

Estimation from Quantized Gaussian Measurements: When and How to Use Dither

Joshua Rapp, *Student Member, IEEE*, Robin M. A. Dawson, *Member, IEEE*, and Vivek K Goyal, *Fellow, IEEE*

Abstract—Subtractive dither is a powerful method for removing the signal dependence of quantization noise for coarsely-quantized signals. However, estimation from dithered measurements often naively applies the sample mean or midrange, even when the total noise is not well described with a Gaussian or uniform distribution. We show that the generalized Gaussian distribution approximately describes subtractively-dithered, quantized samples of a Gaussian signal. Furthermore, a generalized Gaussian fit leads to simple estimators based on order statistics that match the performance of more complicated maximum likelihood estimators requiring iterative solvers. The order statistics-based estimators outperform both the sample mean and midrange for nontrivial sums of Gaussian and uniform noise. Additional analysis of the generalized Gaussian approximation yields rules of thumb for determining when and how to apply dither to quantized measurements. Specifically, we find subtractive dither to be beneficial when the ratio between the Gaussian standard deviation and quantization interval length is roughly less than $1/3$. If that ratio is also greater than $0.822/K^{0.930}$ for the number of measurements $K > 20$, we present estimators more efficient than the midrange.

Index Terms—Quantization, subtractive dither, generalized Gaussian distribution, order statistics, L-estimator, alpha-trimmed mean, midrange

I. INTRODUCTION

Estimation of the mean of a Gaussian distribution from independent and identically distributed (i.i.d.) samples is a canonical problem in statistics, yet it has important subtleties when the samples are quantized. Without quantization, the sample mean is an unbiased, efficient, and consistent estimator. With uniformly quantized samples, the situation is immediately more complicated: The sample mean is an unbiased estimate only when the true mean falls on the quantizer’s reproduction grid or asymptotically in the limit of fine quantization [1]; and in the opposite extreme of very coarse quantization, all the samples are identical, so the estimates do not even improve with increasing numbers of data samples.

The use of subtractive dither changes the situation substantially. The sample mean is then an unbiased and consistent estimator—like in the unquantized case—but it may be arbitrarily far from minimizing the mean-squared error

(MSE). For example, when the population variance vanishes, the sample mean estimator has MSE inversely proportional to the number of samples, whereas the MSE achieved by the midrange estimator is inversely proportional to the *square* of the number of samples [2].

In this paper, we develop estimators for cases where the quantization is neither extremely fine nor extremely coarse. The motivation for this work stemmed from a series of experiments performed by the authors and colleagues with single-photon lidar. In [3], temporally spreading a narrow laser pulse, equivalent to adding non-subtractive Gaussian dither, was found to reduce the effects of the detector’s coarse temporal resolution on ranging accuracy. Later work on a similar system showed that implementing subtractive dither could likewise reduce the effects of coarse quantization [4], [5]. Our aim was then to compare the two approaches by determining performance limits, optimal estimators, and when one method might be preferable over the other. The estimators we develop in this work are based on a generalized Gaussian (GG) approximation for the combination of sample variation and quantization noise, which the authors first proposed in [5]. While the benefit of the GG approximation did not yield improved results for the lidar data due to model mismatch, our framework is valid for a more general set of problems in which quantization of a Gaussian scalar signal occurs. We propose a number of estimators for additive GG noise and find a clear computational advantage with negligible loss in accuracy for simple estimators based on order statistics.

A. Main Contributions

This paper makes the following contributions:

1) *Estimation Efficiency*: We demonstrate the inefficiency of the mean and midrange estimators for subtractively-dithered measurements of a Gaussian signal by deriving the maximum likelihood estimator and Cramér-Rao bound.

2) *Generalized Gaussian Approximation*: We expand upon the generalized Gaussian approximation introduced in [5] for the sum of Gaussian and uniform random variables that arises from subtractively dithered quantization of a Gaussian signal, using the approximation to determine three distinct regimes of estimator behavior.

3) *Estimator Proposal*: We consider a family of location estimators based on the GG approximation, in particular linear combinations of the measurement order statistics. We introduce a version of the trimmed mean estimator with trimming determined by the GG approximation that is simple, computationally efficient, and performs as well as the maximum

This work was supported in part by a Draper Fellowship under Navy contract N00030-16-C-0014, the U.S. National Science Foundation under Grant 1422034, and the DARPA REVEAL program under Contract No. HR0011-16-C-0030. This work was presented in part at the IEEE International Conference on Image Processing, Athens, Greece, October 2018.

J. Rapp and V. K. Goyal are with the Department of Electrical and Computer Engineering, Boston University, Boston, MA, 02215 USA (e-mail: jrapp@bu.edu; v.goyal@ieee.org).

J. Rapp and R. Dawson are with the Charles Stark Draper Laboratory, Cambridge, MA, 02139 USA.

likelihood (ML) estimator. Monte Carlo estimator comparisons are shown versus the number of measurements K and versus σ_Z/Δ , the ratio of the Gaussian standard deviation to the quantization bin size.

4) *Rules of Thumb*: We determine several key rules of thumb for deciding when and how to use subtractive dither. For instance, we find dither is not beneficial roughly for $\sigma_Z/\Delta > 1/3$; below this value, however, applying subtractive dither and a GG-based estimator lowers the MSE. Moreover, if the quantization is coarser than $\sigma_Z/\Delta = 0.822/K^{0.930}$ and $K > 20$, then the midrange is a good estimator.

B. Outline

This paper is organized as follows. Section II sets up the problem of measuring a Gaussian signal with a subtractively-dithered quantizer and explores the fact that the mean and midrange are inefficient estimators. Section III motivates the use of the generalized Gaussian distribution and estimators based on order statistics. Section IV discusses several estimator implementations for our noise model. Section V introduces mean-squared error expressions as a guide to better understanding the results of numerical simulation presented in Section VI, which tests several estimators and compares the use of quantized data with and without dithering. Finally, Section VII presents our conclusions regarding which estimators to use and when to apply dither.

II. FORMULATION, BACKGROUND, AND MOTIVATION

A. Quantized Measurement

We begin by presenting and expanding upon the signal acquisition model introduced in [5]. Suppose we have an unknown constant signal μ_X corrupted by additive, zero-mean Gaussian noise $Z \sim \mathcal{N}(0, \sigma_Z^2)$. Then estimation of μ_X from K independent samples

$$X_i = \mu_X + Z_i, \quad i = 1, 2, \dots, K,$$

is straightforward, as the sample mean $\bar{\mu} = (1/K) \sum_{i=1}^K X_i$ can easily be shown to be an efficient estimator of the mean of a Gaussian distribution. However, all measurement instruments perform some quantization. For instance, consider a uniform midtread quantizer $q(\cdot)$ with bin size Δ applied to X_i when $\sigma_Z \ll \Delta$. Except when μ_X is close to a quantizer threshold, it will be the case that $U_i = q(X_i)$ is identical for all i , so that the “quantized-sample mean” given as

$$\hat{\mu}_Q = \frac{1}{K} \sum_{i=1}^K U_i, \quad (1)$$

is no more informative an estimate of μ_X than any single measurement. For σ_Z not too small compared to Δ , estimation error can be reduced by properly accounting for the quantization and the underlying distribution, e.g., via the maximum

likelihood estimator for quantized samples of a Gaussian signal [6], [7]:

$$\hat{\mu}_{QML} = \arg \max_{\mu_X} \sum_{i=1}^K \log \left[\frac{\Phi \left(\frac{u_i - \mu_X + \frac{\Delta}{2}}{\sigma_Z} \right) - \Phi \left(\frac{u_i - \mu_X - \frac{\Delta}{2}}{\sigma_Z} \right)}{\sigma_Z} \right], \quad (2)$$

where $\Phi(\cdot)$ is the cumulative distribution function (CDF) of a standard normal random variable. Still, $\hat{\mu}_{QML}$ is no more accurate than $\hat{\mu}_Q$ when all of the samples have the same value. Because of the coarse quantization mapping every value in $[j\Delta - \Delta/2, j\Delta + \Delta/2]$ to $j\Delta$ for $j \in \mathbb{Z}$, the resolution of an estimate $\hat{\mu}_X$ is limited by the bin size Δ and the quantization error is signal-dependent.

Statisticians have long recognized that working with rounded data is not the same as working with underlying continuous-valued data. Let X_{hist} be the continuous random variable with density constant on intervals $((j - \frac{1}{2})\Delta, (j + \frac{1}{2})\Delta)$ with $P(X_{\text{hist}} \in ((j - \frac{1}{2})\Delta, (j + \frac{1}{2})\Delta)) = P(U = k\Delta)$, for all $j \in \mathbb{Z}$. Because of the piecewise-constant form, X_{hist} is said to have a *histogram density* [8]. The widely known *Sheppard's corrections* introduced in [9], [10] relate the moments of U and the moments of X_{hist} [11]. From the construction of X_{hist} , it is immediate that these corrections are zero for odd moments. See [12] for a review of Sheppard's corrections and [13] for results for autoregressive and moving average processes and more recent references.

The moments of X_{hist} being close to the moments of X depends on continuity arguments and Δ being small. In contrast, our interest here is in situations where the quantization is coarse relative to the desired precision in estimating μ_X . Quantization may be coarse because of limitation of instruments, such as the fundamental trade-offs in analog-to-digital converters [14] or the time resolution in time-correlated single photon counting [15], which may be coarse relative to the resolution desired for time-of-flight ranging.

When quantizing X_{hist} , the quantization error $E_{\text{hist}} = q(X_{\text{hist}}) - X_{\text{hist}}$ is uniformly distributed on $[-\Delta/2, \Delta/2]$ and independent of X_{hist} . In general, however, quantization error being uniformly distributed and independent of the input does not extend to the quantization of X ; approximating quantization error as such—without regard to whether the input has a histogram density—is often called the “additive-noise model,” “quantization-noise model,” or “white-noise model.” A substantial literature is devoted to understanding the validity of this approximation, e.g. [16]–[20].

One approach to improving the precision of estimates from quantization-limited measurements is the use of dither, a small signal introduced before the discretization to produce enough variation in the input such that it spans multiple quantization levels. By combining multiple dithered measurements, estimates can achieve resolution below the least-significant bit and the result may also have desirable statistical and perceptual properties, such as whitened noise. Early applications empirically demonstrating the benefits of dither include control systems [21], [22], image display [23], [24], and audio [25], with numerous contributions to the statistical theory developed

in [16], [17], [26]–[30], among others. More recent work has focused on varying the quantizer thresholds primarily for 1-bit measurements in wireless sensing networks, including [31]–[37].

