High-performance quantization for spectral super-resolution

C. Sinan Güntürk
Courant Institute of Mathematical Sciences
New York University
New York, NY, USA
gunturk@cims.nyu.edu

Weilin Li
Courant Institute of Mathematical Sciences
New York University
New York, NY, USA
weilinli@cims.nyu.edu

***Abstract*—We show that the method of distributed noise-shaping beta-quantization offers superior performance for the problem of spectral super-resolution with quantization whenever there is redundancy in the number of measurements. More precisely, if the (integer) oversampling ratio λ is such that $\lfloor M/\lambda \rfloor - 1 \geq 4/\Delta$, where M denotes the number of Fourier measurements and Δ is the minimum separation distance associated with the atomic measure to be resolved, then for any number $K \geq 2$ of quantization levels available for the real and imaginary parts of the measurements, our quantization method guarantees reconstruction accuracy of order $O(\lambda^{3/2} K^{-\lambda/2})$, up to constants which are independent of K and λ. In contrast, memoryless scalar quantization offers a guarantee of order $O(K^{-1})$ only.**

I. Introduction

Analog-to-digital conversion is inherently lossy. It typically consists of two stages, sampling and quantization, where the sampling stage produces a stream of scalar samples and the quantization stage replaces each sample with an element of a discrete set, called the quantization alphabet. Ideally the sampling stage is lossless (or the distortion is negligible) so that the distortion is only (or primarily) caused by quantization.

The naive method of quantization is to round each scalar measurement to the nearest available level in the quantization alphabet; this is known as memoryless scalar quantization (MSQ). While MSQ has many advantages in hardware implementation (e.g. simplicity and robustness), its rate-distortion performance is highly suboptimal when the sampling map is redundant, meaning it collects more measurements than the minimal number needed for perfect reconstruction. This is due to the fact that a redundant sampling system increases the ambient dimension of the measurement vectors but not the intrinsic dimension of the manifold on which these vectors lie, and therefore, memoryless scalar quantized vectors, lying in a small neighborhood of this manifold, occupy only an asymptotically vanishing portion of the rectangular lattice of available quantization vectors. More efficient quantization methods achieve improved rate-distortion performance by utilizing some (or all) of the remaining quantization vectors as well. Noise-shaping quantizers (such as $\Sigma\Delta$ modulators) fall into this category.

This paper provides a new approach of quantizing non-harmonic Fourier measurements for the spectral super-resolution problem. Super-resolution has received considerable attention in the past several years (e.g. [3], [4]). The goal is to accurately estimate an unknown discrete measure

$$\mu = \sum_{j=1}^{S} a_j \delta_{t_j} \tag{1}$$

defined on $\mathbb{T} := [0,1)$, from its noisy samples

$$\widetilde{y}_k = y_k + z_k, \qquad k = 0, \ldots, M-1,$$

where

$$y_k := \widehat{\mu}(k) := \int_0^1 e^{-2\pi i k t} \, d\mu(t) = \sum_{j=1}^{S} a_j e^{2\pi i k t_j} \tag{2}$$

is the k-th Fourier coefficient of μ and the unknown noise vector $z := (z_k)_0^{M-1}$ satisfies $\|z\|_2 \leq \varepsilon$ for some known $\varepsilon > 0$. We emphasize that the total number of spikes S, the amplitudes $a \in \mathbb{C}^S$, and the support set $T = \{t_j\}_{j=1}^S$ are unknown.

This problem is ill-conditioned if there are points in T that are too close to one another (e.g. [10], [9], [13]). However, assuming a lower bound to their minimum separation, it has been shown that μ can be recovered from its measurements in a robust way when M is sufficiently large, meaning that the reconstruction error, when measured in a suitable metric, is controlled by the noise energy in a graceful manner, and typically linearly (e.g. [12]).

