

Low-Complexity Vector Quantized Compressed Sensing via Deep Neural Networks

MARKUS LEINONEN¹ (Member, IEEE), AND MARIAN CODREANU² (Member, IEEE)

¹Centre for Wireless Communications—Radio Technologies, University of Oulu, 90014 Oulu, Finland

²Department of Science and Technology, Linköping University, 581 83 Linköping, Sweden

CORRESPONDING AUTHOR: M. LEINONEN (e-mail: markus.leinonen@oulu.fi)

This work was supported in part by Infotech Oulu, in part by the Academy of Finland under Grant 323698, and in part by the Academy of Finland 6Genesis Flagship under Grant 318927. The work of Markus Leinonen was supported in part by the Academy of Finland under Grant 319485. The work of Marian Codreanu was supported by the European Union's Horizon 2020 Research and Innovation Programme through the Marie Skłodowska-Curie Grant under Agreement 793402 (COMPRESS NETS). Preliminary results of this work were presented in [1].

ABSTRACT Sparse signals, encountered in many wireless and signal acquisition applications, can be acquired via compressed sensing (CS) to reduce computations and transmissions, crucial for resource-limited devices, e.g., wireless sensors. Since the information signals are often continuous-valued, digital communication of compressive measurements requires quantization. In such a quantized compressed sensing (QCS) context, we address remote acquisition of a sparse source through vector quantized noisy compressive measurements. We propose a deep encoder-decoder architecture, consisting of an encoder deep neural network (DNN), a quantizer, and a decoder DNN, that realizes low-complexity vector quantization aiming at minimizing the mean-square error of the signal reconstruction for a given quantization rate. We devise a supervised learning method using stochastic gradient descent and backpropagation to train the system blocks. Strategies to overcome the vanishing gradient problem are proposed. Simulation results show that the proposed non-iterative DNN-based QCS method achieves higher rate-distortion performance with lower algorithm complexity as compared to standard QCS methods, conducive to delay-sensitive applications with large-scale signals.

INDEX TERMS Compressed sensing, data compression, feedforward neural networks, supervised learning, vector quantization.

I. INTRODUCTION

IN A MYRIAD of wireless applications and signal acquisition tasks, information signals are *sparse*, i.e., they contain many zero-valued elements, either naturally or after a transformation [2]. Sparse signals are encountered in, e.g., environmental monitoring [3], source localization [4], spectrum sensing [5], and signal/anomaly detection [6]. Sparsity can be utilized by the joint sampling and compression paradigm, *compressed sensing* (CS) [7]–[9], which enables accurate reconstruction of a sparse signal from few (random) linear measurements. Due to its simple encoding, allowed by more computationally intensive decoding [10], CS is well-suited for communication applications with resource-limited encoder devices, e.g., low-power wireless sensors.

Early CS works considered continuous-valued signals and treated CS as a *dimensionality reduction* technique.

However, the information signals are often continuous-valued, and thus, digital communication/storage of compressive measurements requires quantization. Quantization resolution may range from simple 1-bit quantization [11] to multi-bit quantization [12] which balances between the performance and complexity, crucial in, e.g., wireless sensors. Applying CS for digital transmission/storage initiated the *quantized compressed sensing* (QCS) framework [13]. QCS accomplishes *source compression* in the information-theoretic sense; it compresses continuous-valued signals to finite representations. Due to *indirect* observations of a source, the compression in QCS falls into *remote source coding* [14], [15, Sec. 3.5].

Standard decoders designed for non-quantized CS yield inferior rate-distortion performance in QCS [13]. Hence, it is of the utmost importance to (re)design the recovery

methods for QCS to handle the impact of highly non-linear quantization, especially for low-rate schemes. The overarching idea of numerous QCS algorithms is to explicitly accommodate the presence of quantization in the encoder/decoder. First QCS works used *scalar quantizers* (SQs¹) and optimized either the (quantization-aware) encoder or decoder [16], [17]. At the cost of increased complexity, enhanced rate-distortion performance is achieved by *vector quantization* (VQ) [18]–[20]. For empirical performance comparison of various QCS methods, see, e.g., [21], [2, Sec. 7.6].

The existing QCS methodology has two bottlenecks: 1) high encoding or/and decoding complexity, and 2) high decoding latency. Namely, although SQ permits simple encoding, the decoder runs a (quantization-aware) greedy/polynomial-complexity CS algorithm [22] that may become prohibitive for large-scale data and cause unacceptable delays in real-time applications. On the other hand, VQ yields superior rate-distortion performance – even approaching the information-theoretic limit with the aid of entropy coding [20] – but the encoding complexity grows infeasible. One remedy for this complexity-performance hindrance in QCS is *deep learning*: realizing the CS decoder by a *deep neural network* (DNN), along with a simple encoder/quantizer. The crux is that if such a *non-iterative* signal reconstruction method meets the desired rate-distortion performance after trained *offline*, the *online* communications enjoy an extremely *fast, low-complexity* encoding-decoding process. This would allow a resource-limited encoder device limitations, to compress and communicate large-scale data in a timely fashion.

