



HAL
open science

Goal-Oriented 1-Bit Quantization With Uncertain Distortion Measures

Malcolm Egan

► **To cite this version:**

Malcolm Egan. Goal-Oriented 1-Bit Quantization With Uncertain Distortion Measures. IEEE Communications Letters, In press. <hal-05253410>

HAL Id: hal-05253410

<https://hal.science/hal-05253410v1>

Submitted on 14 Sep 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Goal-Oriented 1-Bit Quantization With Uncertain Distortion Measures

Malcolm Egan

Abstract—Tailoring low resolution quantization to specific tasks is an important aspect of goal-oriented communications, with applications ranging from feedback for multiple-input multiple-output (MIMO) systems to compression of neural network weights. However, target distortion functions associated with a task may either be imperfectly known or intractable to compute, leading to uncertainty in the distortion function. In this paper, we propose new criteria based on risk measures to design 1-bit quantizers with uncertainty in the distortion function. While information-theoretic limits of compression with risk measure constraints have recently been established, quantizers based on risk measure criteria have not yet been developed. We introduce an algorithm for quantizer design with our risk measure criteria based on the cross-entropy method, and apply our approach to an environmental monitoring dataset.

I. INTRODUCTION

Low resolution quantization plays an important role in communications and learning. For example, channel feedback for MIMO at mmWave frequencies requires quantizers with only a small number of a bits, or even a single bit [1]. In compression of large deep neural networks for network functions on memory-constrained devices, neuron weights may also be compressed with a low resolution quantizer [2].

The choice of a quantizer depends on the probability distribution of the data, and a measure of distortion between the data samples and their quantizations. For real-valued data, a standard choice of the distortion measure is the mean-square error [3]. However, the choice of the mean-square error is largely due to its simplicity and the fact that many distortion measures are Lipschitz continuous and therefore can be bounded in terms of the Euclidean norm. Indeed, application-dependent measures of distortion can yield improved performance in what is often known as goal- or task-oriented compression [4], [5]. For example, in speech compression the mean distortion measures based on the p -norm, weighted square error amongst others are common alternatives to the mean-square error [6]. In [7], the mean optimality loss has been considered for resource allocation problems, and in [2] mean p -norm distortion functions have been applied for compression of neural network weights.

Goal-oriented quantization can outperform quantizers based on the mean-square error, but only when the distortion measure is appropriately tailored to a target application. In speech and vision applications, a target distortion measure often is

based on human visual or aural perception experiments [8], [9]. Measurements in these experiments are often imperfect, resulting in uncertainty in the quality of the distortion measure. In the context of resource allocation problems, the mean optimality loss may be an appropriate distortion measure [7]. However, it can be extremely costly to evaluate the optimality loss for a large number of samples. As such, approximations of the optimality loss may be required to construct quantizers by, for example, Lloyd-type algorithms [6]. As for speech and vision applications, there is uncertainty in the quality of the distortion measure.

In this paper, we propose a method to select a quantizer when there is uncertainty in the distortion measure with a focus on the important case of 1-bit quantization [10]. We first investigate the impact of uncertainty in the distortion measure in the case of symmetric 1-bit quantization with Gaussian data. For two families of mean distortion measures, one of which can also model uncertainty in the data distribution, we numerically study dependence of performance on the choice of quantization level. We observe that uncertainty in the mean distortion measure can lead to significant variations in the choice of quantization level.

We then introduce our proposed method. The core idea is that uncertainty in the mean distortion measure can be accounted for by increasing the cost associated to large values of the distortion function. We show that this modification leads to a risk measure rather than the standard mean distortion measure [11]. Recently, we have established information-theoretic limits for fixed-length compression with risk measure constraints [12]; however, the design of quantizers with risk measure criteria has not been investigated.

While our method can be viewed as a form of robust quantization, it can be applied when the data has an unbounded support where worst-case distortion criteria [3] cannot be utilized. Our approach also contrasts with perception-distortion criteria [13], which incorporate the Kullback-Leibler divergence in addition to the mean distortion.