Since no structure is assumed on the noise, robustness also becomes the key property that allows for quantization. In particular, it guarantees that sufficiently high-resolution quantizers will always produce sufficiently high quality approximations. However, the question of achievable limits of quantization accuracy is open. The answer depends on the interplay between the given (fixed) parameters, such as the number of measurements, the number of quantization levels per measurement, and the minimum separation distance of the measures of interest, as well as the quantization method which itself is a design parameter.

For simplicity, we will assume in this paper that quantization is the only source of perturbation. Other (generally uncontrolled) sources of perturbations can be incorporated into quantization as well; any such perturbations typically provide a noise floor. We assume that the real and imaginary parts of each Fourier measurement is replaced by an element of a (signal-independent) quantization alphabet $\mathcal{A}$ with K levels.

The linear dependence of the reconstruction error on the noise energy implies that with simple rounding, i.e. MSQ, it is straightforward to achieve reconstruction accuracy of order $O(K^{-1})$. However, since the noise energy is measured in ℓ_2, its bound ε grows with the number of measurements, and therefore it is not even clear if there is any advantage of using any additional measurements. For example, the popular reconstruction method total-variation minimization (TV-min, see Section III) appears to be indifferent to oversampling in practice.

Our main result in this paper is that there is an alternative quantization method, called the distributed noise-shaping β-encoder (or β-quantization in short) which, together with an accompanying alternative recovery method derived from TV-min, is able to exploit any available redundancy. More precisely, if the (integer) oversampling ratio λ is such that $\lfloor M/\lambda \rfloor - 1 \geq 4/\Delta$, where M denotes the number of Fourier measurements and Δ is the minimum separation distance associated with μ, then for any number K of quantization levels, our quantization method guarantees reconstruction accuracy of order $O(\lambda^{3/2}K^{-\lambda/2})$, up to constants which are independent of K and λ. In principle our method can work with other robust recovery methods, too.

The paper is organized as follows. Section II reviews the TV-min method and discusses MSQ for spectral super-resolution. In Section II, we introduce the proposed quantization method and the main ingredients needed for its error performance analysis for spectral super-resolution, which is done in Section IV. Finally, we provide a sample numerical experiment to demonstrate the practical performance of our proposed quantization method in Section V.

II. A review of TV-min for super-resolution and MSQ

There are a number of robust recovery algorithms for super-resolution. For convenience and concreteness, we will focus on TV-min, also known under the name BLASSO (Beurling Lasso).

Let us denote by $\mathcal{F}_M$ the operator which maps the measure μ to $(\widehat{\mu}(k))_0^{M-1}$, i.e. with our notation of (2), we have $y = \mathcal{F}_M\mu$, where $y := (y_k)_0^{M-1}$. Given noisy data $\widetilde{y} \in \mathbb{C}^M$ and a bound $\varepsilon > 0$ on the noise, the TV-min algorithm outputs an estimate $\widetilde{\mu}$ of μ given by

$$\widetilde{\mu} := \arg\min\{\|\nu\|_{TV} : \|\mathcal{F}_M\nu - \widetilde{y}\|_2 \leq \varepsilon\}. \tag{3}$$

This is a convex program whose feasibility is guaranteed by the assumption $\|y - \widetilde{y}\|_2 \leq \varepsilon$. The solution may not be unique, but it is known that there is at least one minimizer which is a discrete measure which we will identify with $\widetilde{\mu}$. (See [4], [3], [11], [2] for this and other results.)

The performance of TV-min depends on the minimum separation of the measure, defined as

$$\Delta(\mu) := \min_{s \neq t, s,t \in \mathrm{supp}(\mu)} |s-t|_{\mathbb{T}}. \tag{4}$$

