times \mathbb{R}^m \to \mathcal{M} = \{1, \dots, M\}$$
(1)

maps the signal realizations $x_1, \ldots, x_n \in \mathbb{R}^m \times \cdots \times \mathbb{R}^m$ onto an index $j \in \mathcal{M}$ and the decoder

$$g(\cdot): \mathcal{M} = \{1, \dots, M\} \to \mathbb{R}^k \tag{2}$$

maps the index $j \in \mathcal{M}$ onto an estimate $\hat{\theta} \in \mathbb{R}^k$ of the true underlying parameter $\theta \in \mathbb{R}^k$. The operation of this encoder-decoder pair

$$\theta \longrightarrow x_{1,x_{2},...,x_{n}} \longrightarrow \begin{array}{c} \text{Encoder} \\ f(\cdot) \end{array} \xrightarrow{j \in \{1,...,M\}} \begin{array}{c} \text{Decoder} \\ g(\cdot) \end{array}$$

Fig. 1. The rate-distortion setting.

will be described via an M-level (vector) quantization function as follows:

$$Q^{M}(x_{1},\ldots,x_{n}) = g(f(x_{1},\ldots,x_{n})) = \hat{\theta}_{i}, x_{1},\ldots,x_{n} \in S_{i}$$
(3)

where $\hat{\theta}_1, \ldots, \hat{\theta}_M$ is the reproduction alphabet and S_1, \ldots, S_M is a partition of $\mathbb{R}^m \times \cdots \times \mathbb{R}^m$, i.e. $\cup_i S_i = \mathbb{R}^m \times \cdots \times \mathbb{R}^m$ and $S_i \cap S_j = \emptyset \ (i \neq j).$

Our goal is to provide insight into the fundamental tradeoff between the number of quantization levels M and the quantization distortion D associated with an optimal quantizer, i.e. the quantizer that minimizes D for a certain M. We assume a squared-error distortion metric, so that

$$D = \mathbb{E}\left\{d(\theta, \hat{\theta})\right\} = \mathbb{E}\left\{\mathsf{tr}\left((\theta - \hat{\theta})(\theta - \hat{\theta})^t\right)\right\}$$
(4)

We also assume that the signal parameter $\theta \in \mathbb{R}^k$ is drawn according to the distribution $p_{\theta}(\theta)$ and that the signal realizations $x_1, \ldots, x_n \in \mathbb{R}^m \times \cdots \times \mathbb{R}^m$ conditioned on the signal parameter $\theta \in \mathbb{R}^k$ are drawn according to the distribution $p_{x_1,\ldots,x_n|\theta}(x_1,\ldots,x_n|\theta)$. Finaly, the quantizer knows $p_{\theta}(\theta)$ and $p_{x_1,\ldots,x_n|\theta}(x_1,\ldots,x_n|\theta)$ (but not the θ itself).

We note in passing that a straightforward approach to address these class of problems would involve optimal encoding and decoding of the signal samples, and then estimation of the parameter from the decoded samples. It will be shown that our rate-distortion approach can lead to considerably gains in relation to this other conventional approach.

3. RATE-DISTORTION TRADEOFF: ASYMPTOTIC LARGE-RATE CHARACTERIZATION

We now provide a bound to the fundamental tradeoff between the number of quantization levels and the quantization distortion that is applicable to the asymptotic regime where the number of quantization levels approaches infinity $(M \to \infty)$.

Our approach relies on reduction of the original indirect quantization problem into a standard vector quantization problem, by leveraging a well-known result that illuminates the structure of the optimal quantizer associated with squared-error distortion metrics for indirect rate-distortion problems [15] (see also [16], [17]).