Quantizer optimization for mean distortion measures is typically solved via Lloyd-type algorithms [3]. However, these algorithms cannot be directly applied for risk measure criteria. Alternative algorithms based on global search are required. As such, we utilize the cross-entropy method [14] to optimize the quantizer. Using real temperature data from the Twin House Experiment¹, we show that our risk measure criteria can improve quantization of extremes, which is important for classification of extreme events based on quantized data.

M. Egan is with Inria, INSA Lyon, CITI, UR3720, 69621 Villeurbanne, France (email: malcom.egan@inria.fr). This work was partially supported by the French National Agency for Research (ANR) via the project ANR-22-PEFT-0010 of the France 2030 program PEPR réseaux du Futur, and the project ANR-24-CE25-2256 TCDTP.

¹<https://tinyurl.com/32d5nf>

II. UNCERTAIN MEAN DISTORTION MEASURES

A. 1-Bit Quantization

Scalar- and vector-valued data arises naturally from data sets containing images, time series for speech, or channel coefficients in multiple-antenna communication systems. We therefore consider a data source $P_{\mathbf{X}}$ with data $\mathbf{X} \in \mathbb{R}^d$ with $d \geq 1$. The data source is assumed to be known, either from a physical generative model (e.g., ray-tracing models in wireless communications [15]), or approximated by the empirical distribution obtained from a historical data set.

A 1-bit quantizer of data \mathbf{X} consists of a reproduction set $\mathcal{Y} = \{\mathbf{y}_1, \mathbf{y}_2\} \subset \mathbb{R}^d$, and a quantization function $Q : \mathbb{R}^d \rightarrow \mathcal{Y}$. The quantization function Q maps the data \mathbf{X} to the quantization $\hat{\mathbf{X}} \in \mathcal{Y}$. The quality of a 1-bit quantizer is typically characterized by the mean distortion, depending on a distortion function $D : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$; e.g., the square error distortion function $D_{SE}(\mathbf{x}, \hat{\mathbf{x}}) = \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2$. In particular, the performance of a 1-bit quantizer has the form

$$\mathcal{E} = \mathbb{E}[D(\mathbf{X}, \hat{\mathbf{X}})]. \quad (1)$$

As the choice of a 1-bit quantizer depends on the data distribution $P_{\mathbf{X}}$ and the distortion function D , an uncertainty in either can lead to poor performance.

In the sequel, we introduce new criteria for 1-bit quantizer design that form alternatives to mean distortion criteria and mitigates the impact of uncertainty in $P_{\mathbf{X}}$ and D . Before introducing our criteria, we first highlight that uncertainty in $P_{\mathbf{X}}$ can be viewed as uncertainty in D , and illustrate the performance degradation due to uncertainty in D .

B. Distortion Function Uncertainty

Suppose that the data source $P_{\mathbf{X}}$ admits a density $p_{\mathbf{X}}$. In this case,

$$\mathbb{E}[D(\mathbf{X}, \hat{\mathbf{X}})] = \int_{\mathbb{R}^d} D(\mathbf{x}, Q(\mathbf{x})) p_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}. \quad (2)$$

When there is uncertainty in $p_{\mathbf{X}}$, and only an approximate density $\hat{p}_{\mathbf{X}}$ with the same support as $p_{\mathbf{X}}$ is known, we then have

$$\begin{aligned} \mathbb{E}[D(\mathbf{X}, \hat{\mathbf{X}})] &= \int_{\mathbb{R}^d} D(\mathbf{x}, Q(\mathbf{x})) \frac{p_{\mathbf{X}}(\mathbf{x})}{\hat{p}_{\mathbf{X}}(\mathbf{x})} \hat{p}_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} \\ &= \mathbb{E}_{\mathbf{X} \sim \hat{p}_{\mathbf{X}}} [K(\mathbf{X}) D(\mathbf{X}, \hat{\mathbf{X}})], \end{aligned} \quad (3)$$