Here, $|s-t|_{\mathbb{T}} := \min_{n\in\mathbb{Z}} \; |s-t-n|$ is the "wrap-around" metric on $\mathbb{T}$. If the number of samples M is sufficiently large so that

$$\Delta(\mu) \geq \frac{4}{M-1}, \tag{5}$$

then $\widetilde{\mu}$ provides an accurate estimate of μ in the following sense: any spurious spikes in $\widetilde{\mu}$ are smaller than $O(\varepsilon)$, the remaining spikes in $\widetilde{\mu}$ are within $O(1/M)$ of the true spikes, and the recovered amplitudes are within $O(\varepsilon)$ of the true ones. More precisely, given the representation

$$\widetilde{\mu} = \sum_{k=1}^{\widetilde{S}} \widetilde{a}_k \delta_{\widetilde{t}_k}, \tag{6}$$

along with the "neighborhood" index sets

$$\mathcal{I}_j^M := \left\{k : |\widetilde{t}_k - t_j|_{\mathbb{T}} \leq 2\cdot 0.1649\,(M-1)^{-1}\right\}, \;\; j = 1, \dots, S, \tag{7}$$

and the residual index set

$$\mathcal{I}_0^M := \{1, \dots, \widetilde{S}\} \setminus \bigcup_{j=1}^{S} \mathcal{I}_j^M, \tag{8}$$

the following bounds are guaranteed:

$$\Big|a_j - \sum_{k\in\mathcal{I}_j^M} \widetilde{a}_k\Big| \;\; \leq \;\; C_1\varepsilon, \quad j = 1, \dots, S, \tag{9}$$

$$\sum_{k\in\mathcal{I}_j^M} |\widetilde{a}_k|\, |t_j - \widetilde{t}_k|_{\mathbb{T}}^2 \;\; \leq \;\; C_2 M^{-2}\varepsilon, \quad j = 1, \dots, S, \tag{10}$$

$$\sum_{k\in\mathcal{I}_0^M} |\widetilde{a}_k| \;\; \leq \;\; C_3\varepsilon. \tag{11}$$

Additional details can be found in [12, Theorem 1.2]; see also [1, Theorems 2.1 and 2.2] for related results.

This result provides an immediate error bound for MSQ. For each integer $K \geq 2$, let $\mathcal{Z}_K$ denote the K-term origin-symmetric arithmetic progression of integers with spacing 2 and $\mathcal{A}_K := K^{-1}\mathcal{Z}_K \subset (-1, 1)$. It is then clear that for all $u \in [-1, 1]$ there exists $q \in \mathcal{A}_K$ such that $|u - q| \leq 1/K$. Assuming that $\|\mu\|_{TV} \leq 1$ so that $\|y\|_\infty \leq 1$, it follows that for each complex measurement y_k, there is an element $q_k \in \mathcal{A}_K + i\mathcal{A}_K$ (found by separately rounding the real and the imaginary parts of y_k to elements of $\mathcal{A}_K$) such that $|y_k - q_k| \leq \sqrt{2}/K$. Consequently, we have $\|y - q\|_2 \leq \sqrt{2M}/K$. Setting $\widetilde{y} = q$ and $\varepsilon = \sqrt{2M}/K$ in (3) guarantees, in view of (9)-(11), an overall reconstruction accuracy of $O(\sqrt{M}/K)$.

We can produce a lower bound on the worst-signal reconstruction error of MSQ as follows: Even if we knew the support T of μ, where $|T| = S$, memoryless scalar quantization of the M linear measurements of all possible coefficient vectors $a \in \mathbb{C}^S$ chosen from any fixed ball results in a partition of this ball using at most $O(MK)$ hyperplanes, and therefore into at most $(cMK/S)^{2S}$ cells. (Here c is an absolute constant.) Consequently, there will always be a cell of diameter at least $O(S/MK)$ whose elements are all mapped to the same quantized vector.

Therefore, suppressing the dependence on M, it follows that MSQ cannot offer error performance better than $O(K^{-1})$.

III. Proposed quantization method

Our proposed quantization approach in this paper is based on the general framework of distributed noise-shaping β-encoding developed in [6], [7] (see also [5] and [8] for prior versions). However, the specialization for the spectral super-resolution problem requires some new choices and adaptations.