B. Subtractively-Dithered Quantization

When not only the dither distribution but also the exact value is known for each sample, it is possible to achieve an additional reduction in estimation error. Papadopoulos et al. showed that a sawtooth waveform was optimal for known dither, whereas further improvement could be achieved with adaptively-selected dither values [31]. However, a deterministic-length dither sequence is ill-suited to applications such as single-photon lidar, for which the number of measurements is random, and adaptive dither may be infeasible to implement for such high-speed systems. Instead, we consider *subtractively-dithered quantization*, which uses a random dither signal that is measured for every sample. The ML estimator for subtractive dither can be shown to be equivalent to generic quantization schemes with known dither values, e.g., [33], [37]. However, principled selection of the subtractive dither signal distribution makes the quantization error uniformly distributed and independent of the input, which we show can lead to simple, non-iterative estimators.

Define the dither signal D_i , $i = 1, \dots, K$ as a sequence of i.i.d. random variables, independent of the noisy quantizer input X_i . The output of a subtractively-dithered quantizer is

$$Y_i = q(X_i + D_i) - D_i, \quad (3)$$

with the quantization error defined as

$$W_i = Y_i - X_i. \quad (4)$$

Define the characteristic function of the dither signal probability density function (PDF) as

$$M_D(ju) = \mathbb{E}[e^{juD}]. \quad (5)$$

Then Schuchman's condition [28] is the property of the dither PDF that

$$M_D(j2\pi\ell/\Delta) = 0, \quad \ell \in \mathbb{Z} \setminus 0. \quad (6)$$

As long as the quantizer has a sufficient number of levels so that it does not overload, by [17], [29] the Schuchman condition is necessary and sufficient for X_i to be independent of W_j for all i, j , with i.i.d. $W_i \sim \mathcal{U}[-\Delta/2, \Delta/2]$. Subtractive dither often uses a uniform dither signal with $D \sim \mathcal{U}[-\Delta/2, \Delta/2]$ because its characteristic function

$$M_D(ju) = \frac{\sin(u\Delta/2)}{u\Delta/2}$$

meets Schuchman's condition (6).

The rest of this paper considers only when Schuchman's condition is met, with an i.i.d. input signal of the form $X_i = \mu_X + Z_i$, $Z \sim \mathcal{N}(0, \sigma_Z^2)$, an i.i.d. dither signal $D_i \sim \mathcal{U}[-\Delta/2, \Delta/2]$ independent of the input signal, and a non-overloading uniform quantizer.¹

¹Since Gaussian noise has infinite support, in principle the quantizer must have an infinite number of levels.

Then the dithered measurements take the form

$$Y_i = \mu_X + Z_i + W_i, \quad (7)$$

and the problem of estimating μ_X simply becomes one of mitigating independent additive noise. The sum of the Gaussian and uniform terms can be combined into a single total noise term to obtain

$$Y_i = \mu_X + V_i, \quad (8)$$

where $V_i = Z_i + W_i$ are i.i.d. Then the means and variances simply add so that $\mu_V = 0$ and $\sigma_V^2 = \sigma_Z^2 + \Delta^2/12$.

For convenient shorthand, we refer to measurements from a quantizer without dither as “quantized” and measurements from a subtractively-dithered quantizer as “dithered.” The usual approach to estimating μ_X from K dithered measurements Y_i , $i = 1, 2, \dots, K$, is via the sample mean

$$\hat{\mu}_{\text{mean}} = \frac{1}{K} \sum_{i=1}^K Y_i. \quad (9)$$

The MSE of the sample mean is

$$\text{MSE}(\text{mean}) = \sigma_V^2/K, \quad (10)$$

which is $O(K^{-1})$. Although using the sample mean is logical when $\sigma_Z \gg \Delta$ so that the contribution of the uniform noise component is negligible, the sample mean is not in general an efficient estimator. For example, in an alternative case of $\sigma_Z = 0$, a maximum likelihood (ML) estimator is the midrange²

$$\hat{\mu}_{\text{mid}} = \frac{1}{2} (Y_{(1)} + Y_{(K)}), \quad (11)$$

where $Y_{(1)} \leq Y_{(2)} \leq \dots \leq Y_{(K)}$ are the order statistics of the K measured samples. Whereas the MSE of the sample mean for $\sigma_Z = 0$ is $\Delta^2/(12K)$, the MSE of the midrange is

$$\text{MSE}(\text{mid}) = \Delta^2/[2(K+1)(K+2)], \quad (12)$$

which is $O(K^{-2})$ and hence better than the sample mean by an unbounded factor [39]. Nevertheless, the midrange is not a good estimator in the general case of $\sigma_Z > 0$, as it relies on the finite support of the uniform distribution. If instead σ_Z is much larger than Δ , rendering the uniform component negligible, then the MSE of the midrange would only improve as $O(1/\log(K))$ [40]. As others have noted for quantization of a Gaussian signal without dither [6], [7], the key figure of merit for determining estimator performance is then σ_Z/Δ , a measure of the relative sizes of the noise components. We observe that normalizing the MSE by Δ^2 removes the separate dependence on σ_Z and Δ , resulting in

$$\text{NMSE}(\text{mean}) = [(\sigma_Z/\Delta)^2 + 1/12]/K, \quad (13)$$

and

$$\text{NMSE}(\text{mid}) = 1/[2(K+1)(K+2)]. \quad (14)$$

Except in trivial cases ($\sigma_Z \gg \Delta$ or $\sigma_Z \ll \Delta$), V has neither Gaussian nor uniform distribution, so the conventional mean

²Any statistic in $[Y_{(n)} - \Delta/2, Y_{(1)} + \Delta/2]$ is an ML estimator for the mean of a uniform distribution with known variance [38, p. 282]. The midrange is commonly used because it is unbiased and the minimum-variance estimator among linear functions of order statistics [39]. However, no uniformly minimum-variance unbiased estimator exists [38, p. 331].

and midrange estimators are expected to be suboptimal. Furthermore, existing nonlinear processing schemes for dithered measurements do not adapt to best suit the noise statistics [41]. A first approach to finding a better estimator for arbitrary σ_Z/Δ is to derive the ML estimator for the dithered noise model. From the definitions of the random variables, the PDF of W is

$$f_W(w) = \begin{cases} 1/\Delta, & w \in [-\Delta/2, \Delta/2] \\ 0, & \text{otherwise,} \end{cases}$$

and the PDF of Z is $f_Z(z) = \phi(z/\sigma_Z)/\sigma_Z$, where $\phi(x)$ is the standard normal PDF. Since the total noise is the sum of independent noise terms, the PDF of the samples is given by the convolution

$$\begin{aligned} f_V(v) &= f_Z(z) * f_W(w) \\ &= \frac{1}{\Delta} \int_{-\frac{\Delta}{2}}^{\frac{\Delta}{2}} f_Z(v - \tau) d\tau \\ &= \frac{1}{\Delta} \left[\Phi\left(\frac{v + \frac{\Delta}{2}}{\sigma_Z}\right) - \Phi\left(\frac{v - \frac{\Delta}{2}}{\sigma_Z}\right) \right]. \end{aligned} \quad (15)$$

For i.i.d. samples from a dithered quantizer, the likelihood function is then

$$\begin{aligned} \mathcal{L}(\{y_i\}_{i=1}^K; \mu_X) &= \prod_{i=1}^K f_V(y_i - \mu_X) \\ &= \prod_{i=1}^K \frac{1}{\Delta} \left[\Phi\left(\frac{y_i - \mu_X + \frac{\Delta}{2}}{\sigma_Z}\right) - \Phi\left(\frac{y_i - \mu_X - \frac{\Delta}{2}}{\sigma_Z}\right) \right]. \end{aligned} \quad (16)$$

From the log-likelihood, the dithered-sample ML estimator of μ_X is

$$\hat{\mu}_{\text{DML}} = \arg \max_{\mu_X} \sum_{i=1}^K \log \left[\Phi\left(\frac{y_i - \mu_X + \frac{\Delta}{2}}{\sigma_Z}\right) - \Phi\left(\frac{y_i - \mu_X - \frac{\Delta}{2}}{\sigma_Z}\right) \right]. \quad (17)$$

The ML estimator is notably identical to (2), except the dithered measurements are not discrete-valued as are the samples used for $\hat{\mu}_{\text{QML}}$.

To determine the efficiency of the mean, midrange, and DML estimators, we derive the Cramér-Rao bound (CRB), which is a limit on the MSE that an unbiased estimator can achieve [42, Chapter 4.2.2]. The normalized CRB is derived in Appendix A for one dithered measurement to be

$$\text{NCRB}(\mu_X) = \frac{(\sigma_Z/\Delta)^2}{\int \frac{\left[\phi\left(\frac{u-1/2}{\sigma_Z/\Delta}\right) - \phi\left(\frac{u+1/2}{\sigma_Z/\Delta}\right) \right]^2}{\Phi\left(\frac{u+1/2}{\sigma_Z/\Delta}\right) - \Phi\left(\frac{u-1/2}{\sigma_Z/\Delta}\right)} du}, \quad (18)$$

which can be evaluated via numerical integration. Note that the uniform PDF does not meet the regularity condition required for the CRB to apply, so (18) is not expected to be meaningful for $\sigma_Z/\Delta = 0$.




Fig. 1: Applied to subtractively-dithered measurements, the midrange is approximately optimal only for small σ_Z/Δ (Regime I, red), and the sample mean is approximately optimal only for large σ_Z/Δ (Regime III, blue). For confirmation, in Regime I the midrange approaches $\text{NMSE} = 1/[2(K+1)(K+2)] \approx 3.12 \times 10^{-5}$ and the mean approaches $\text{NMSE} = (1/12)/K \approx 6.67 \times 10^{-4}$, and in Regime III the mean approaches $\text{NMSE} = (\sigma_Z/\Delta)^2/K$, which is 8.00×10^{-3} at $\sigma_Z/\Delta = 1$. We seek an estimator simpler than the dithered-sample maximum likelihood that performs at least as well as the mean and midrange for intermediate values of σ_Z/Δ (Regime II, green).

Fig. 1 illustrates the suboptimality of the mean and midrange estimators compared to $\hat{\mu}_{\text{DML}}$ for intermediate values of σ_Z/Δ . In a Monte Carlo simulation with $T = 20000$ trials, $K = 125$ measurements were generated according to (3), where both μ_X and D were selected uniformly at random over $[-\Delta/2, \Delta/2]$. Computing the normalized MSE of the $\hat{\mu}_{\text{mean}}$, $\hat{\mu}_{\text{mid}}$, and $\hat{\mu}_{\text{DML}}$ estimates as

$$\text{NMSE}(\hat{\mu}_X) = \frac{1}{T} \sum_{t=1}^T \left(\frac{\mu_X - \hat{\mu}_X}{\Delta} \right)^2 \quad (19)$$

reveals how the performance of each estimator changes as a function of σ_Z/Δ .

Fig. 1 highlights three distinct regimes of estimator behavior. In Regime I (red), the Gaussian noise component is negligible, so the ML estimator and the midrange are nearly identical and outperform the mean. In Regime III (blue), the uniform noise component is negligible, so the ML estimator and the mean are nearly identical and outperform the midrange. In Regime II (green), neither the uniform nor the Gaussian component dominates, and the DML estimator performs significantly better than both the mean and midrange. Still, $\hat{\mu}_{\text{DML}}$ does not achieve the CRB for small σ_Z/Δ , indicating that an efficient estimator of μ_X does not exist; however, $\hat{\mu}_{\text{DML}}$ is asymptotically efficient in K for $\sigma_Z/\Delta > 0$ [43, Theorem 7.1].