where $K(\mathbf{x}) = \frac{p_{\mathbf{X}}(\mathbf{x})}{\hat{p}_{\mathbf{X}}(\mathbf{x})}$. In other words, the impact of the uncertainty is a weighting of the original distortion function of the form $K(\mathbf{X})D(\mathbf{X}, \hat{\mathbf{X}})$. An important family of distortion functions are therefore given by

$$D_K(\mathbf{X}, \hat{\mathbf{X}}) = K(\mathbf{X})D(\mathbf{X}, \hat{\mathbf{X}}), \quad (4)$$

where $K : \mathbb{R}^d \rightarrow \mathbb{R}_+$.

Another popular family of distortion functions are based on the p -norm [2], [6]. In particular, the distortion function is given by

$$D_p(\mathbf{X}, \hat{\mathbf{X}}) = \sum_{i=1}^d |x_i - \hat{x}_i|^p, \quad (5)$$

where $p > 0$. By increasing p , the cost of large distortions is also increased.

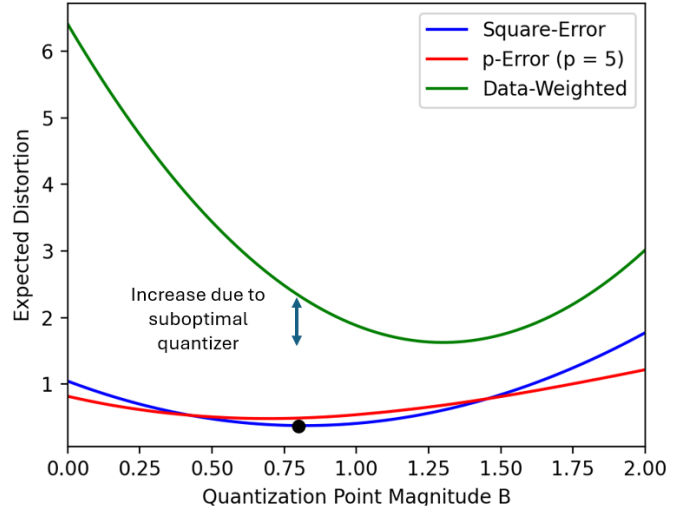


Fig. 1. Impact of distortion functions for $X \sim \mathcal{N}(0, 1)$, $p = 5$, and $D_K(x, \hat{x}) = \exp(|x|)(x - \hat{x})^2$ for the data-weighted scenario.

C. Impact of Distortion Uncertainty for Symmetric Quantizers

To illustrate the impact of uncertainty in the data-weighted and p -error distortion functions, we consider the case of symmetric quantizers for scalar data where $\mathcal{Y} = \{-B, B\}$ for $B \in \mathbb{R}_+$. The quantization point is chosen via

$$Q(x) = \arg \min_{\hat{x} \in \{-B, B\}} (x - \hat{x})^2. \quad (6)$$

Fig. 1 plots the impact on the true distortion metric for a Gaussian source $X \sim \mathcal{N}(0, 1)$ for three scenarios: D_{SE} ; D_p in (5) with $p = 5$; and the weighted distortion function in (4) with $D_K(x, \hat{x}) = \exp(|x|)(x - \hat{x})^2$. Observe that the optimal quantization point magnitude, B , differs significantly in each of the three scenarios. As such, a quantizer based on D_{SE} leads to performance reduction with a mismatched distortion function. On the other hand, if the target distortion function is known, it can be utilized to construct a tailored quantizer.

III. RISK-AVERSE 1-BIT QUANTIZATION

Fig. 1 highlights that design of a quantizer based on an approximation mean distortion criterion can lead to a significant performance reduction. In this section, we develop a new design strategy for 1-bit quantization to account for uncertainty in the distortion function.