Theorem 1. [15] Let

$$\tilde{\theta} = \tilde{\theta}(x_1, \dots, x_n) = E\{\theta | x_1, \dots, x_n\}$$
(5)

be the conditional mean estimate of the parameter given the signal realizations. Let also

$$\tilde{D}^* = E\left\{d(\theta, \tilde{\theta})\right\} + \min_{\tilde{Q}^M(\cdot)} E\left\{d(\tilde{\theta}, \tilde{Q}^M(\tilde{\theta}))\right\}$$
(6)

and

$$D^* = \min_{Q^M(\cdot)} E\left\{d(\theta, Q^M(x_1, \dots, x_n))\right\}$$
(7)

where $\tilde{Q}^{M}(\tilde{\theta})$ is an *M*-level quantizer that operates on $\tilde{\theta} \in \mathbb{R}^{k}$ and $Q^{M}(x_{1}, \ldots, x_{n})$ is the original *M*-level quantizer that operates on $x_{1}, \ldots, x_{n} \in \mathbb{R}^{m} \times \cdots \times \mathbb{R}^{m}$. Then,

$$\tilde{D}^* = D^* \tag{8}$$

Theorem 1 – which is valid for any M, n, m and k – states that optimal quantization in the setting of Fig. 1 decomposes into two operations:

- An estimation operation that yields the optimal mean-squared error estimate (MSE) of the parameter given the signal realizations;
- A quantization operation that yields the optimal quantization of the MSE estimate of the parameter given the signal realizations.

Theorem 1 also states that the minimal distortion between the source parameter and the source parameter representation decomposes into the sum of two components: *i*) The first component of the distortion corresponds to the estimation error (i.e. minimum mean-squared error (MMSE) associated with the estimation of the source parameter given the signal realizations); *ii*) The second component of the distortion corresponds to the quantization error. A similar decomposition was also identified in [3].

The characterization of the structure of the optimal quantizer embodied in Theorem 1 leads to bound to the fundamental ratedistortion tradeoff applicable to large-rate settings.

Theorem 2. Let

$$\tilde{\theta} = \tilde{\theta}(x_1, \dots, x_n) = E\{\theta | x_1, \dots, x_n\}$$
(9)

be the conditional mean estimate of the parameter given the signal realizations and $p_{\tilde{\theta}}(\tilde{\theta})$ its distribution. Then, in the regime of large-rate,

$$D^* \gtrsim E\left\{ \operatorname{tr}\left((\theta - \tilde{\theta})(\theta - \tilde{\theta})^t\right) \right\} + \frac{k}{k+2} C_k^{-2/k} M^{-2/k} \left[\int p_{\tilde{\theta}}(\tilde{\theta})^{\frac{k}{k+2}} d\tilde{\theta} \right]^{\frac{k+2}{k}}$$
(10)

where $C_k = (2\Gamma(1/2)^k)/(k\Gamma(k/2)).$

Proof. This theorem follows by adapting the argument in [18]. It is omitted for brevity. $\hfill \Box$

Theorem 2 captures the interplay between the number of quantization levels and distortion in the asymptotic regime of large M (it is applicable however for any k, m and n). The estimation error – first term on the right hand side of (10) – does not depend on the number of quantization levels M, so it acts as a fundamental distortion *plateau*. In contrast, the quantization error depends on the number of quantization levels M; it also depends on the distribution of the conditional mean estimate of the parameter in *lieu* of the distribution of the quantization error approach zero as $n \to \infty$ and $M \to \infty$, respectively.

4. EXAMPLES

We now specialize the bound to the distortion-rate function in Theorem 2 to concrete scenarios in order to expose fundamental limitations in parameter-tailored rate-distortion problems.

4.1. Representation of the Mean of a Multivariate Gaussian Source

Consider a scenario where the signal realizations are drawn according to a multivariate Gaussian distribution:

$$p(x_1, \dots, x_n | \mu, \Sigma) = \prod_{i=1}^k \mathcal{N}(x_i; \mu, \Sigma)$$
(11)

where $\mu \in \mathbb{R}^k$ and $\Sigma \in \mathbb{R}^{k \times k}$ represent the mean and covariance of the multivariate Gaussian distribution. We assume that the mean μ is unknown to the quantizer but the covariance Σ is known. We also assume that the mean follows a multivariate Gaussian distribution:¹

$$p(\mu) = \mathcal{N}(\mu; 0, \Sigma_{\mu}) \tag{12}$$

with zero-mean and covariance $\Sigma_{\mu} \in \mathbb{R}^{k \times k}$.