From the results in Fig. 1, it may seem obvious that $\hat{\mu}_{\text{DML}}$ is a better choice than $\hat{\mu}_{\text{mean}}$ or $\hat{\mu}_{\text{mid}}$ for any value of σ_Z/Δ . However, $\hat{\mu}_{\text{DML}}$ requires iterative solution, thus making it far more computationally complex than the mean and midrange. In this work, one of our primary aims is to find a computationally simple estimator that can likewise outperform the mean and midrange in Regime II. We show that a generalized Gaussian approximation to the total noise of a

dithered quantizer gives rise to order statistics-based estimators that approach the performance of $\hat{\mu}_{\text{DML}}$. In addition, we compare their results to those from quantized measurements without dither, leading to design rules for when to use dither and which estimator to apply.

C. Mixed Measurements of Vector Signals

Whereas this paper is focused on the estimation of a single scalar value from multiple direct noisy measurements, many other estimation problems involving quantized data have been studied extensively. In particular, interest in linear inverse problems—both undersampled and oversampled—has resulted in work on estimating vectors from quantized linearly mixed measurements, with and without subtractive dither.

Quantized, mixed, noisy measurements of a vector $\mathbf{x}$ can be represented as $\mathbf{y} = Q(A\mathbf{x} + \mathbf{z})$, where $A \in \mathbb{C}^{K \times N}$ is a linear operator, $\mathbf{z}$ is additive noise, and $Q(\cdot)$ represents scalar quantization. Theoretical results on how well $\mathbf{x}$ can be estimated generally depend on the structure of A , such as being an oversampled inverse discrete Fourier transform—thus modeling oversampled analog-to-digital conversion (OADC)—or being large and random.

For noiseless ($\mathbf{z} = 0$) OADC, $O(K^{-2})$ upper- and lower-bounds on MSE using deterministic analyses [44], [45] are reminiscent of (12) and similarly rooted in quantized values providing hard constraints on $\mathbf{x}$. For general A , such constraints can be expressed with a linear program [46], and introduction of subtractive dither makes the optimal $O(K^{-2})$ rate provably achievable by a very simple algorithm [47]. The compressive case of $K < N$ is addressed, for example, in [48].

For Gaussian $\mathbf{z}$, [49] provides a method applicable with nonuniform quantization to provide an ℓ_1 -regularized estimate of $\mathbf{x}$. More general priors and quantizers (potentially non-regular as well as nonuniform) can be incorporated in the method of [50]. Like these earlier works, this paper also addresses the case of Gaussian $\mathbf{z}$, but A being a $K \times 1$ matrix of 1s makes it qualitatively different, in part because the estimate of a single scalar from $K \gg 1$ measurements need not be regularized.

III. GENERALIZED GAUSSIAN APPROXIMATION AND ESTIMATION

In order to find a simple estimator for Regime II, we begin by examining the other two regimes and the simple forms of the ML estimator there. We notice that the uniform and Gaussian noise distributions in Regimes I and III are special cases of the generalized Gaussian distribution (GGD), which has PDF [51]

$$f_{\hat{V}}(v; \mu, \sigma, p) = \frac{1}{2\Gamma(1 + 1/p)A(p)} \exp \left\{ - \left(\frac{|v - \mu|}{A(p)} \right)^p \right\}, \quad (20)$$

where $A(p) = \sqrt{\sigma^2 \Gamma(1/p) / \Gamma(3/p)}$ and $\Gamma(\cdot)$ is the Gamma function. In addition to mean and variance parameters μ and σ^2 , the GG density has a third parameter p that controls the exponential decay of its tails. When $p = 2$ or $p \rightarrow \infty$, the GGD simplifies to the Gaussian or uniform distributions,

respectively. Another special case of the GGD is the Laplace distribution for $p = 1$.

For each of the special cases, we further notice that the ML estimator (median, mean, and midrange for $p = 1, 2, \infty$) is a linear combination of order statistics. When $p = 1$, only the middle order statistic has nonzero weight, whereas the reverse is true for $p \rightarrow \infty$, with all weight on the two extreme samples. For $p = 2$, all of the order statistics are equally weighted. With these two observations in mind, we hypothesize that, if there is a value of p that approximates intermediate combinations of uniform and Gaussian noise, then there may be a corresponding order statistics-based estimator that approaches the performance of $\hat{\mu}_{\text{DML}}$.

A. Approximation

For our stated purpose, it would be ideal if proper selection of p exactly represented nontrivial sums of Gaussian and uniform terms. Unfortunately the sum of any two independent GG random variables (GGRVs) is another GGRV only when $p = 2$ for each addend³ [52]. Nevertheless, the sum of independent GGRVs has many of the same properties as a GGRV, and can be well-approximated as a GGRV through a number of approximation methods. A simple approach from [53] matches the mean, variance, and kurtosis of the GG approximation to the corresponding moments of the true noise distribution as follows. Defining $\hat{V}$ as the GG approximation to $V = Z + W$, then since the uniform and Gaussian noise components are independent random variables, the mean and variance parameters of the GG noise approximation are simply given as $\mu_{\hat{V}} = \mu_W + \mu_Z$ and $\sigma_{\hat{V}}^2 = \sigma_Z^2 + \sigma_W^2$. The shape parameter approximation $\hat{p}_{\hat{V}}$ for the special case of uniform and Gaussian addends is computed as the unique solution to

$$\frac{\Gamma(1/\hat{p}_{\hat{V}})\Gamma(5/\hat{p}_{\hat{V}})}{\Gamma(3/\hat{p}_{\hat{V}})^2} = 3 - \frac{6}{5} \frac{1}{\left[12 \left(\frac{\sigma_Z}{\Delta}\right)^2 + 1\right]^2} \quad (21)$$

(see derivation in Appendix B). We thus see that $\hat{p}_{\hat{V}}$ depends on σ_Z/Δ , with the relationship plotted in Fig. 2. Solving (21) is fast, and the values of $\hat{p}_{\hat{V}}$ for a range of σ_Z/Δ values could be precomputed and stored in a table if necessary. A rough approximation and good initial value for a solver is $\hat{p}_{\hat{V}}^{(0)} = \max\{2, \Delta/\sigma_Z\}$.

To verify the quality of the generalized Gaussian approximation to the output noise distribution using the kurtosis match, Fig. 3 shows comparisons between the true density, computed numerically according to (15), and its GG approximation from (20). We test $\sigma_Z/\Delta = 0.004, 0.04$, and 0.4 , maintaining $\Delta = 1$ for consistency. For $\sigma_Z \ll \Delta$, the distribution is close to uniform, and at $\sigma_Z \approx \Delta$, the distribution is almost Gaussian. In the intermediate regime, however, the distribution combines attributes of each component, with the flat top of the uniform distribution and exponential tails of the Gaussian distribution. The GGD appears to be a good approximation of the true noise distribution, almost perfectly matching the shape behavior.

³The limiting distribution of the sum of i.i.d. GGRVs is Gaussian by the Central Limit Theorem [38, Chapter 5.4.2], but the sum of any finite number of GGRVs will only be approximately Gaussian unless each term is Gaussian.




Fig. 2: The value of $\hat{p}_{\tilde{V}}$ goes to infinity as σ_Z/Δ decreases and converges to $\hat{p}_{\tilde{V}} = 2$ as σ_Z/Δ increases, with convergence beginning around $\sigma_Z/\Delta = 1/3$ matching the anticipated behavior.




Fig. 3: The three plots show the noise PDF calculated numerically from the true density (15) (solid red) and via the GG approximation (20) (dashed black). The close agreement suggests the GGD is a good approximation for the noise.

B. Estimation

For i.i.d. samples of a GG distribution, the likelihood function is

$$\mathcal{L}(\{v_i\}_{i=1}^K; \mu, \sigma, p) = \prod_{i=1}^K f_{\tilde{V}}(v_i; \mu, \sigma, p). \quad (22)$$

By differentiating the log of (22) with respect to μ , the ML estimator $\hat{\mu}_{\text{GGML}}$ for the mean of a GGRV is given in [51] as the solution to

$$\sum_{i=1}^K \text{sgn}(y_i - \hat{\mu}_{\text{GGML}}) |y_i - \hat{\mu}_{\text{GGML}}|^{p-1} = 0, \quad (23)$$

and is shown to be asymptotically normal and efficient in K for $p \geq 2$, which is the regime of interest. The asymptotic variance of $\hat{\mu}_{\text{GGML}}$ normalized by Δ^2 is given by

$$\text{NVar}(\text{GGML}) = \frac{\beta(p)[(\sigma_Z/\Delta)^2 + 1/12]}{K}, \quad (24)$$




Fig. 4: As p increases beyond $p = 2$, $\beta(p)$ becomes much less than 1, implying that $\text{MSE}(\hat{\mu}_{\text{GGML}})$ is much lower than $\text{MSE}(\hat{\mu}_{\text{mean}})$.

where

$$\beta(p) = \frac{\Gamma^2(1/p)}{p^2 \Gamma\left(\frac{2p-1}{p}\right) \Gamma(3/p)}. \quad (25)$$

We notice that (24) decreases as $O(K^{-1})$, but the coefficient $\beta(p)$, which is plotted in Fig. 4, is much less than 1 for large p , suggesting that $\hat{\mu}_{\text{GGML}}$ should outperform $\hat{\mu}_{\text{mean}}$ for $\hat{p}_{\tilde{V}} > 2$. Since the GGD closely approximates the total noise distribution, it would be ideal if $\hat{\mu}_{\text{GGML}}$ reduced to a computationally simple estimator such as one based on order statistics for all p . Unfortunately, the ML estimator does not generally have a closed-form expression, except in special cases such as $p = 1, 2, \infty$ (an explicit expression has also recently been derived for $p = 4$ [54]), so iterative solution would again be necessary.

We have already observed that the ML estimators for $p = 1, 2, \infty$ all belong to a class of linear combinations of order statistics called L-estimates [55], which are attractive because they have closed-form definitions of the form

$$\hat{\mu}_X = \sum_{i=1}^K a_i Y_{(i)}. \quad (26)$$

We thus consider how to obtain the coefficients a_i for L-estimates that perform well for GG noise when p is not one of the special cases.

An effective L-estimate should weight the order statistics in accordance with the noise distribution. Notable past approaches include that of Lloyd, who derived the minimum-variance unbiased estimator among linear functions of the order statistics in [39]. This formulation is impractical, however, as it requires the correlations of the order statistics for a given distribution, which are often not known even for common special cases like the Gaussian distribution. Bovik et al. [56], [57] further specified the minimum variance unbiased L-estimate and then numerically computed results for several values of p from samples of a GGD with $K = 3$.