Launching a fresh view on sparse signal reconstruction, the work [23] was the first to apply deep learning for (non-quantized) CS. The authors trained stacked denoising autoencoders to learn sparsity structures, nonlinear adaptive measurements, and a decoder to reconstruct sparse signals (images). Despite remarkable advances in non-quantized CS, only a few works have applied deep learning for QCS. The first end-to-end QCS design was proposed in [24], where the devised method optimizes a binary measurement acquisition realized by a DNN, a compander-based non-uniform quantizer, and a DNN decoder to estimate neural spikes.

DNN designs applied for non-quantized CS do not resolve the peculiarities that the discretized measurements bring about. Direct adoption of “non-quantized” learning techniques in a QCS setup is inapplicable; the non-differentiable quantization induces a *vanishing gradient problem* [25]–[27], precluding the use of standard *stochastic gradient descent* (SGD) [28, Ch. 5.9] and *backpropagation* [29] in the DNN training. Since quantization becomes a pronounced factor in degrading the rate-distortion performance for low bit resolutions, the presence of a quantizer needs to be *integrated* in the design. Putting all these into

a QCS context, wherein the decoder receives the compressive measurements from the encoder only in digital form, **the pertinent design task is to jointly optimize the encoder and decoder to obtain accurate signal estimates under coarsely quantized measurements-which is the main focus of our paper.**

We address low-complexity remote acquisition of a sparse source through vector quantized noisy compressive measurements. To tackle the pertinent *source compression* problem, we propose a **deep encoder – decoder architecture** for QCS, where 1) the encoder realizes low-complexity VQ of measurements by the cascade of an *encoder DNN* and a (non-uniform) SQ, and 2) the decoder feeds the quantized measurements into a *decoder DNN* to estimate the source. The key driver to our proposed DNN architecture is the fact the optimal compression in QCS is achieved by constructing a minimum mean square error estimate of the sparse source at the encoder and compressing the estimate with a VQ [14], [20], [30]. The objective is to train the proposed scheme to minimize the mean square error (MSE) of the signal reconstruction for a given measurement matrix and quantization rate. We use SGD and backpropagation to develop a practical supervised learning algorithm to train *jointly* the encoder DNN, quantizer, and decoder DNN. The main design driver is that once trained *offline*, the non-iterative QCS method provides an extremely fast and low-complexity encoding-decoding stage for *online* communications, conducive to delay-sensitive applications with large-scale signals. As the key technique to ameliorate the training, we adopt *soft-to-hard quantization* (SHQ) [31] at the encoder DNN to mitigate the vanishing gradient problem. The core idea is to adjust a *continuous* SHQ mapping during the training to *asymptotically* replicate the behavior of a continuous-to-discrete SQ implemented after training. Simulation results show that the proposed method obtains superior rate-distortion performance with faster algorithm running time compared to standard QCS methods.

A. CONTRIBUTIONS

To summarize, the main contributions of our paper are:

- We propose a deep encoder-decoder architecture for QCS, consisting of the encoder DNN, quantizer, and decoder DNN for efficient compression of sparse signals.
- The proposed encoder – the cascade of the encoder DNN and SQ – realizes low-complexity VQ, enhancing the rate-distortion performance.
- We provide a comprehensive treatment of the SGD optimization steps and develop a practical supervised learning algorithm to train the proposed method.
- We propose practical asymptotic quantizer and gradient approximation strategies for the SHQ stage to facilitate the SGD optimization and improve the performance.
- Extensive numerical experiments illustrate that the proposed QCS method obtains superior rate-distortion performance with a low-complexity, fast

1. “SQ” is used interchangeably to refer to “scalar quantizer” and “scalar quantization”; the similar convention holds for “VQ”.

encoding-decoding process in comparison to several conventional QCS methods relying on either SQ or VQ.

- Owing to the use of VQ, the proposed DNN-based QCS method is empirically shown to be capable to approach the finite block length compression limits of QCS.

To the best of our knowledge, DNN-based vector quantized CS has not been addressed earlier. Overall, our work gives a thorough view on the quantization aspects and challenges in the DNNs from the *source compression* viewpoint, which is less explored so far. Hence, the paper is endeavored to open avenues for further developments in the related DNN context.