The following Theorem provides a (large-rate) bound to the fundamental interplay between rate and distortion under the squarederror distortion metric given by:

$$d(\mu,\hat{\mu}) = \operatorname{tr}\left(\left(\mu - \hat{\mu}\right) \cdot \left(\mu - \hat{\mu}\right)^{t}\right).$$
(13)

Theorem 3. A (high-rate) bound to the distortion-rate function associated with the representation of the mean of the Gaussian source is given by:

$$D^* \gtrsim \frac{1}{n} \operatorname{tr} \left(\Sigma^{-1} + \frac{1}{n} \Sigma^{-1}_{\mu} \right)^{-1} + 2\pi C_k^{-2/k} M^{-2/k} \left(\frac{k+2}{k} \right)^{\frac{k}{2}} \times \det \left(\Sigma^{-1}_{\mu} + \frac{1}{n} \Sigma^{-1}_{\mu} \Sigma \Sigma^{-1}_{\mu} \right)^{-\frac{1}{k}}.$$
 (14)

Proof. In view of (11) and (12) it follows immediately that

$$p(\mu|x_1,\ldots,x_n) = \mathcal{N}\left(\mu',\Sigma'\right) \tag{15}$$

where $\mu' = (\Sigma_{\mu}^{-1} + n\Sigma^{-1})^{-1} \Sigma^{-1} \sum_{i=1}^{n} x_i$ and $\Sigma' = (\Sigma_{\mu}^{-1} + n\Sigma^{-1}) \overline{w}$ here $B(\cdot, \cdot)$ represents the beta function. Then, the optimal minimum MSE estimate of the parameter given Proof. In view of (11) and (18) it follows the signal realizations is given by

$$\tilde{\mu} = \sum_{i=1}^{n} \left(\Sigma_{\mu}^{-1} + n \Sigma^{-1} \right)^{-1} \Sigma^{-1} x_i$$
(16)

and its distribution is given by

$$p(\tilde{\mu}) = \mathcal{N}\left(\tilde{\mu}; 0, \left(\Sigma_{\mu}^{-1} + \frac{1}{n}\Sigma_{\mu}^{-1}\Sigma\Sigma_{\mu}^{-1}\right)^{-1}\right).$$
(17)

The result then follows immediately by using (12) and (17) in the characterization in (10).

The first and second terms on the right hand side of (14) correspond to the estimation and quantization errors, respectively. The estimation error approaches zero as $n \to \infty$ and the error associated with the quantization of the optimal minimum MSE estimate of the parameter approaches the error associated with the (optimal) quantization of the parameter itself as $n \to \infty$. This is due to the fact that - in the regime where n is large - it is possible to obtain a nearly perfect estimate of the underlying source parameter.

The quantization error - in addition to depending on the number of quantization levels - also depends directly on the geometric mean of the eigenvalues of the covariance of the optimal minimum MSE estimate (a proxy to its variability): the larger this geometric mean the larger the quantization error (for fixed M and n).

4.2. Representation of the Eigen-Spectrum of a Multi-Variate **Gaussian Source**

Consider now the case where the signals are again drawn according to (11) but we now assume that the mean μ is known but the covariance $\Sigma = U\Lambda U^t$ is such that U is known and Λ is unknown. Here, $U \in \mathbb{R}^{k \times k}$ is an orthogonal matrix that contains the eigenvectors of Σ and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_k) \in \mathbb{R}^{k \times k}$ is a diagonal matrix that contains the eigenvalues of Σ .² We also assume that the eigenvalues follow the distribution: ³

$$p(\lambda_1, \dots, \lambda_k) = \prod_{i=1}^k \mathsf{IG}(\lambda_i; \alpha_i, \beta_i)$$
(18)

with $\alpha_i > 2, \forall i$, where IG $(\cdot; \alpha, \beta)$ represents the inverse-gamma distribution with shape parameter α and scale parameter β .