A number of approximations to Lloyd's formulation exist to more simply compute near-optimal coefficients for linear combinations of order statistics, including [58], [59]. Öten and de Figueiredo introduced one such method using Taylor expansion approximations to get around the difficulties of knowing distributions of order statistics [60]. This method does still require knowledge of the inverse CDF of the noise distribution, and while there is no closed form expression

for the GGD, the necessary values can be pre-computed numerically.

Simpler L-estimates have much longer histories, with consideration of trimming extreme or middle order statistics at least as old as [61] (credited to Gergonne in [62]), with the first known mathematical analysis by Daniell, who called such an estimate the “discard-average” [63], [64]. The method now known as the α -trimmed mean and popularized by Tukey [65], [66] avoids extensive computation of the weights by trimming a fixed fraction α from the extremes of the order statistics. Restrepo and Bovik defined a complementary α -“outer” trimmed-mean [67], which retains a fraction α of the data by trimming the middle order statistics and is suitable for distributions with short tails within the range from Gaussian to uniform distributions. They tabulated several instances of the trimmed mean for GGDs with multiple combinations of K and p .

Lastly, Beaulieu and Guo introduced an estimator specifically for the GGD but using nonlinear combinations of the order statistics [68]. The weighting of the order statistics depends on p via a heuristically-justified function and is shown to perform almost identically to $\hat{\mu}_{\text{GGML}}$. This estimator is unbiased and exactly matches the ML estimator for the special cases of $p = 2$ and ∞ .

In the following section, we consider three of the most computationally-efficient order statistics-based estimators to use for the GG approximation: the nearly-best L-estimate $\hat{\mu}_{\text{NB}}$ of [60], the trimmed-mean estimator $\hat{\mu}_\alpha$ modeled on [67], and the non-linear estimator $\hat{\mu}_{\text{NL}}$ of [68]. Each estimator takes the form of (26) with different computations of the coefficients a_i . While $\hat{\mu}_{\text{NL}}$ is specifically designed for use with GG noise, we modify the more general $\hat{\mu}_{\text{NB}}$ and $\hat{\mu}_\alpha$ to match the GG approximation. For $\hat{\mu}_{\text{NB}}$, we use the PDF and inverse CDF (computed numerically) of the GG approximation to determine the coefficients. One could alternatively compute the coefficients for $\hat{\mu}_{\text{NB}}$ directly from the true noise distribution in (15); however, additional numerical evaluation would be required for the inverse CDF, which we eschew in our search for computationally efficient estimators. There is no explicit distribution assumed by $\hat{\mu}_\alpha$, but we propose a choice of the trimmed fraction α based on the estimated $\hat{p}_{\hat{V}}$ value to implicitly link the estimator to the GGD.

IV. ESTIMATOR IMPLEMENTATIONS

A. ML Estimators

An EM algorithm for obtaining the quantized-sample ML estimate $\hat{\mu}_{\text{QML}}$ was introduced by Papadopoulos et al. [31, Appendix E]:

$$\hat{\mu}_{\text{QML}}^{(j+1)} = \hat{\mu}_{\text{QML}}^{(j)} + \frac{\sigma_Z}{K\sqrt{2\pi}} \sum_{i=1}^K m(u_i), \quad (27)$$

where

$$m(u_i) = \frac{\exp\left(-\frac{[u_i - \frac{\Delta}{2} - \hat{\mu}_{\text{QML}}^{(j)}]^2}{2\sigma_Z^2}\right) - \exp\left(-\frac{[u_i + \frac{\Delta}{2} - \hat{\mu}_{\text{QML}}^{(j)}]^2}{2\sigma_Z^2}\right)}{\Phi\left(\frac{u_i + \frac{\Delta}{2} - \hat{\mu}_{\text{QML}}^{(j)}}{\sigma_Z}\right) - \Phi\left(\frac{u_i - \frac{\Delta}{2} - \hat{\mu}_{\text{QML}}^{(j)}}{\sigma_Z}\right)}. \quad (28)$$

Zymnis et al. derived a gradient descent algorithm equivalent to that in (27) for the special case of a repeated scalar input and no mixing (i.e., the mixing matrix is a column of 1s) [49]. They also showed the negative log-likelihood to be convex, so the EM algorithm converges to the ML estimate. A good initialization is $\hat{\mu}_{\text{QML}}^{(0)} = \hat{\mu}_Q$, since the estimators are equal for $\sigma_Z = 0, \infty$. Since $\hat{\mu}_{\text{DML}}$ has the same formulation as $\hat{\mu}_{\text{QML}}$, the same algorithm also works for continuous-valued dithered measurements:

$$\hat{\mu}_{\text{DML}}^{(j+1)} = \hat{\mu}_{\text{DML}}^{(j)} + \frac{\sigma_Z}{K\sqrt{2\pi}} \sum_{i=1}^K m(y_i). \quad (29)$$

We initialize with $\hat{\mu}_{\text{DML}}^{(0)} = \hat{\mu}_{\text{mid}}$, since the midrange is known to be the ML estimator for $\sigma_Z = 0$. A solver for $\hat{\mu}_{\text{GGML}}$ was likewise initialized with $\hat{\mu}_{\text{GGML}}^{(0)} = \hat{\mu}_{\text{mid}}$.

B. Order Statistics-Based Estimators

To evaluate the GG noise approximation and find the best non-iterative estimator, we compared the three simplest estimators based on the order statistics: the nearly-best L-estimate, the α -outer mean, and the nonlinear combination from [68]. Since the GGD is symmetric, the coefficients of an unbiased order statistics-based estimator are defined symmetrically and only half must be uniquely computed. It is thus useful to define $M = \lfloor K/2 \rfloor$ and $N = \lceil K/2 \rceil$ using the floor and ceiling functions, respectively.

To derive the nearly-best L-estimate of [60]

$$\hat{\mu}_{\text{NB}} = \sum_{i=1}^K a_i^{\text{NB}} Y_{(i)}, \quad (30)$$

we first compute

$$b_1 = f_{\hat{V}}(c_1)[-2f_{\hat{V}}(c_1) + f_{\hat{V}}(c_2)], \quad (31a)$$

$$b_i = f_{\hat{V}}(c_i)[f_{\hat{V}}(c_{i-1}) - 2f_{\hat{V}}(c_i) + f_{\hat{V}}(c_{i+1})], \quad (31b)$$

$$i = 2, \dots, N-1,$$

$$b_N = f_{\hat{V}}(c_N)[f_{\hat{V}}(c_{N-1}) - f_{\hat{V}}(c_N)], \quad (31c)$$

where $c_i = F_{\hat{V}}^{-1}(i/(K+1))$, and $F_{\hat{V}}^{-1}$ is the inverse of the GG CDF. From this, the weights are derived for $i = 1, \dots, N$ as

$$a_i^{\text{NB}} = a_{K-i+1}^{\text{NB}} = \begin{cases} b_i / (2 \sum_{i=1}^N b_i), & K \text{ even;} \\ b_i / (b_N + 2 \sum_{i=1}^M b_i), & K \text{ odd.} \end{cases} \quad (32)$$

For the simulations in Python, the inverse CDF was numerically computed with the `stats.gennorm.ppf` GGD percentile function in `scipy`, as no closed-form expression exists.

For the α -outer mean estimate

$$\hat{\mu}_\alpha = \sum_{i=1}^K a_i^\alpha Y_{(i)}, \quad (33)$$

the order statistics' weights a_i^α are only given in [67] for a symmetric filter applied to an odd number of samples:

$$a_i^\alpha = a_{K-i+1}^\alpha = \begin{cases} \frac{1}{K\alpha}, & i \leq \lfloor \frac{1}{2}K\alpha \rfloor \\ \frac{\frac{1}{2}K\alpha - \lfloor \frac{1}{2}K\alpha \rfloor}{K\alpha}, & i = \lfloor \frac{1}{2}K\alpha \rfloor + 1, \\ \frac{K\alpha - 2\lfloor \frac{1}{2}K\alpha \rfloor}{K\alpha}, & i = \lfloor \frac{1}{2}K\alpha \rfloor + 1, \\ 0, & \text{otherwise.} \end{cases} \quad \alpha \in [0, 1 - 1/K]; \quad (34)$$

Since an even number of measurements is also possible, we similarly define, for all $\alpha \in [0, 1]$,

$$a_i^\alpha = a_{K-i+1}^\alpha = \begin{cases} \frac{1}{K\alpha}, & i \leq \lfloor \frac{1}{2}K\alpha \rfloor; \\ \frac{\frac{1}{2}K\alpha - \lfloor \frac{1}{2}K\alpha \rfloor}{K\alpha}, & i = \lfloor \frac{1}{2}K\alpha \rfloor + 1; \\ 0, & \text{otherwise.} \end{cases} \quad (35)$$

Note that the outer mean is equivalent to $\hat{\mu}_{\text{mean}}$ when $\alpha = 1$ and reduces to $\hat{\mu}_{\text{mid}}$ for $\alpha = 0$. To match the GGD behavior, we thus propose to define $\alpha = 2/\hat{p}_{\hat{V}}$, which yields the ML estimate for both $\hat{p}_{\hat{V}} = 2$ and $\hat{p}_{\hat{V}} = \infty$.

Finally, the nonlinear estimator of [68] is given as

$$\hat{\mu}_{\text{NL}} = \sum_{i=1}^K a_i^{\text{NL}} Y_{(i)}, \quad (36)$$

where the data-dependent coefficients are given for $i = 1, \dots, M$ by

$$a_i^{\text{NL}} = a_{K-i+1}^{\text{NL}} = \frac{1}{2} \frac{[Y_{(K-i+1)} - Y_{(i)}]^{p-2}}{\sum_{j=1}^M [Y_{(K-j+1)} - Y_{(j)}]^{p-2}}. \quad (37)$$

Note that if K is odd, the median term ($i = N$) is ignored, as it would correspond to a numerator of zero.

V. DITHER NOISE REGIMES

To better understand the dither noise behavior, we have previously described three regimes of the dither noise distribution, with Regimes I and III corresponding to approximately uniform and Gaussian noise, respectively. We have furthermore proposed the GGD with $p \in (2, \infty)$ as an approximation for the noise distribution in Regime II. However, the boundaries of these regions are imprecise, and we aim to more rigorously define them in this section. We first define ξ_1 and ξ_2 as the values of the ratio σ_Z/Δ separating the regimes such that the noise distribution is approximately uniform for $\sigma_Z/\Delta < \xi_1$, GG for $\xi_1 \leq \sigma_Z/\Delta < \xi_2$, and Gaussian for $\sigma_Z/\Delta \leq \xi_2$. In each regime, we have an expression for the expected MSE or asymptotic variance of the ML estimator, so we use the intersection or approximate point of convergence of these expressions to define ξ_1 and ξ_2 .