A. Proposed Criteria

Consider a family $\mathcal{D} = \{D_n : n \in \mathcal{N}\}$ of distortion functions $D_n : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ index by $n \in \mathcal{N}$. Elements in the index set \mathcal{N} can correspond to different weighting functions K for the weighted distortion function in (4), or values of p in the case of p -norm distortion function (5).

Suppose there is a prior distribution $P_{\mathcal{N}}$ on the index set \mathcal{N} ; i.e., prior information about the target distortion function. In this case, it is natural to construct a quantizer via

$$\min \mathbb{E}_{N, \mathbf{X}} [D_N(\mathbf{X}, Q(\mathbf{X}))]. \quad (7)$$

The problem (7) corresponds to Bayesian robust quantizer design [16]. However, (7) is intractable in general and also requires reliable knowledge of the prior distribution P_N . As such, it is necessary to consider tractable approximations.

The simplest, but widely utilized, approximation is to ignore knowledge of P_N and simply choose a reference distortion function $D_{\text{ref}} \in \mathcal{D}$; e.g., $D_{\text{ref}} = D_{\text{SE}}$. An alternative approach is to observe that

$$\begin{aligned} & \mathbb{E}_{N, \mathbf{X}}[D_N(\mathbf{X}, Q(\mathbf{X}))] \\ &= \mathbb{E}_N \left[\int_0^\infty \mathbb{P}(D_N(\mathbf{X}, Q(\mathbf{X})) > u | N) du \right] \\ &= \int_0^\infty \mathbb{E}_N[\mathbb{P}(D_N(\mathbf{X}, Q(\mathbf{X})) > u | N)] du, \end{aligned} \quad (8)$$

which follows from the fact that the expectation of a positive-valued random variable Z satisfies $\mathbb{E}[Z] = \int_0^\infty \mathbb{P}(Z > u) du$.

Let $h : [0, 1] \rightarrow [0, 1]$ be a concave non-decreasing function satisfying $h(0) = 0$ and $h(1) = 1$, which implies that $h(w) \geq w$, $w \in [0, 1]$ (see [11], [12]). It follows that

$$\begin{aligned} & \mathbb{E}_N[\mathbb{P}(D_N(\mathbf{X}, Q(\mathbf{X})) > u | N)] \\ & \leq \mathbb{E}_N[h(\mathbb{P}(D_N(\mathbf{X}, Q(\mathbf{X})) > u | N))] \end{aligned} \quad (9)$$

for all $u \in [0, \infty)$. A conservative estimate of $\mathbb{E}_{N, \mathbf{X}}[D_N(\mathbf{X}, Q(\mathbf{X}))]$ is then given by

$$\begin{aligned} & \mathbb{E}_{N, \mathbf{X}}[D_N(\mathbf{X}, Q(\mathbf{X}))] \\ & \leq \int_0^\infty \mathbb{E}_N[h(\mathbb{P}(D_N(\mathbf{X}, Q(\mathbf{X})) > u | N))] du. \end{aligned} \quad (10)$$

To remove dependence on the prior P_N , we make the approximation

$$\begin{aligned} & \mathbb{E}_N[h(\mathbb{P}(D_N(\mathbf{X}, Q(\mathbf{X})) > u | N))] \\ & \approx h(\mathbb{P}(D_{\text{ref}}(\mathbf{X}, Q(\mathbf{X})) > u)). \end{aligned} \quad (11)$$

This approximation corresponds to the standard approach for mean distortion criteria in the case $h(w) = w$, $w \in [0, 1]$. Nevertheless, unlike the standard approximation, the approximation in (11) can be made arbitrarily accurate by setting $h(w) \approx 1$ for all $w \in [0, 1]$.