The following Theorem provides a (large-rate) bound to the fundamental interplay between rate and distortion under the squarederror distortion metric:

$$d\left(\Lambda,\hat{\Lambda}\right) = \operatorname{tr}\left(\left(\Lambda-\hat{\Lambda}\right)\cdot\left(\Lambda-\hat{\Lambda}\right)^{t}\right) = \sum_{i=1}^{k} \left(\lambda_{i}-\tilde{\lambda}_{i}\right)^{2}.$$
 (19)

Theorem 4. A (high-rate) bound to the distortion-rate function associated with the representation of the eigen-spectrum of the Gaussian source is given by:

$$D^* \gtrsim \sum_{i=1}^{k} \frac{\beta_i}{\alpha_i - 1 + n/2} \cdot \frac{\beta_i}{(\alpha_i - 1)(\alpha_i - 2)} + \frac{k}{k+2} C_k^{-2/k} M^{-2/k} \times \prod_{i=1}^{k} \frac{\beta_i^2}{(\alpha_i - 1 + n/2)^2} \cdot \frac{B^{\frac{k+2}{k}}\left(\frac{k\alpha_i - 2}{k+2}, \frac{kn/2 + 2}{k+2}\right)}{B(\alpha_i, n/2)}.$$
 (20)

Proof. In view of (11) and (18) it follows immediately that

$$p(\lambda_1, \dots, \lambda_k | x_1, \dots, x_n) = \prod_{i=1}^k \mathsf{IG}\left(\lambda_k; \alpha'_k, \beta'_k\right)$$
(21)

where $\alpha'_{k} = \alpha_{k} + \frac{n}{2}$ and $\beta'_{k} = \beta_{k} + \frac{1}{2} e_{k}^{t} U^{t} \left(\sum_{i=1}^{n} (x_{i} - \mu) (x_{i} - \mu)^{t} \right) U e_{k}$. Then, the optimal minimum MSE estimate of the unknown parameters are given by

$$\left(\tilde{\lambda}_1, \dots, \tilde{\lambda}_k\right) = \left(\frac{\beta_1'}{\alpha_1' - 1}, \dots, \frac{\beta_k'}{\alpha_k' - 1}\right)$$
 (22)

and its distribution is given by

$$p\left(\tilde{\lambda}_{1},\ldots,\tilde{\lambda}_{k}\right) = \prod_{i=1}^{k} \frac{1}{B(\alpha_{i},n/2)} \times \frac{1}{\delta_{i}} \times \left(\frac{\tilde{\lambda}_{i}-\delta_{i}}{\delta_{i}}\right)^{\frac{n}{2}-1} \times \left(1+\frac{\tilde{\lambda}_{i}-\delta_{i}}{\delta_{i}}\right)^{-\frac{n}{2}-\alpha_{i}}$$
(23)

with $\tilde{\lambda}_1 \geq \delta_1, \ldots, \tilde{\lambda}_k \geq \delta_k$, where $\delta_i = \beta_i/(\alpha_i + n/2 - 1)$. The result follows by using (18) and (23) in the characterization in (10).

¹We choose for analytical convenience the distribution of the parameter to be the conjugate prior to a Gaussian distribution with unknown mean and known covariance.

²Note that this applies to scenarios where one has additional structural information about the covariance matrix. For example, if the covariance matrix is known to be circulant, then U is a DFT matrix.

³We again choose for analytical convenience the distribution of the parameter to be the conjugate prior to a Gaussian distribution with known mean, known eigenvectors and unknown eigenvalues.



Fig. 2. Evaluation of the encoding of the (a) mean and the (b) eigenspectrum of a multivariate Gaussian distribution.