A. Defining ξ_1

We define ξ_1 as the value of σ_Z/Δ where $\text{NMSE}(\text{mid})$ and $\text{NVar}(\text{GGML})$ intersect, which from (14) and (24) is the solution to

$$\beta(\hat{p}_{\hat{V}})[(\sigma_Z/\Delta)^2 + 1/12] = \frac{K/2}{K^2 + 3K + 2} \quad (38)$$




Fig. 5: The value of ξ_1 substantially decreases and ξ_2 slowly increases as K increases, expanding Regime II. (a) A log-log-cubic fit can be used to compute a close approximation to ξ_1 for all K , while a log-log-linear fit suffices for $K > 20$. (b) The square root of a log-quadratic fit closely approximates ξ_2 .

for a given K . We remind the reader that $\hat{p}_{\hat{V}}$ is also dependent on σ_Z/Δ as shown in (21). Fig. 5a shows that ξ_1 decreases as K increases, since the probability of observing an “outlier” measurement due to the exponential tails increases with K , so a lower σ_Z/Δ value (i.e., with shorter tails) is needed for the midrange estimator to achieve nearly-optimal performance. The figure shows the exact values of ξ_1 computed by solving (38) as well as a log-log-cubic least-squares fit

$$\log \xi_1 \approx 0.0104(\log K)^3 - 0.1760(\log K)^2 + 0.0274(\log K) - 1.8511, \quad (39)$$

which can be used quickly to calculate an approximation for a desired value of K . Since the relationship appears fairly linear for $K > 20$, the simple log-log-linear fit

$$\log \xi_1 \approx -0.9301(\log K) - 0.1963, \quad (40)$$

which can be rewritten as $\xi_1 \approx 0.8217/K^{0.9301}$, is also useful for quick computation. The natural logarithm is used in each case.

B. Defining ξ_2

Since $\text{NMSE}(\text{mean})$ and $\text{NVar}(\text{GGML})$ both have $1/K$ factors, they converge where $\beta(\hat{p}_{\hat{V}}) = 1$, which is only the case for $\hat{p}_{\hat{V}} = 2$. This suggests that equality requires the noise to be exactly Gaussian, which only occurs for $\sigma_Z/\Delta \rightarrow \infty$. Instead, we can look for a point where $\text{NVar}(\text{GGML})$ and $\text{NMSE}(\text{mean})$ can reasonably be considered to have converged (i.e., the GG is close enough to a Gaussian). We propose that a reasonable definition of ξ_2 is the value of σ_Z/Δ that minimizes $\text{NMSE}(\text{Q})$, the expected normalized MSE of $\hat{\mu}_{\text{Q}}$. Intuitively, as σ_Z/Δ increases from ξ_2 , the Gaussian variance will dominate for both quantized and dithered measurements, so that the effect of the quantization error is negligible, whether signal-independent for dithered measurements or signal-dependent without dither. Thus the point at which $\text{NMSE}(\text{Q})$ is minimized indicates where the Gaussian variance begins to dominate and is a reasonable place to consider a

GG approximation to be sufficiently Gaussian. We derive in Appendix C that NMSE(Q) is given as

$$\begin{aligned} \text{NMSE(Q)} &= \mathbb{E}[(\hat{\mu}_Q - \mu_X)^2]/\Delta^2 \\ &= \frac{1}{12} + \frac{1}{K} \int_{-1/2}^{1/2} \sum_{m=-M}^M m^2 \Psi(m, \mu_X) d\mu_X \\ &\quad + \frac{K-1}{K} \int_{-1/2}^{1/2} \left(\sum_{m=-M}^M m \Psi(m, \mu_X) \right)^2 d\mu_X \\ &\quad - 2 \int_{-1/2}^{1/2} \mu_X \sum_{m=-M}^M m \Psi(m, \mu_X) d\mu_X, \end{aligned} \quad (41)$$

where

$$\Psi(m, \mu_X) = \Phi\left(\frac{m + 1/2 - \mu_X}{\sigma_Z/\Delta}\right) - \Phi\left(\frac{m - 1/2 - \mu_X}{\sigma_Z/\Delta}\right). \quad (42)$$

Defining

$$\xi_2 = \arg \min_{\sigma_Z/\Delta} \mathbb{E}[(\hat{\mu}_Q - \mu_X)^2]/\Delta^2 \quad (43)$$

and solving via a Nelder-Mead algorithm [69] and numerical integration, we show in Fig. 5b that the value of ξ_2 changes only slightly as a function of K . This range of values is notably very close to the value $\sigma_Z/\Delta = 1/2$ recommended by Vardeman and Lee [6], or the value $\sigma_Z/\Delta = 1/3$ at which Moschitta et al. suggest that the loss of information from quantizing samples of a Gaussian distribution becomes negligible in estimation of the mean [7]. For quick computation, ξ_2 can be approximated by the square root of a log-quadratic fit:

$$\xi_2 \approx \sqrt{-0.000756(\log K)^2 + 0.328 \log K}. \quad (44)$$

We notice that the Regime boundary definitions are inconsistent for $K < 3$, as $\xi_1 > \xi_2$; however, the Regimes are meaningless for $K = 1$ or 2 anyway, as symmetric order statistics-based estimators (e.g., mean, median, midrange) are all equivalent for such small numbers of measurements, so there is no advantage to distinguishing between noise distributions. We notice also that since ξ_1 decreases monotonically and ξ_2 increases monotonically with K , Regime II grows as K increases, since small mismatches between the assumed and true PDFs become easier to observe. Intuitively, ξ_1 decreases much faster than ξ_2 increases because the difference between a PDF with finite support ($\sigma_Z/\Delta = 0$) and one with infinite support ($\sigma_Z/\Delta > 0$) is more significant for large K than the difference between finite σ_Z/Δ (e.g., GG approximation with $\hat{p}_{\tilde{V}} > 2$ and $\sigma_Z/\Delta \rightarrow \infty$ corresponding to $\hat{p}_{\tilde{V}} = 2$).

VI. NUMERICAL RESULTS

Monte Carlo simulations were performed to compare the NMSE performance of the generalized Gaussian and order statistics-based estimators ($\hat{\mu}_{\text{NB}}$, $\hat{\mu}_{\text{NL}}$, $\hat{\mu}_{\alpha}$) against the ML estimators ($\hat{\mu}_{\text{DML}}$, $\hat{\mu}_{\text{GGML}}$) and the conventional sample mean ($\hat{\mu}_{\text{mean}}$) and midrange ($\hat{\mu}_{\text{mid}}$). Estimates were also computed applying the sample mean ($\hat{\mu}_Q$) and ML estimator ($\hat{\mu}_{\text{QML}}$) to the quantized data to determine under which conditions subtractive dithering actually provides an advantage. As in the motivating example in Section II, for each Monte Carlo

trial, μ_X was chosen uniformly at random from $[-\Delta/2, \Delta/2]$, and K samples of signal noise $Z \sim \mathcal{N}(0, \sigma_Z^2)$ and dither $D \sim \mathcal{U}(-\Delta/2, \Delta/2)$ were generated for (3). The quantization bin size was maintained at $\Delta = 1$ throughout. The normalized MSE was computed for $T = 20,000$ trials.

A. Normalized MSE vs. σ_Z/Δ

We begin by discussing the plots in Fig. 6 of NMSE as a function of σ_Z/Δ for $K = 5, 25$, and 125. Because nine separate estimators and four NMSE bounds are displayed in each plot, we acknowledge that the figures can be difficult to follow due to the overlapping curves. A flowchart is included in Fig. 7 that summarizes the results and provides a decision-making process for whether to use dither and which estimator to choose.

1) *Generalized Gaussian Estimators*: The GG-based estimators ($\hat{\mu}_{\text{NB}}$, $\hat{\mu}_{\text{NL}}$, $\hat{\mu}_{\alpha}$, $\hat{\mu}_{\text{GGML}}$) have effectively identical performance and match that of $\hat{\mu}_{\text{DML}}$. The actual differences in performance vary on the order of a few percent over large ranges of K and σ_Z/Δ , compared to the orders of magnitude differences for the mean and midrange. The negligible performance difference further validates approximating the total noise with the GGD. For this same reason, we collectively discuss $\hat{\mu}_{\text{DML}}$ and the GG-based estimators in the following sections.

The GG-based estimators meet or exceed the performance of all other estimators for all σ_Z/Δ and for all K . More specifically, the GG estimators converge to and match the performance of the midrange in Regime I and likewise converge to and match the performance of the mean in Regime III. In Regime II, the GG estimators outperform both the mean and the midrange. Thus, a GG estimator should be the default estimator choice for any σ_Z/Δ .

Given the approximate equivalence of the GG estimators, we argue that the trimmed-mean $\hat{\mu}_{\alpha}$ is the best choice of general-purpose estimator for dithered data. The other estimators either require iterative solvers ($\hat{\mu}_{\text{DML}}$, $\hat{\mu}_{\text{GGML}}$), rely on numerical computation for the GG inverse CDF ($\hat{\mu}_{\text{NB}}$), or are data-dependent ($\hat{\mu}_{\text{NL}}$). On the other hand, $\hat{\mu}_{\alpha}$ has a simple closed-form solution that can be tabulated if needed.

2) *Performance by Regime—Dithered Measurements*: The plots in Fig. 6 validate the concept of three distinct regimes of noise behavior. In the plots, the approximate regime boundaries are computed to be $\xi_1 = \{0.1098, 3.85 \times 10^{-2}, 9.56 \times 10^{-3}\}$ and $\xi_2 = \{0.2296, 0.3132, 0.3737\}$ for $K = \{5, 25, 125\}$, respectively, confirming that Regime II expands as K increases. In Regime I, the NMSE performance of all estimators on the dithered data is basically flat and equal to NMSE(mid). This suggests that for a practical system where σ_Z/Δ can be tuned, once the system is operating in Regime I (dependent on a fixed K), there is no benefit from further decreasing σ_Z/Δ ; performance can only be improved by increasing K . In Regime II, the GG-based estimators approach NCRB(μ_X), especially for large K . We note that while NVar(GGML) and NCRB(μ_X) are close in Regime II, NCRB(μ_X) is a tighter bound, as it is based on the true noise distribution, although NVar(GGML) may be easier to




Fig. 6: The performance of the estimators is evaluated for $K =$ (a) 5, (b) 25, and (c) 125 to show the range of behavior as σ_Z/Δ varies. The ML estimator for dithered measurements $\hat{\mu}_{\text{DML}}$ and the estimators based on the GGD ($\hat{\mu}_{\text{GGML}}$, $\hat{\mu}_{\text{NB}}$, $\hat{\mu}_{\text{NL}}$, and $\hat{\mu}_{\alpha}$) achieve the lowest NMSE for each σ_Z/Δ regime (curves are overlapping). Results are shown for 20000 Monte Carlo trials.