The quantity

$$\rho(D_{\text{ref}}(\mathbf{X}, Q(\mathbf{X}))) = \int_0^\infty h(\mathbb{P}(D_{\text{ref}}(\mathbf{X}, Q(\mathbf{X})) > u)) du \quad (12)$$

is a type of *risk measure* [11]. The impact of choosing $h(w) \geq w$ is a higher cost associated to large values of $D_{\text{ref}}(\mathbf{X}, Q(\mathbf{X}))$. Quantizers corresponding to the solution of

$$\min_{Q, \mathcal{Y}} \rho(D_{\text{ref}}(\mathbf{X}, Q(\mathbf{X}))) \quad (13)$$

are therefore *risk averse*.

We propose to utilize the criterion in (13) to design 1-bit quantizers in the presence of uncertainty in the distortion function (or the data distribution $P_{\mathbf{X}}$, as discussed in Sec. II-B). In contrast to standard goal-oriented compression schemes, instead of assuming that the distortion function D_{ref} is perfectly adapted to a task, we incorporate robustness to uncertainty in the distortion function via h . In the sequel, we design optimal 1-bit quantizers with risk measure criteria by solving (13).

B. Optimal 1-Bit Quantization

To obtain an optimal 1-bit quantizer, it is necessary to select both the reproduction set $\mathcal{Y} = \{\mathbf{y}_1, \mathbf{y}_2\} \subset \mathbb{R}^d$, and the quantization function $Q : \mathbb{R}^d \rightarrow \mathcal{Y}$. We first consider the quantization function Q .

Let $\mathcal{Y} = \{\mathbf{y}_1, \mathbf{y}_2\}$ be fixed and $Q : \mathbb{R}^d \rightarrow \mathcal{Y}$ be arbitrary. Observe that

$$\begin{aligned} & \int_0^\infty h(\mathbb{P}(D_{\text{ref}}(\mathbf{X}, Q(\mathbf{X})) > u)) du \\ & \geq \int_0^\infty h\left(\mathbb{P}\left(\min_{\hat{\mathbf{x}} \in \{\mathbf{y}_1, \mathbf{y}_2\}} D_{\text{ref}}(\mathbf{X}, \hat{\mathbf{x}}) > u\right)\right) du \end{aligned} \quad (14)$$

since for each $u \geq 0$,

$$\mathbb{P}(D_{\text{ref}}(\mathbf{X}, Q(\mathbf{X})) > u) \geq \mathbb{P}\left(\min_{\hat{\mathbf{x}} \in \{\mathbf{y}_1, \mathbf{y}_2\}} D_{\text{ref}}(\mathbf{X}, \hat{\mathbf{x}}) > u\right), \quad (15)$$

and h is a concave non-decreasing function. As such, the optimal quantization function is given by

$$Q^*(\mathbf{x}) = \min_{\hat{\mathbf{x}} \in \{\mathbf{y}_1, \mathbf{y}_2\}} D_{\text{ref}}(\mathbf{X}, \hat{\mathbf{x}}). \quad (16)$$

Note that the optimal quantization function is the same as for the mean distortion criterion corresponding to $h(w) = w$.

The optimal reproduction set \mathcal{Y} is therefore given by

$$\min_{\mathbf{y}_1, \mathbf{y}_2 \in \mathbb{R}^d} \int_0^\infty h\left(\mathbb{P}\left(\min_{\hat{\mathbf{x}} \in \{\mathbf{y}_1, \mathbf{y}_2\}} D_{\text{ref}}(\mathbf{X}, \hat{\mathbf{x}}) > u\right)\right) du. \quad (17)$$

Unlike for mean distortion criteria, a solution to (17) cannot be obtained via a Lloyd-type algorithm [6] due to the non-linearity introduced by h . To obtain a solution we exploit a global optimization method. In particular, we utilize the cross-entropy method with a Gaussian sampling [14]. Pseudocode for the algorithm² is given in Alg. 1.