The first and second terms on the right hand side of (20) also correspond to the estimation and quantization errors, respectively. It is clear that the estimation error tends to zero as the number of signal realizations tends to infinity and – in view of the fact that $B(x, y) \sim \Gamma(y)x^{-y}$ for x large and y fixed – the quantization error tends to:

$$\frac{k}{k+2}C_k^{-2/k}M^{-2/k}\prod_{i=1}^k\beta_i^2\frac{\Gamma\left(\frac{k\alpha_i-2}{k+2}\right)}{\Gamma(\alpha_i)}\left(\frac{k}{k+2}\right)^{-\frac{k\alpha_i-2}{k}}$$
(24)

as the number of signal realizations tends to infinity, i.e. the quantization error approaches the error associated with the direct quantization of the source eigen-spectrum itself. In addition, the quantization error increases with the increase of the scale parameter (for a fixed M and n), in view of the fact that this parameter measures the variability of the inverse-gamma random variable (recall that the variance of an inverse-gamma random variable grows quadratically with the scale parameter β for fixed α).

5. RESULTS

We conduct experiments to compare our approach, where one encodes the MMSE estimate of the parameter using $\log_2 M$ bits, against the approach where the encoder encodes n samples into $n \log_2 M$ bits and the decoder recovers the parameter from the

quantized samples. In our experiments we fix the dimensionality of the source samples to k = 2 and consider n = 100 or n = 1000 signal samples.

In the first set of experiments, we consider the encoding of the mean of a multivariate Gaussian distribution. We generate μ values according to (12), where $\Sigma_{\mu} = J \times J^T$ and $J \in (0, 1)^{k \times k}$ is a matrix of which the entries are uniformly-distributed random numbers in (0,1). Per μ value, we generate n source samples (x_1,\ldots,x_n) according to (11), where Σ is drawn similarly to Σ_{μ} . Adhering to our approach, we derive the MMSE estimate $\tilde{\mu}$ of the mean using (16) and encode it using an M level vector quantizer, trained using the LBG algorithm [19] on the distribution given in (17). In the alternative approach, each source sample is first quantized with an M-level vector quantizer⁴ that is trained for samples following (11). At the decoder the parameter is estimated via the sample mean of the recovered data samples. It is worth mentioning that, in the second approach, one needs to train a quantizer per trial (corresponding to each mean value μ), meaning that, in practice, the obtained codebook needs to be communicated to the decoder (this extra rate is even not taken into account in our results).

In the second set of experiments, we consider the encoding of the eigen-spectrum of a multivariance Gaussian source. The setup is similar to the previous one: we draw a k-dimensional eigenvalue vector according to (18), where U is equal to the identity matrix and $\alpha_i \in (5, 6)$ and $\beta_i \in (0.33, 0.5)$, $k = \{1, 2\}$, drawn uniformly at random. Given n generated source samples, we encode either the MMSE estimate [calculated via (22)] or the samples.

Figure 2 presents the performance of the two approaches in terms of the rate (in bits) versus the MSE distortion—presented in a logarithmic scale—, both averaged over 200 Monte-Carlo trials. It can be observed that quantization of the underlying parameter can lead to substantially lower distortions than quantization of the samples themselves. It can also be observed that the derived characterization is sharp in the high-rate regime, thereby providing the means to gauge the rate-distortion tradeoff associated with our approach.

6. CONCLUSION

Emerging problems in signal acquisition involve the representation of parameters associated with a signal source based on the observation of a number of realizations of the source signal, subject to a certain rate and distortion.

This paper provides a bound to the fundamental interplay between rate and distortion for this indirect quantization problems that is applicable to the large-rate regime. This characterization—which is also specialized to two concrete settings—states that the distortion consists of two components: one component corresponds to the error associated with the estimation of the parameter from the samples and the other component is associated with the error associated with the quantization of the parameter estimate itself.