Fig. 7: The results of our Monte Carlo simulations lead to a simplified decision process for when and how to use dither. If $\sigma_Z/\Delta > \xi_2$ ($\approx 1/3$), there is no benefit to using anything but $\hat{\mu}_Q$ applied to the quantized measurements, whereas dither leads to reduce estimation error when $\sigma_Z/\Delta \leq 1/3$. If subtractive dithering is not possible, the best performance can be achieved by adding Gaussian noise to set $\sigma_Z/\Delta \approx 1/3$ and applying $\hat{\mu}_{\text{QML}}$ (although $\hat{\mu}_Q$ can be used if simplicity is required). However, larger performance improvements can be achieved with a subtractively-dithered quantizer. For K subtractively-dithered measurements, compute ξ_1 from either (39) or (40) to determine whether to use $\hat{\mu}_{\text{mid}}$ (in Regime I) or $\hat{\mu}_{\alpha}$ (in Regime II).

compute for a rough estimate of performance. In Regime III, the NMSE performance of all estimators on the dithered data is equal to $\text{NMSE}(\text{mean})$. In both Regimes II and III, the NMSE decreases as σ_Z/Δ decreases. Performance likewise improves with increasing K .

3) Performance by Regime—Quantized Measurements:

While the three Regimes were technically defined for dithered measurements in particular, they are also informative of the behavior of estimators applied to quantized measurements. In Regime I, σ_Z/Δ is so small that, unless μ_X lies on the boundary between quantization bins, all measurements are quantized to the same value. As a result, the NMSE of both $\hat{\mu}_Q$ and $\hat{\mu}_{\text{QML}}$ is dominated by the squared bias term, which

is $1/12$ (the variance is zero). Further decreasing σ_Z/Δ or increasing K provides no benefit.

In Regime II, σ_Z/Δ is large enough that there is often some variation in the measurements due to signal even without the addition of dither. This phenomenon is sometimes referred to in the literature as self-dithering, equivalent to adding nonsubtractive Gaussian dither to a constant signal μ_X [70]. Within Regime II, both $\hat{\mu}_Q$ and $\hat{\mu}_{\text{QML}}$ improve as σ_Z/Δ increases because the increased signal variation reduces the bias term of the NMSE faster than the variance increases. The NMSE is minimized for $\hat{\mu}_Q$ by definition at ξ_2 , and then the NMSE increases as σ_Z/Δ increases in Regime III. This suggests that if σ_Z/Δ is small and subtractive dither cannot be used, then quantized measurements benefit from adding nonsubtractive Gaussian dither such that $\sigma_Z/\Delta = \xi_2$, which is approximately $1/3$. It is in Regime II that $\hat{\mu}_{\text{QML}}$ shows the largest improvement in performance over $\hat{\mu}_Q$, with the ML estimator accounting for the form of the signal variation for quantized measurements.

In Regime III, the NMSE of $\hat{\mu}_Q$ and $\hat{\mu}_{\text{QML}}$ matches that of the best estimators applied to dithered data. Clearly, σ_Z/Δ is large enough that even the quantized measurements contain sufficient information about the signal variation. This suggests that dither provides no benefit in Regime III, since equal performance can be achieved without dither. Again for both Regimes II and III, the NMSE decreases as K is increased.

B. Order Statistics-Based Estimator Coefficients

To better understand why the order statistics-based estimators have essentially identical performance, in Fig. 8 we plot the coefficients a_i from (26) for each estimator. The top row shows example measurements for $K = 20$, $\Delta = 1$, and $\sigma_Z/\Delta = 0.004$, 0.04 , and 0.4 , respectively, with the samples spreading out as the Gaussian variance increases. The second row of plots depicts the resulting coefficients for $\hat{\mu}_{\text{NB}}$, $\hat{\mu}_{\text{NL}}$, and $\hat{\mu}_{\alpha}$ using the estimated value $\hat{p}_{\tilde{V}}$. Fig. 8d shows the




Fig. 8: Example dithered measurements are shown in (a-c) for $\sigma_Z/\Delta = 0.004, 0.04$, and 0.4 with $K = 20$. Plots (d-f) show the resulting coefficient values for the GG estimators given the estimated value of $\hat{p}_{\hat{V}}$ above. In (g), $\sigma_Z/\Delta = 0.04$ and $K = 100$, highlighting how the coefficients change as K increases. Note that the coefficients of the order statistics for the NL estimator depend on the measured data sequence shown above.

coefficients are equivalent to those of $\hat{\mu}_{\text{mid}}$ for small σ_Z/Δ . In Figs. 8e and 8f, the coefficients of the various estimators are no longer identical. However, the coefficients follow the same trends for each estimator, with zero weight on the middle order statistics for small σ_Z/Δ and more evenly-distributed weights as the noise model approaches a Gaussian. We note that the coefficients for $\hat{\mu}_{NL}$ vary depending on the particular set of measurements shown in the top row, and that different sample realizations can result in coefficients more or less similar to those of $\hat{\mu}_{NB}$ and $\hat{\mu}_{\alpha}$. To show the behavior of the coefficients as K increases, we also plot $\{a_i\}_{i=1}^K$ for $K = 100$ and $\sigma_Z/\Delta = 0.04$ in Fig. 8g. This plot underscores that the coefficients for $\hat{\mu}_{\alpha}$ are basically indicators of the most significant non-zero coefficients of $\hat{\mu}_{NB}$ and $\hat{\mu}_{NL}$. Using the simple formulation of $\hat{\mu}_{\alpha}$ as a guide, in the limit as $K \rightarrow \infty$, only $K\alpha$ coefficients would have nonzero weight. Since the performance of all three order statistics-based estimators is similar, this further suggests that the selection of which order statistics are used is more important than exactly how much they are weighted.

C. Normalized MSE vs. K

To better understand how the number of measurements affects the estimators' performance, we plot results for three fixed values of σ_Z/Δ in each regime (0.004, 0.04, 0.4) while varying K in Fig. 9.

In Fig. 9a, $\hat{\mu}_{\text{mid}}$ follows NMSE(mid) as expected for Regime I until $K \approx 200$. At that point, the NMSE of the

midrange begins to diverge, with slower improvement as K increases. Similarly, the GG estimators follow NMSE(mid) until $K \approx 200$ and then switch to NCRB(μ_X). This suggests that $\sigma_Z/\Delta = 0.004$ is in Regime I for $K < 200$ and in Regime II for $K > 200$. This switch between regimes occurs near the intersection of NMSE(mid) and NVar(GGML) as a function of K , further validating these bounds as useful demarcations of estimator performance. For all K in the plotted range, the midrange and GG estimators outperform the mean. The quantized estimators show almost no improvement as K increases.

In Fig. 9b, the midrange performance is similar to that in Fig. 9a, with $\hat{\mu}_{\text{mid}}$ following NMSE(mid) until the intersection of NMSE(mid) and NVar(GGML) and then improving more slowly as a function of the number of measurements, eventually being outperformed by $\hat{\mu}_{\text{mean}}$ for large K . The GG estimators likewise follow NMSE(mid) for small K and switch to following NCRB(μ_X) after the intersection. For large K , the NMSE of the GG estimators is a constant factor lower than that of $\hat{\mu}_{\text{mean}}$, with this factor approximately given by $\beta(\hat{p}_{\hat{V}})$. The NMSE of the quantized estimators decreases slowly as K increases, with marginally better performance for $\hat{\mu}_{QML}$ than $\hat{\mu}_Q$.

Figures 9a and 9b help answer the question of how the order statistics-based estimators “between” the midrange and the mean would perform as a function of K . The results suggest that these estimators ultimately have $O(K^{-1})$ NMSE reduction, although this reduction is faster for small values of




Fig. 9: The performance of the order statistics estimators is evaluated for $\sigma_Z/\Delta =$ (a) 0.004, (b) 0.04, and (c) 0.4 to show the full range of behavior as the number of measurements K increases. The plots show the MSE normalized to $\Delta = 1$ from 20000 trials per data point. Dashed lines show the theoretical NMSE of the mean and midrange, and the asymptotic variance of the ML estimator of the GGD mean. The ML estimator for dithered measurements $\hat{\mu}_{\text{DML}}$ and the estimators based on the GGD ($\hat{\mu}_{\text{GGML}}$, $\hat{\mu}_{\text{NB}}$, $\hat{\mu}_{\text{NL}}$, and $\hat{\mu}_{\alpha}$) achieve the lowest NMSE for all K (curves are overlapping).

K .

In Fig. 9c where $\sigma_Z/\Delta = 0.4$, the noise can be sufficiently described as Gaussian for $K < 359$ because $\sigma_Z/\Delta > \xi_2$ as shown in Fig. 5b; however, $\sigma_Z/\Delta < \xi_2$ for $K \geq 359$, so the noise distribution transitions from Regime III to Regime II. The midrange has poor performance for all K , while the other dithered estimators and the quantized estimators have essentially identical performance for $K < 359$. Those estimators follow NMSE(mean), NVar(GGML), and NCRB(μ_X), which have converged. In Regime III, it is clear that there is no benefit to using dither, as there is minimal improvement in performance even for large K . In fact, implementing a dithered quantizer is likely more complicated in practice and is discouraged for Regime III.

VII. CONCLUSION

This work studied the task of estimating the mean of a Gaussian signal from quantized measurements. By applying subtractive dither to the measurement process, the noise becomes signal-independent but no longer has a Gaussian distribution. We showed that the generalized Gaussian distribution is a close and useful approximation for the Gaussian plus uniform total noise distribution. Estimators using the generalized Gaussian approximation effectively match the performance of the ML estimator for the total noise, which is a significant improvement over the conventional mean and midrange estimators in Regime II. Due to its computational simplicity and efficient performance, we recommend the trimmed mean $\hat{\mu}_{\alpha}$. From further comparison against estimators for quantized measurements, we determined simple design rules for deciding whether and how to use quantized measurements. In short, there is value in using dither in Regimes I and II, and a GG-based estimator should be used in Regime II. Future work will address variations on the measurement model, including

non-Gaussian signal distributions and different dither implementations.