Algorithm 1 Quantizer Optimization Algorithm

- 1: **Input:** Maximum number of iterations T_{max} , samples $\{\mathbf{x}_j\}_{j=1}^S$ from $P_{\mathbf{X}}$, samples for CEM N , reference distortion function D_{ref} , risk measure parameter function h in (13), initial search parameters $\boldsymbol{\mu} \in \mathbb{R}^2$, $\sigma \in \mathbb{R}_+$.
 - 2: Initialize $t = 0$, $\mathcal{Y}_{\text{best}}$, and $\rho_{\text{best}} = \infty$.
 - 3: **for** $t < T_{\text{max}}$ **do**
 - 4: $t \leftarrow t + 1$.
 - 5: Sample quantizers $\mathcal{Y}_i \sim \mathcal{N}(\boldsymbol{\mu}, \sigma^2 \mathbf{I})$, $i = 1, \dots, N$.
 - 6: Estimate risk measure ρ_i in (13) from data $\{\mathbf{x}_j\}_{j=1}^S$ for each \mathcal{Y}_i , $i = 1, \dots, N$.
 - 7: **for** $i = 1, \dots, N$ **do**
 - 8: **if** $\rho_i < \rho_{\text{best}}$ **then**
 - 9: $\rho_{\text{best}} \leftarrow \rho_i$, $\mathcal{Y}_{\text{best}} \leftarrow \mathcal{Y}_i$.
 - 10: **end if**
 - 11: Update $\boldsymbol{\mu}, \sigma$ via [14, Algorithm 4.1 Eq. (19-21)].
 - 12: **end for**
 - 13: **end for**
 - 14: **Output:** $\mathcal{Y}_{\text{best}}$.
-

²Code to implement the algorithm along with parameters for the cross-entropy method are available at https://github.com/malcolmalexegan/risk_averse_quantization/tree/master.

IV. NUMERICAL RESULTS

In rare-event estimation with scalar measurements, such as extreme temperature events, it is desirable to reduce the distortion associated to samples with large values. This can be modeled via a distortion function $K(X)|X - \hat{X}|^2$ with K monotonically increasing. In this case, misspecification of K , such as the choice $D_{\text{ref}}(X, \hat{X}) = |X - \hat{X}|^2$, can lead to poor performance.

We investigate our 1-bit quantization strategy using temperature data from the Twin House Experiment³. In this setting, 1-bit quantization reduces the communication requirements of low-power sensors, relevant for classification of extreme temperature events condition monitoring applications. We investigate the design of quantizers based on the criteria

$$\rho_\alpha = \int_0^\infty \left(\mathbb{P}(D_{\text{SE}}(X, \hat{X}) > u) \right)^\alpha du, \quad \alpha \in (0, 1]. \quad (18)$$

When $\alpha = 1$, $\rho = \mathbb{E}[D_{\text{SE}}(X, \hat{X})]$, while for small values of α , the quantizer attempts to reduce the distortion for lower probability events. This has the byproduct of reducing quantization errors for extreme temperature events. Fig. 2 shows the temperature time series and corresponding quantized data obtained from the optimized quantizer with $\alpha = 0.1$.

Table I shows the optimized quantizer using Alg. 1, decision boundary, and corresponding values of risk measures with varying α . Here, ρ_α denotes the value of the risk measure in (18) with parameter α . Observe that quantization levels y_1, y_2 in \mathcal{Y} and the decision boundary for the optimized quantizer increase as α decreases. In other words, increasing the cost associated with low probability events associated with extreme events leads to quantization levels that approach the values of the extreme events, reducing their distortion. For example, the square-error distortion associated with a temperature of 40 degrees Celsius is 20.1 with $\alpha = 0.1$. In contrast, the square error is 68.2 for the standard $\mathbb{E}[D_{\text{SE}}(X, \hat{X})]$ criterion. As such, the quantizer has a lower distortion for the extreme values of the temperature.