Let $\lambda \geq 1$ be an integer which should be thought of as a lower bound on the oversampling ratio. For the simplicity of discussion, we assume M is divisible by λ and set $m := M/\lambda$. Let $\beta > 1$ be a parameter which shall be chosen later, and consider the $m \times M$ matrix $V := V_\beta$

$$V := \begin{bmatrix} I_m & \beta^{-1} I_m & \cdots & \beta^{-\lambda+1} I_m \end{bmatrix}, \tag{12}$$

where I_m denotes the $m \times m$ identity matrix. Observe that

$$(Vy)_\ell = \sum_{k=0}^{\lambda-1} \beta^{-k} y_{mk+\ell} = \sum_{j=1}^{S} a_j w_j e^{-2\pi i \ell t_j}$$

for $\ell = 1, \ldots, m$, where

$$w_j := \frac{1 - \beta^{-\lambda} e^{-2\pi i m \lambda t_j}}{1 - \beta^{-1} e^{-2\pi i m t_j}}, \quad j = 1, \ldots, S. \tag{13}$$

In other words, we have the relation $Vy = \mathcal{F}_m \mu_V$ where

$$\mu_V := \sum_{j=1}^{S} b_j \delta_{t_j} \quad \text{and} \quad b_j := a_j w_j. \tag{14}$$

Observe that μ and μ_V have identical supports, but different amplitudes. However, the weights w_j satisfy

$$\frac{1}{c_\beta} \leq |w_j| \leq c_\beta \quad \text{where} \quad c_\beta := \frac{1 + \beta^{-1}}{1 - \beta^{-1}}. \tag{15}$$

Let us define $H := H_\beta$ to be the $M \times M$ matrix where

$$H_{j,k} := \begin{cases} 1, & \text{if } j = k, \\ -\beta, & \text{if } j = k+m \text{ and } 1 \leq k \leq M-m. \end{cases} \tag{16}$$

The following is a special case of [7, Lemma 2]:

Lemma 1. *Let $K \geq 2$ be an integer, and suppose the parameters $\alpha, \beta, \delta > 0$ satisfy the inequality*

$$\beta + \alpha \delta^{-1} \leq K, \tag{17}$$

and consider the quantization alphabet $\mathcal{A} := \delta(\mathcal{Z}_K + i\mathcal{Z}_K)$. Then, for any $y \in \mathbb{C}^M$ with $\|y\|_\infty \leq \alpha$, there exists $q \in \mathcal{A}^M$ and $u \in \mathbb{C}^M$ with $\|u\|_\infty \leq \sqrt{2}\delta$ satisfying the relationship

$$y - q = Hu. \tag{18}$$

The mapping $y \mapsto q$ implied by the above lemma can be implemented by means of a simple recursive algorithm. We omit the details and refer to [6], [7].

The significance of V and H is that VH is very small when β or λ is large. Indeed, as shown in [6], we have $\|VH\|_{\infty \to 2} = \sqrt{m}\beta^{-\lambda+1}$. The immediate consequence is that

$$\|Vy - Vq\|_2 \leq \|VH\|_{\infty \to 2} \|u\|_\infty \leq \sqrt{2m}\ \beta^{-\lambda+1}\delta. \tag{19}$$

The above findings provide the core strategy of our proposed quantization and recovery method. First, we note that for any $1 < \beta < K$ and any $0 \leq \alpha < \infty$, there exists $\delta > 0$ such that the condition (17) is satisfied. Hence the existence of the mapping $y \mapsto q$ (with a fixed quantization alphabet $\mathcal{A}$) is guaranteed over any bounded set of inputs y. Next, recall that with $y = \mathcal{F}_M \mu$, we have $Vy = \mathcal{F}_m \mu_V$. Since Vq is now a small perturbation of Vy, we can obtain a close approximation $\widetilde{\mu}_V$ of μ_V by means of any robust super-resolution recovery method, such as the TV-min algorithm. Then, since μ and μ_V have identical supports, we can define an approximate recovery $\widetilde{\mu}$ by means of approximate weights $\widetilde{w}_j$ derived from the approximate support $\widetilde{T}$ of $\widetilde{\mu}_V$.

Let us summarize the proposed quantization method.