Numerical results demonstrate that this quantization approach can lead to marked gains over an alternative approach where the parameter is estimated directly from the quantized signal samples; moreover, the results also reveal that the characterizations are sharp in the large-rate regime, providing the means to gauge a rate to achieve a certain target distortion for this new class of signal acquisition problems.

⁴The source sample quantizer is different from the quantizer for the MMSE estimate. The training is performed using the LGB algorithm [19].

7. REFERENCES

- [1] Y. C. Eldar. Sampling Theory: Beyond Bandlimited Systems. Cambridge University Press, 2015
- [2] T. M. Cover and J. A. Thomas. Elements of Information Theory, 2nd Edition. Wiley, 2006.
- [3] A. Kipnis, A. J. Goldsmith, Y. C. Eldar and T. Weissman, "Distortion-rate function of sub-Nyquist sampled Gaussian sources," *IEEE Transactions on Information Theory*, vol. 62, no. 1, pp. 401–429, 2016.
- [4] E. Candes and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5406-5425, 2006.
- [5] D. Donoho, "Compressed sensing," *IEEE Transactions on In*formation Theory, vol. 52, no. 4, pp. 1289-1306, 2006.
- [6] R. G. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, pp. 253263, 2008.
- [7] M. Mishali and Y. C. Eldar, "From theory to practice: Sub-Nyquist sampling of sparse wideband analog signals," *IEEE Journal of Selected Topics on Signal Processing*, vol. 4, no. 2, pp. 375-391, Apr. 2010
- [8] M. Mishali and Y. C. Eldar, "Sub-Nyquist sampling: Bridging theory and practice", *IEEE Signal Processing Magazine*, vol. 28, no. 6, pp. 98-124, 2011.
- [9] Y. C. Eldar and G. Kutyniok. *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2012.
- [10] W. Dai, H. V. Pham and O. Milenkovic. Quantized compressive sensing. arXiv:0901.0749v2, 2009.
- [11] D. Romero, D. D. Ariananda, Z. Tian, and G. Leus, "Compressive covariance sensing: Structure-based compressive sensing beyond sparsity," *IEEE Signal Processing Magazine*, vol. 33, no. 1, pp. 78–93, 2016.
- [12] D. Cohen and Y. C. Eldar, "Sub-Nyquist sampling for power spectrum sensing in cognitive radios: A unified approach," *IEEE Transactions on Signal Processing*, vol. 62, no. 15, pp. 3897– 3910, 2014.
- [13] D. Cohen and Y. C. Eldar, "Sub-Nyquist cyclostationary detection for cognitive radio", submitted to *IEEE Transactions on Signal Processing*, April 2016.
- [14] S. Stein, O. Yair, D. Cohen and Y. C. Eldar, "CaSCADE: Compressed Carrier and DOA Estimation", submitted to *IEEE Transactions on Signal Processing*, April 2016.
- [15] Y. Ephraim and R. M. Gray, "A unified approach for encoding clean and noisy sources by means of waveform and autoregressive model vector quantization," *IEEE Transactions on Information Theory*, vol. 34, no. 4, pp. 826–834, 1988.
- [16] D. J. Sakrison, "Source encoding in the presence of random disturbance," *IEEE Transactions on Information Theory*, vol. 14, no. 1, pp. 165–167, 1968.

- [17] J. Wolf and J. Ziv, "Transmission of noisy information to a noisy receiver with minimum distortion," *IEEE Transactions on Information Theory*, vol. 16, no. 4, pp. 406-411, 1970.
- [18] J. Li, N. Chaddha, and R. M. Gray, "Asymptotic performance of vector quantizers with a perceptual distortion measure," *IEEE Transactions on Information Theory*, vol. 45, no. 4, pp. 1082– 1091, 1999.
- [19] Y. Linde, A. Buzo, R. M. Gray, "An algorithm for vector quantizer design,"*IEEE Transactions on Communications*, vol. COM-28, no. 1, pp. 84–95, 1980.