TABLE I
RISK-AWARE QUANTIZATION

α	y_1	y_2	Decision Boundary	$\mathbb{E}[D_{\text{SE}}(X, \hat{X})]$	$\rho_{0.5}$	$\rho_{0.1}$
1	24.44	31.74	28.09	5.44	15.17	52.57
0.5	24.77	33.43	29.1	7.04	12.60	34.72
0.1	25.42	35.51	30.46	12.14	16.88	23.15

V. CONCLUSION

Goal-oriented compression relies on knowledge of an application-dependent distortion function. In practice, there may be uncertainty in the choice of this distortion function. In this paper, we have proposed a new approach to robust quantization to cope with both uncertainty in the data distribution and the distortion function. We establish a connection between robust quantization and risk measures, which leads to new criteria. To optimize the quantizer based on risk measure criteria, we apply the cross-entropy method. Our




Fig. 2. Temperature time series from the Twin House Experiment (sensor in the living room at a height of 10cm) and corresponding quantization based on $\rho_{0.1}$, corresponding to $\alpha = 0.1$ in (18).

approach is illustrated for environmental monitoring data from the Twin House Experiment, which shows improved distortion for temperature extremes.

REFERENCES

- [1] J. Mo and R. Heath, "Capacity analysis of one-bit quantized MIMO systems with transmitter channel state information," *IEEE Transactions on Signal Processing*, vol. 63, no. 20, pp. 5498–5512, Oct. 2015.
- [2] Y. Nahshan *et al.*, "Loss aware post-training quantization," *Machine Learning*, vol. 110, no. 11, pp. 3245–3262, 2021.
- [3] A. Gersho and R. Gray, *Vector Quantization and Signal Compression*. New York: Springer Science & Business Media, 2012.
- [4] D. Gündüz *et al.*, "Beyond transmitting bits: Context, semantics, and task-oriented communications," *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 1, pp. 5–41, Nov. 2022.
- [5] N. Shlezinger, Y. Eldar, and M. Rodrigues, "Hardware-limited task-based quantization," *IEEE Transactions on Signal Processing*, vol. 67, no. 20, pp. 5223–5238, 2019.
- [6] Y. Linde, A. Buzo, and R. Gray, "An algorithm for vector quantizer design," *IEEE Transactions on communications*, vol. 28, no. 1, pp. 84–95, Jan. 1980.
- [7] H. Zou, C. Zhang, S. Lasaulce, L. Saludjian, and H. Poor, "Goal-oriented quantization: Analysis, design, and application to resource allocation," *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 1, pp. 42–54, Nov. 2022.
- [8] N. Jayant, J. Johnston, and R. Safranek, "Signal compression based on models of human perception," *Proceedings of the IEEE*, vol. 81, no. 10, pp. 1385–1422, Oct. 1993.
- [9] M. Eckert and A. Bradley, "Perceptual quality metrics applied to still image compression," *Signal Processing*, vol. 70, no. 3, pp. 177–200, Nov. 1998.
- [10] A. Magnani, A. Ghosh, and R. Gray, "Optimal one-bit quantization," in *IEEE Data Compression Conference*, 2005.
- [11] E. Sereda *et al.*, "Distortion risk measures in portfolio optimization," *Handbook of Portfolio Construction*, pp. 649–673, 2010.
- [12] M. Egan, "Risk-aware estimation from compressed data beyond the bayes risk," in *IEEE International Symposium on Information Theory (ISIT)*, 2025.
- [13] Y. Blau and T. Michaeli, "Rethinking lossy compression: the rate-distortion-perception tradeoff," in *International Conference on Machine Learning (ICML)*, 2019.
- [14] Z. Botev *et al.*, "The cross-entropy method for optimization," *Handbook of Statistics*, vol. 31, pp. 35–59, 2013.
- [15] J. Hoydis *et al.*, "Sionna RT: Differentiable ray tracing for radio propagation modeling," in *IEEE Globecom Workshops (GC Wkshps)*, 2023.
- [16] A. Vempaty *et al.*, "On quantizer design for distributed Bayesian estimation in sensor networks," *IEEE Transactions on Signal Processing*, vol. 62, no. 20, pp. 5359–5369, Aug. 2014.

³<https://tinyurl.com/32djoy5nf>