System parameters and assumptions:
- M (number of Fourier measurements),
- α (upper bound on TV-norm of the input measures),
- Δ (lower bound on minimum separation distance),
- $M = \lambda m$, $\lambda \geq 1$, $m - 1 \geq 4/\Delta$,
- K (number of quantization levels),
- β and δ such that (17) holds.

Encoding (quantization) stage:
- Input to quantizer: y such that $\|y\|_\infty \leq \alpha$,
- V and H defined via (12) and (16),
- Quantization alphabet: $\mathcal{A} := \delta(\mathcal{Z}_K + i\mathcal{Z}_K)$,
- Output of quantizer: $q \in \mathcal{A}^M$ such that $\|Vy - Vq\|_2 \leq \sqrt{2m}\ \beta^{-\lambda+1}\delta$.

Decoding (recovery) stage:
- Input to decoder: q,
- Compute a minimum TV-norm measure $\widetilde{\mu}_V$ of the form $\sum_{k=1}^{\widetilde{S}} \widetilde{b}_k \delta_{\widetilde{t}_k}$ satisying $\|\mathcal{F}_m \widetilde{\mu}_V - Vq\|_2 \leq \varepsilon_V$ where
$$\varepsilon_V := \sqrt{2m}\ \beta^{-\lambda+1}\delta;$$
abort if it cannot be found (e.g. invalid measurements),
- Set
$$\widetilde{w}_k = \frac{1 - \beta^{-\lambda} e^{-2\pi i m \lambda \widetilde{t}_k}}{1 - \beta^{-1} e^{-2\pi i m \widetilde{t}_k}}, \quad k = 1, \cdots, \widetilde{S},$$
- Output of decoder: $\widetilde{\mu} = \sum_{k=1}^{\widetilde{S}} \widetilde{a}_k \delta_{\widetilde{t}_k}$ with $\widetilde{a}_k := \widetilde{b}_k / \widetilde{w}_k$.

IV. Error analysis

Let us start by noting that when the input to the quantizer y equals $\mathcal{F}_M \mu$ for some μ of the form (1) with $\Delta(\mu) \geq \Delta$ and $\|\mu\|_{TV} \leq \alpha$, then the decoder will always output a measure $\widetilde{\mu}$, thanks to the fact that $Vy = \mathcal{F}_m \mu_V$ where μ_V is defined by (14) which guarantees that μ_V is a feasible measure for the TV-min program.

Let us now proceed to find an error bound for $\widetilde{\mu}$. We start by comparing $\widetilde{\mu}_V$ to μ_V. With the error bounds of the general TV-min method reviewed in Section II, we have

$$\Big| b_j - \sum_{k \in \mathcal{I}_j^m} \widetilde{b}_k \Big| \leq C_1 \varepsilon_V, \quad j = 1, \ldots, S, \tag{20}$$

$$\sum_{k \in \mathcal{I}_j^m} |\widetilde{b}_k|\, |t_j - \widetilde{t}_k|_{\mathbb{T}}^2 \leq C_2 m^{-2} \varepsilon_V, \quad j = 1, \ldots, S, \tag{21}$$

$$\sum_{k \in \mathcal{I}_0^m} |\widetilde{b}_k| \leq C_3 \varepsilon_V. \tag{22}$$

where the index sets $\mathcal{I}_j^m$, $j = 0, \dots, S$ are as in (7) and (8), only for m measurements. Note that for all $j \in \{1, \dots, S\}$,

$$\Big|a_j - \sum_{k\in\mathcal{I}_j^m} \widetilde{a}_k\Big| \le \frac{1}{|w_j|}\Big|b_j - \sum_{k\in\mathcal{I}_j^m} \widetilde{b}_k\Big| + \sum_{k\in\mathcal{I}_j^m} |\widetilde{b}_k| \Big|\frac{1}{w_j} - \frac{1}{\widetilde{w}_k}\Big|. \tag{23}$$