APPENDIX A CRAMÉR-RAO BOUND

The Cramér-Rao Bound is a lower bound on the variance of an unbiased estimator [42], given by

$$\text{CRB}(\mu_X) = 1/I(\mu_X),$$

where $I(\mu_X)$ is the Fisher information computed as

$$\begin{aligned} I(\mu_X) &= \mathbb{E} \left[\left(\frac{\partial \log f_Y(y; \mu_X, \sigma_Z, \Delta)}{\partial \mu_X} \right)^2 \right] \\ &\stackrel{(a)}{=} \int \frac{\left(\frac{\partial}{\partial \mu_X} f_Y(y; \mu_X, \sigma_Z, \Delta) \right)^2}{f_Y(y; \mu_X, \sigma_Z, \Delta)} dy \\ &\stackrel{(b)}{=} \frac{1}{\sigma_Z^2 \Delta} \int \frac{\left[\phi \left(\frac{v - \Delta/2}{\sigma_Z} \right) - \phi \left(\frac{v + \Delta/2}{\sigma_Z} \right) \right]^2}{\Phi \left(\frac{v + \Delta/2}{\sigma_Z} \right) - \Phi \left(\frac{v - \Delta/2}{\sigma_Z} \right)} dv \\ &\stackrel{(c)}{=} \frac{1}{\sigma_Z^2} \int \frac{\left[\phi \left(\frac{u - 1/2}{\sigma_Z/\Delta} \right) - \phi \left(\frac{u + 1/2}{\sigma_Z/\Delta} \right) \right]^2}{\Phi \left(\frac{u + 1/2}{\sigma_Z/\Delta} \right) - \Phi \left(\frac{u - 1/2}{\sigma_Z/\Delta} \right)} du, \end{aligned} \quad (45)$$

where step (a) uses the definition of expectation and the chain rule, (b) differentiates (15) with respect to μ_X for $v = y - \mu_X$, and (c) changes variables to $u = v/\Delta$. Normalizing by Δ^2 removes the separate dependence on σ_Z or Δ , so we define the normalized CRB as

$$\begin{aligned} \text{NCRB}(\mu_X) &= \text{CRB}(\mu_X)/\Delta^2 \\ &= \frac{(\sigma_Z/\Delta)^2}{\int \frac{\left[\phi \left(\frac{u - 1/2}{\sigma_Z/\Delta} \right) - \phi \left(\frac{u + 1/2}{\sigma_Z/\Delta} \right) \right]^2}{\Phi \left(\frac{u + 1/2}{\sigma_Z/\Delta} \right) - \Phi \left(\frac{u - 1/2}{\sigma_Z/\Delta} \right)} du}. \end{aligned} \quad (46)$$

Finally, Fisher information is additive for independent observations, so for K independent samples, the lower bound on the NCRB is $1/K$ times that for one observation.

APPENDIX B KURTOSIS MATCHING

The *kurtosis* of a random variable B is the standardized fourth central moment [71], defined as

$$\kappa(B) = \frac{\mathbb{E}[(B - \mu_B)^4]}{\{\mathbb{E}[(B - \mu_B)^2]\}^2} = \frac{\mu_4(B)}{\sigma_B^4}. \quad (47)$$

The *excess kurtosis* $\gamma(B) = \kappa(B) - 3$ is often used to simplify computations. Define $A = B + C$, where B and C are independent random variables. The kurtosis of the sum can be computed by expanding (47) as follows:

$$\begin{aligned} \kappa(A) &= \frac{\mathbb{E}[(A - \mu_A)^4]}{\{\mathbb{E}[(A - \mu_A)^2]\}^2} = \frac{\mathbb{E}\{[(B - \mu_B) + (C - \mu_C)]^4\}}{\{\mathbb{E}\{[(B - \mu_B) + (C - \mu_C)]^2\}\}^2} \\ &= \frac{\mu_4(B) + \mu_4(C) + 6\sigma_B^2\sigma_C^2}{(\sigma_B^2 + \sigma_C^2)^2}, \end{aligned}$$

where independence eliminates the odd cross terms. Then the excess kurtosis is

$$\gamma(A) = \frac{\sigma_B^4\gamma(B) + \sigma_C^4\gamma(C)}{\sigma_A^4}. \quad (48)$$

The kurtosis of Gaussian and uniform random variables is well-known and straightforward to compute from the definition; the excess kurtosis is 0 for a Gaussian and $-6/5$ for a uniform distribution. From [53], we have that the excess kurtosis⁴ of a GGRV V with shape parameter p_v is

$$\gamma(V) = \frac{\Gamma(1/p_v)\Gamma(5/p_v)}{[\Gamma(3/p_v)]^2} - 3. \quad (49)$$

To fit the GGD to the sum of uniform and Gaussian random variables, we set the kurtosis of the approximation to match the kurtosis of the sum using (48)

$$\begin{aligned} \frac{\Gamma(1/p_v)\Gamma(5/p_v)}{[\Gamma(3/p_v)]^2} &= 3 + \frac{\sigma_Z^4 \cdot 0 + \sigma_W^4(-6/5)}{(\sigma_W^2 + \sigma_Z^2)^2} \\ &= 3 - \frac{6}{5} \frac{1}{[1 + 12(\frac{\sigma_Z}{\Delta})^2]^2}, \end{aligned} \quad (50)$$

where $\sigma_W^2 = \Delta^2/12$.

APPENDIX C MEAN SQUARED ERROR OF $\hat{\mu}_Q$

We use iterated expectation to compute the MSE of $\hat{\mu}_Q$ as

$$\mathbb{E}[(\hat{\mu}_Q - \mu_X)^2] = \mathbb{E}[\mathbb{E}[(\hat{\mu}_Q - \mu_X)^2 | \mu_X]], \quad (51)$$

with no prior knowledge on the true value so that we assume $\mu_X \sim \mathcal{U}[-\Delta/2, \Delta/2]$ within a bin. Define a function $g: \mathbb{R} \rightarrow \mathbb{R}$ as $g(x) := \mathbb{E}[(\hat{\mu}_Q - \mu_X)^2 | \mu_X = x]$, then

$$\begin{aligned} g(x) &= \mathbb{E} \left[\left(\frac{1}{K} \sum_{i=1}^K q(x + Z_i) - x \right)^2 \right] \\ &= x^2 + \frac{1}{K^2} \left(\sum_{i=1}^K \mathbb{E}[(q(x + Z_i))^2] \right. \\ &\quad \left. + \sum_{i=1}^K \sum_{j \neq i}^K \mathbb{E}[q(x + Z_i)] \mathbb{E}[q(x + Z_j)] \right) \\ &\quad - \frac{2x}{K} \sum_{i=1}^K \mathbb{E}[q(x + Z_i)] \\ &= x^2 + \frac{1}{K} \mathbb{E}[(q(x + Z))^2] + \frac{K-1}{K} (\mathbb{E}[q(x + Z)])^2 \\ &\quad - 2x \mathbb{E}[q(x + Z)]. \end{aligned} \quad (52)$$

Using the definition

$$\Psi(m, x) = \Phi\left(\frac{m + 1/2 - x}{\sigma_Z/\Delta}\right) - \Phi\left(\frac{m - 1/2 - x}{\sigma_Z/\Delta}\right), \quad (53)$$

note that

$$\begin{aligned} \mathbb{E}[q(x + Z)] &= \lim_{M \rightarrow \infty} \sum_{m=-M}^M m \Delta \int_{m\Delta-\Delta/2}^{m\Delta+\Delta/2} \frac{1}{\sigma_Z} \phi\left(\frac{z-x}{\sigma_Z}\right) dz \\ &\approx \Delta \sum_{m=-M}^M m \Psi(m, x) \end{aligned}$$

for some large number M . Similarly,

$$\mathbb{E}[(q(x + Z))^2] \approx \Delta^2 \sum_{m=-M}^M m^2 \Psi(m, x).$$

The MSE normalized by Δ^2 then follows as

$$\begin{aligned} \mathbb{E}[(\hat{\mu}_Q - \mu_X)^2]/\Delta^2 &= \frac{1}{12} + \frac{1}{K} \int_{-1/2}^{1/2} \sum_{m=-M}^M m^2 \Psi(m, x) dx \\ &\quad + \frac{K-1}{K} \int_{-1/2}^{1/2} \left(\sum_{m=-M}^M m \Psi(m, x) \right)^2 dx \\ &\quad - 2 \int_{-1/2}^{1/2} x \sum_{m=-M}^M m \Psi(m, x) dx. \end{aligned} \quad (54)$$

ACKNOWLEDGMENT

The authors would like to thank Dr. Yanting Ma for assistance with the derivations in Appendix C and Dr. Donggeek Shin for helpful discussions on the role of nonsubtractive dither in improving estimation accuracy. Computing resources provided by Boston University's Research Computing Services are also gratefully appreciated.

⁴Note that the definition of kurtosis in [53] corresponds to the excess kurtosis in this work.