With (15) and (20), the first term is bounded by $c_\beta C_1 \varepsilon_V$. For the second term, we note that

$$\Big|w_j^{-1} - \widetilde{w}_k^{-1}\Big| \le m\, C_{\beta,\lambda}\, |t_j - \widetilde{t}_k|_{\mathbb{T}} \tag{24}$$

where $C_{\beta,\lambda}$ stands for the Lipschitz constant of the map

$$t \mapsto \frac{1 - \beta^{-1} e^{-2\pi i t}}{1 - \beta^{-\lambda} e^{-2\pi i \lambda t}}, \quad t \in \mathbb{T}.$$

It can be shown $C_{\beta,\lambda} \le 4\pi\lambda\beta(\beta-1)^{-2}$.

Using (24) and Cauchy-Schwarz, we see that the second term in (23) is bounded by

$$m\, C_{\beta,\lambda} \Big(\sum_{k\in\mathcal{I}_j^m} |\widetilde{b}_k|\, |t_j - \widetilde{t}_k|_{\mathbb{T}}^2\Big)^{1/2} \Big(\sum_{k\in\mathcal{I}_j^m} |\widetilde{b}_k|\Big)^{1/2}$$

Note that $\|\widetilde{b}\|_1 \le c_\beta \alpha$ since $\|\widetilde{b}\|_1 = \|\widetilde{\mu}_V\|_{TV}$ and

$$\|\widetilde{\mu}_V\|_{TV} \le \|\mu_V\|_{TV} \le \|w\|_\infty \|a\|_1 \le c_\beta \|\mu\|_{TV} \le c_\beta \alpha.$$

Hence, with (21) we deduce

$$\sum_{k\in\mathcal{I}_j^m} |\widetilde{b}_k|\, \Big|\frac{1}{w_j} - \frac{1}{\widetilde{w}_k}\Big| \le C_{\beta,\lambda} \sqrt{c_\beta \alpha} \sqrt{C_2 \varepsilon_V}. \tag{25}$$

Injecting (25) into (23) we have

$$\Big|a_j - \sum_{k\in\mathcal{I}_j^m} \widetilde{a}_k\Big| \le c_\beta C_1 \varepsilon_V + C_{\beta,\lambda} \sqrt{c_\beta \alpha} \sqrt{C_2 \varepsilon_V}. \tag{26}$$

Finally, we also have

$$\sum_{k\in\mathcal{I}_j^m} |\widetilde{a}_k|\, |t_j - \widetilde{t}_k|_{\mathbb{T}}^2 \le c_\beta C_2 m^{-2} \varepsilon_V, \tag{27}$$

$$\sum_{k\in\mathcal{I}_0^m} |\widetilde{a}_k| \le c_\beta C_3 \varepsilon_V. \tag{28}$$

which follow readily from (15), (21) and (22).

At this point, we note the following elementary fact: For any $K \ge 2$, setting $\beta := K(\lambda+1)/(\lambda+2)$ and $\delta := (\lambda+2)\alpha/K$ results in $\beta + \alpha\delta^{-1} = K$ and $\delta\beta^{-\lambda+1} < e\alpha(\lambda+1)K^{-\lambda}$. (See, e.g. [6, Lemma 3.2] and [7, Lemma 1].) This choice of parameters results in $\varepsilon_V \le e\alpha\sqrt{2m}(\lambda+1)K^{-\lambda}$. Furthermore it is readily seen that $\beta \ge 4/3$, $c_\beta \le 7$ and $C_{\beta,\lambda} \le 12\pi\lambda$. Hence it follows from (26), (27) and (28) that μ is approximated by $\widetilde{\mu}$ up to resolution $O(\sqrt{M}\lambda^{3/2}K^{-\lambda/2})$.

V. Numerical results

We compare the reconstruction error, quantified by the term on the left hand side of (26), when the Fourier samples are quantized using our proposed beta-quantization versus MSQ. More specifically, we set $\Delta = 1/10$ and we randomly select a measure μ such that $\Delta(\mu) \ge \Delta$; the amplitudes are chosen uniformly at random and normalized to have unit ℓ^1 norm.