REFERENCES

- [1] N. F. Gjeddebaek, "Contribution to the study of grouped observations. IV. Some comments on simple estimates," *Biometrics*, vol. 15, no. 3, pp. 433–439, Sep. 1959.
- [2] R. A. Fisher, "On the mathematical foundations of theoretical statistics," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 222, no. 594–604, pp. 309–368, Jan. 1922.
- [3] D. Shin, F. Xu, D. Venkatraman, R. Lussana, F. Villa, F. Zappa, V. K. Goyal, F. N. Wong, and J. H. Shapiro, "Photon-efficient imaging with a single-photon camera," *Nature Commun.*, vol. 7, p. 12046, 2016.
- [4] J. Rapp, R. M. A. Dawson, and V. K. Goyal, "Dither-enhanced lidar," in *Imaging and Applied Optics*, Orlando, FL, Jun. 2018, p. JW4A.38.
- [5] —, "Improving lidar depth resolution with dither," in *Proceedings of the IEEE International Conference on Image Processing*, Oct. 2018, pp. 1553–1557.
- [6] S. B. Vardeman and C.-S. Lee, "Likelihood-based statistical estimation from quantized data," *IEEE Trans. Instrum. Meas.*, vol. 54, no. 1, pp. 409–414, Feb. 2005.
- [7] A. Moschitta, J. Schoukens, and P. Carbone, "Information and statistical efficiency when quantizing noisy DC values," *IEEE Trans. Instrum. Meas.*, vol. 64, no. 2, pp. 308–317, Feb. 2015.
- [8] S. B. Vardeman, "Sheppard's correction for variances and the "quantization noise model,"" *IEEE Trans. Instrum. Meas.*, vol. 54, no. 5, pp. 2117–2119, Oct. 2005.
- [9] W. F. Sheppard, "On the calculation of the average square, cube, of a large number of magnitudes," *J. Roy. Statist. Soc.*, vol. 60, no. 3, pp. 698–703, Sep. 1897.
- [10] —, "On the calculation of the most probable values of frequency-constants, for data arranged according to equidistant divisions of a scale," *Proc. London Math. Soc.*, vol. 29, pp. 353–380, 1898.
- [11] M. G. Kendall, "The conditions under which Sheppard's corrections are valid," *J. Roy. Statist. Soc.*, vol. 101, no. 3, pp. 592–605, 1938.
- [12] D. F. Heitjan, "Inference from grouped continuous data: A review," *Statist. Sci.*, vol. 4, no. 2, pp. 164–179, 1989.
- [13] Z. Bai, S. Zheng, B. Zhang, and G. Hua, "Statistical analysis for rounded data," *J. Statist. Plan. Infer.*, vol. 139, pp. 2526–2542, 2009.
- [14] R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE J. Sel. Areas Comm.*, vol. 17, no. 4, pp. 539–550, Apr. 1999.
- [15] D. V. O'Connor and D. D. Phillips, *Time-Related Single Photon Counting*. Academic Press, 1984.
- [16] B. Widrow, "Statistical analysis of amplitude-quantized sampled-data systems," *Transactions of the American Institute of Electrical Engineers, Part II: Applications and Industry*, vol. 79, no. 6, pp. 555–568, 1961.
- [17] A. B. Sripad and D. L. Snyder, "A necessary and sufficient condition for quantization errors to be uniform and white," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 25, no. 5, pp. 442–448, Oct. 1977.
- [18] T. A. C. M. Claassen and A. Jongpier, "Model for the power spectral density of quantization noise," *IEEE Trans. Acoust. Speech Signal Process.*, vol. ASSP-29, no. 4, pp. 914–917, Aug. 1981.
- [19] H. Viswanathan and R. Zamir, "On the whiteness of high-resolution quantization errors," *IEEE Trans. Inform. Theory*, vol. 47, no. 5, pp. 2029–2038, Jul. 2001.
- [20] D. Marco and D. L. Neuhoff, "The validity of the additive noise model for uniform scalar quantizers," *IEEE Trans. Inform. Theory*, vol. 51, no. 5, pp. 1739–1755, May 2005.
- [21] L. R. A. MacColl, *Fundamental Theory of Servomechanisms*. D. Van Nostrand Company, inc., 1945.
- [22] T. Ishikawa, "Linearization of contactor control systems by external dither signals," Electronics Laboratories, Stanford University, Stanford, CA, Tech. Rep., 1960.
- [23] W. M. Goodall, "Television by pulse code modulation," *Bell Syst. Tech. J.*, vol. 30, no. 1, pp. 33–49, Jan. 1951.
- [24] L. Roberts, "Picture coding using pseudo-random noise," *IEEE Trans. Inform. Theory*, vol. 8, no. 2, pp. 145–154, Feb. 1962.
- [25] N. S. Jayant and L. R. Rabiner, "The application of dither to the quantization of speech signals," *Bell Syst. Tech. J.*, vol. 51, no. 6, pp. 1293–1304, Jul. 1972.
- [26] G. G. Furman, "Improving performance of quantizer feedback systems by use of external dither," M.S. Thesis, Massachusetts Institute of Technology, 1959.
- [27] R. C. Jaffe, "Causal and statistical analyses of dithered systems containing three-level quantizers," M.S. Thesis, Massachusetts Institute of Technology, 1959.
- [28] L. Schuchman, "Dither signals and their effect on quantization noise," *IEEE Trans. Commun.*, vol. 12, no. 4, pp. 162–165, Dec. 1964.
- [29] R. M. Gray and T. G. Stockham, "Dithered quantizers," *IEEE Trans. Inform. Theory*, vol. 39, no. 3, pp. 805–812, 1993.
- [30] R. A. Wannamaker, S. P. Lipshitz, J. Vanderkooy, and J. N. Wright, "A theory of nonsubtractive dither," *IEEE Trans. Signal Process.*, vol. 48, no. 2, pp. 499–516, 2000.
- [31] H. C. Papadopoulos, G. W. Wornell, and A. V. Oppenheim, "Sequential signal encoding from noisy measurements using quantizers with dynamic bias control," *IEEE Trans. Inform. Theory*, vol. 47, no. 3, pp. 978–1002, 2001.
- [32] D. Rousseau and G. V. Anand, "Nonlinear estimation from quantized signals: Quantizer optimization and stochastic resonance," in *Third International Symposium on Physics in Signal and Image Processing*, 2003, pp. 89–92.
- [33] A. Ribeiro and G. B. Giannakis, "Bandwidth-constrained distributed estimation for wireless sensor networks - part I: Gaussian case," *IEEE Trans. Signal Process.*, vol. 54, no. 3, pp. 1131–1143, Mar. 2006.
- [34] O. Dabeer and A. Karnik, "Signal parameter estimation using 1-bit dithered quantization," *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5389–5405, Dec. 2006.
- [35] J. Fang and H. Li, "Distributed adaptive quantization for wireless sensor networks: From delta modulation to maximum likelihood," *IEEE Trans. Signal Process.*, vol. 56, no. 10, pp. 5246–5257, Oct. 2008.
- [36] R. C. Farias and J. M. Brossier, "Scalar quantization for estimation: From an asymptotic design to a practical solution," *IEEE Trans. Signal Process.*, vol. 62, no. 11, pp. 2860–2870, Jun. 2014.
- [37] C. Gianelli, L. Xu, J. Li, and P. Stoica, "One-bit compressive sampling with time-varying thresholds: Maximum likelihood and the Cramér-Rao bound," in *Conf. Rec. Asilomar Conf. on Signals, Syst. & Computers*, Nov. 2016, pp. 399–403.
- [38] A. M. Mood, F. A. Graybill, and D. C. Boes, *Introduction to the Theory of Statistics*, 3rd ed. McGraw-Hill, 1974.
- [39] E. H. Lloyd, "Least-squares estimation of location and scale parameters using order statistics," *Biometrika*, vol. 39, no. 1/2, pp. 88–95, Apr. 1952.
- [40] J. D. Broffitt, "An example of the large sample behavior of the midrange," *The American Statistician*, vol. 28, no. 2, pp. 69–70, 1974.
- [41] P. Carbone, C. Narduzzi, and D. Petri, "Performance of stochastic quantizers employing nonlinear processing," *IEEE Trans. Instrum. Meas.*, vol. 45, no. 2, pp. 435–439, Apr. 1996.
- [42] H. L. Van Trees, K. L. Bell, and Z. Tian, *Detection, Estimation, and Modulation Theory: Detection, Estimation, and Filtering Theory*, 2nd ed. Hoboken, NJ: John Wiley & Sons, Inc., 2013.
- [43] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ: Prentice Hall, 1993.
- [44] N. T. Thao and M. Vetterli, "Reduction of the MSE in R -times oversampled A/D conversion from $O(1/R)$ to $O(1/R^2)$," *IEEE Trans. Signal Process.*, vol. 42, no. 1, pp. 200–203, Jan. 1994.
- [45] —, "Lower bound on the mean-squared error in oversampled quantization of periodic signals using vector quantization analysis," *IEEE Trans. Inform. Theory*, vol. 42, no. 2, pp. 469–479, Mar. 1996.
- [46] V. K. Goyal, M. Vetterli, and N. T. Thao, "Quantized overcomplete expansions in $\mathbb{R}^N$: Analysis, synthesis, and algorithms," *IEEE Trans. Inform. Theory*, vol. 44, no. 1, pp. 16–31, Jan. 1998.
- [47] S. Rangan and V. K. Goyal, "Recursive consistent estimation with bounded noise," *IEEE Trans. Inform. Theory*, vol. 47, no. 1, pp. 457–464, Jan. 2001.
- [48] L. Jacques, D. K. Hammond, and J. M. Fadili, "Dequantizing compressed sensing: When oversampling and non-Gaussian constraints combine," *IEEE Trans. Inform. Theory*, vol. 57, no. 1, pp. 559–571, Jan. 2011.
- [49] A. Zymnis, S. Boyd, and E. Candès, "Compressed sensing with quantized measurements," *IEEE Signal Process. Lett.*, vol. 17, no. 2, pp. 149–152, Feb. 2010.
- [50] U. Kamilov, V. K. Goyal, and S. Rangan, "Message-passing dequantization with applications to compressed sensing," *IEEE Trans. Signal Process.*, vol. 60, no. 12, pp. 6270–6281, Dec. 2012.
- [51] M. K. Varanasi and B. Aazhang, "Parametric generalized Gaussian density estimation," *The Journal of the Acoustical Society of America*, vol. 86, no. 4, pp. 1404–1415, 1989.
- [52] Q. Zhao, H.-W. Li, and Y.-T. Shen, "On the sum of generalized Gaussian random signals," in *Proceedings of the International Conference on Signal Processing*, vol. 1, Aug. 2004, pp. 50–53.
- [53] H. Soury and M. S. Alouini, "New results on the sum of two generalized Gaussian random variables," in *2015 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, Dec. 2015, pp. 1017–1021.
- [54] N. C. Beaulieu and Q. Guo, "True ML estimator for the location parameter of the generalized Gaussian distribution with $p = 4$," *IEEE Comm. Lett.*, vol. 17, no. 1, pp. 155–157, Jan. 2013.

- [55] G. R. Arce, *Nonlinear Signal Processing: A Statistical Approach*. John Wiley & Sons, 2005.
- [56] A. Bovik, T. Huang, and D. Munson, "Nonlinear filtering using linear combinations of order statistics," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Process.*, vol. 7, May 1982, pp. 2067–2070.
- [57] —, "A generalization of median filtering using linear combinations of order statistics," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 31, no. 6, pp. 1342–1350, Dec 1983.
- [58] A. K. Gupta, "Estimation of the mean and standard deviation of a normal population from a censored sample," *Biometrika*, vol. 39, no. 3, pp. 260–273, Dec. 1952.
- [59] G. Blom, "On linear estimates with nearly minimum variance," *Arkiv för Matematik*, vol. 3, no. 4, pp. 365–369, Jan. 1957.
- [60] R. Öten and R. J. P. de Figueiredo, "An efficient method for L-filter design," *IEEE Trans. Signal Process.*, vol. 51, no. 1, pp. 193–203, Jan. 2003.
- [61] Anonymous, "Dissertation sur la recherche du milieu le plus probable, entre les résultats de plusieurs observations ou expériences," *Annales de Mathématiques pures et appliquées*, vol. 12, pp. 181–204, 1821.
- [62] S. M. Stigler, "The anonymous Professor Gergonne," *Historia Mathematica*, vol. 3, no. 1, pp. 71–74, Feb. 1976.
- [63] P. Daniell, "Observations weighted according to order," *American Journal of Mathematics*, vol. 42, no. 4, pp. 222–236, 1920.
- [64] S. M. Stigler, "Simon Newcomb, Percy Daniell, and the history of robust estimation 1885–1920," *Journal of the American Statistical Association*, vol. 68, no. 344, p. 31, 1973.
- [65] J. W. Tukey, "The future of data analysis," *The Annals of Mathematical Statistics*, vol. 33, no. 1, pp. 1–67, 1962.
- [66] J. W. Tukey and D. H. McLaughlin, "Less vulnerable confidence and significance procedures for location based on a single sample: Trimming/Winsorization 1," *Sankhyā: The Indian Journal of Statistics, Series A (1961–2002)*, vol. 25, no. 3, pp. 331–352, 1963.
- [67] A. Restrepo and A. C. Bovik, "Adaptive trimmed mean filters for image restoration," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 36, no. 8, pp. 1326–1337, 1988.
- [68] N. C. Beaulieu and Q. Guo, "Novel estimator for the location parameter of the generalized Gaussian distribution," *IEEE Comm. Lett.*, vol. 16, no. 12, pp. 2064–2067, Dec. 2012.
- [69] J. A. Nelder and R. Mead, "A simplex method for function minimization," *The Computer Journal*, vol. 7, no. 4, pp. 308–313, Jan. 1965.
- [70] P. Carbone and D. Petri, "Effect of additive dither on the resolution of ideal quantizers," *IEEE Trans. Instrum. Meas.*, vol. 43, no. 3, pp. 389–396, Jun. 1994.
- [71] E. J. Dudewicz and S. N. Mishra, "Moments," in *Modern Mathematical Statistics*. New York: Wiley, 1988, ch. 5, pp. 216–228.