Fig. 1. Average reconstruction error of MSQ and β-quantization.

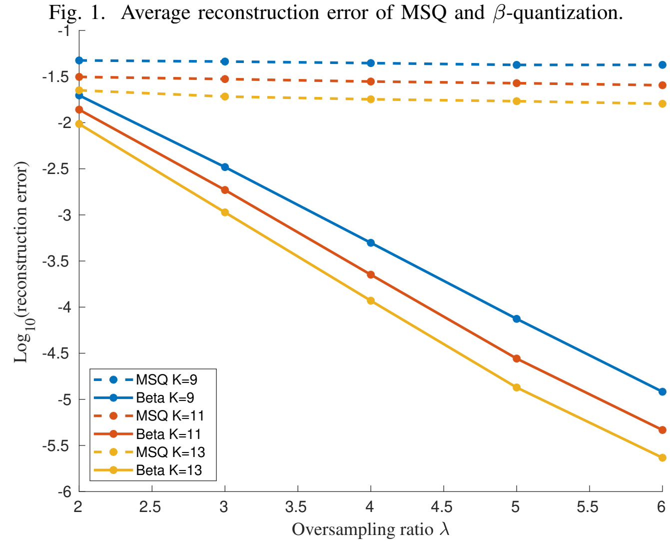

For various choices of λ and K, we quantize the Fourier measurements using both MSQ and β-quantization. Figure 1 displays the reconstruction error as a function of λ, averaged over 110 trials. The experiment validates our theoretical results and also shows that performance of MSQ is suboptimal in the over-sampling regime.

References

[1] Jean-Marc Azaïs, Yohann De Castro, and Fabrice Gamboa. Spike detection from inaccurate samplings. *Applied and Computational Harmonic Analysis*, 38(2):177–195, 2015.

[2] John J. Benedetto and Weilin Li. Super-resolution by means of Beurling minimal extrapolation. *Applied and Computational Harmonic Analysis*, 2018.

[3] Emmanuel J. Candès and Carlos Fernandez-Granda. Super-resolution from noisy data. *Journal of Fourier Analysis and Applications*, 19(6):1229–1254, 2013.

[4] Emmanuel J. Candès and Carlos Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on Pure and Applied Mathematics*, 67(6):906–956, 2014.

[5] Evan Chou. *Beta-duals of frames and applications to problems in quantization*. PhD thesis, New York University, 2013.

[6] Evan Chou and C. Sinan Güntürk. Distributed noise-shaping quantization: I. Beta duals of finite frames and near-optimal quantization of random measurements. *Constructive Approximation*, 44(1):1–22, 2016.

[7] Evan Chou and C. Sinan Güntürk. Distributed noise-shaping quantization: II. Classical frames. In *Excursions in Harmonic Analysis, Volume 5*, pages 179–198. Springer, 2017.

[8] Evan Chou, C. Sinan Güntürk, Felix Krahmer, Rayan Saab, and Özgür Yılmaz. Noise-shaping quantization methods for frame-based and compressive sampling systems. In *Sampling theory, a renaissance*, pages 157–184. Springer, 2015.

[9] Laurent Demanet and Nam Nguyen. The recoverability limit for superresolution via sparsity. *arXiv preprint arXiv:1502.01385*, 2015.

[10] David L. Donoho. Superresolution via sparsity constraints. *SIAM Journal on Mathematical Analysis*, 23(5):1309–1331, 1992.

[11] Vincent Duval and Gabriel Peyré. Exact support recovery for sparse spikes deconvolution. *Foundations of Computational Mathematics*, 15(5):1315–1355, 2015.

[12] Carlos Fernandez-Granda. Support detection in super-resolution. In *Proceedings of the 10th International Conference on Sampling Theory and Applications*, pages 145–148, 2013.

[13] Weilin Li and Wenjing Liao. Stable super-resolution limit and smallest singular value of restricted fourier matrices. *arXiv preprint arXiv:1709.03146*, 2017.