B. RELATED WORK

In this section, we discuss the connections and differences of the related works to our paper.

1) LEARNING-BASED NON-QUANTIZED CS METHODS

The first DNN-based non-quantized CS framework in [23] spurred a vast succession of (convolutional) neural network designs for compressive imaging, giving rise to, e.g., “ReconNet” [32], “DeepCodec” [33], “DeepInverse” [34], “SSDAE_CS” [35], and “ADMM-CSNet” [36]. For wireless neural recording, an autoencoder with a binary measurement matrix was devised in [37]. Some works have enlightened the connections between the standard and learning-aided CS recovery; for example, “Deep ℓ_0 Encoder” was proposed in [38] for approximate ℓ_0 -minimization. An emergent framework [39] employs *generative adversarial models* for compressive signal reconstruction. Besides the above point-to-point cases, deep learning has been applied for distributed CS in [40]. These learning-based CS signal reconstruction techniques do not consider quantization.

2) DNN-BASED QCS METHODS

The most related work to our paper is [24], where the developed “BW-NQ-DNN” method has the following differences. 1) The encoder in [24] has direct access to the source which allows to realize (and optimize) the measurement matrix by a DNN. This is inapplicable in our setup where the encoder observes the source only *indirectly* through CS; herein, the (fixed) measurement matrix is dictated by a physical sub-sampling mechanism [41]. 2) [24] quantizes measurements by a (non-uniform) SQ; we use VQ, the benefits of which are substantiated in Section V. 3) [24] uses a *straight-through estimator* [25]; our SHQ with quantizer and gradient approximation policies is demonstrated to improve the (encoder) training and rate-distortion performance. 4) [24] uses a compander to realize non-uniform SQ; our encoder DNN (a non-linear transformation) surrogates the compander. 5) [24] considers noiseless CS; we consider a noisy setup.

Another related works include [42], [43]. Different to our work, the DNN-based image recovery method in [42]

1) assumes direct access to the source, 2) passes the gradient through the quantizer via approximate rounding, and 3) employs SQ (in a block-by-block fashion). The work [43] designs only the encoder by learning-based optimization of the CS measurement acquisition for a given decoder and uniform SQ.

3) THE “DNN-QUANTIZER-DNN” ARCHITECTURE

While the deep task-based quantization scheme in [31], [44] does not address source compression and CS, there is a connection to our encoder-decoder DNN architecture. The works [31], [44] devised a DNN-based multiple-input multiple-output communication *receiver*, consisting of an analog DNN, a quantizer, and a digital DNN; this cascade is, to some extent, analogous to our encoder DNN, quantizer, and decoder DNN. Thus, [31], [44] address bit-constrained signal acquisition at a DNN-based *decoder*, whereas we consider bit-constrained source compression at a resource-limited *encoder*; our (CS-based) encoder undergoes the stringently bit-constrained quantization stage, imposed by, e.g., a low-resolution analog-to-digital converter (ADC) or/and rate-limited communications. To this end, we devise a deep *joint encoder-decoder* as a first attempt to apply DNN-based VQ in QCS.

4) SOFT-TO-HARD QUANTIZATION

We integrate the SHQ, proposed in [31], [44], in our deep encoder-decoder scheme and provide techniques to optimize the SHQ to ameliorate the training. In DNNs, quantization has primarily been addressed for *DNN quantization*, i.e., discretization of full-precision weights and biases for memory-efficient DNN implementation [26], [27], [45]–[47]. These works have connections to our SHQ design. Similar to our SHQ, differential soft quantization in [47] uses a series of hyperbolic tangents. A soft-to-hard annealing technique was proposed for DNN and data compression in [26]; differently to our work, it uses a softmax-operator. Our gradual gradient approximation is akin to the “alpha-blending” method in [27].

5) AUTOENCODER

Our DNN architecture resembles one special feedforward-type DNN – an *autoencoder* [48]. An autoencoder attempts to copy its input to the output through an encoding and decoding function while undergoing an intermediate “compression/representation” stage [28, Ch. 14]. In light of remote observations, a CS scheme resembles a *denoising autoencoder* [49], [28, Ch. 14.2.2] which amounts to estimate the source from a *corrupted* input. The “ ℓ_1 -AE” autoencoder proposed in [50] learns a linear encoder (i.e., the dimensionality reduction step) for a standard ℓ_1 -decoder. *Uncertainty autoencoders* were employed in [51] to learn the measurement acquisition and recovery stages. Autoencoders have been designed for, e.g., CS reconstruction in [35] and sparse support recovery in [52]. All the above methods consider non-quantized CS.




FIGURE 1. An illustration of the proposed DeepVQCS scheme for a signal of length $N = 10$ for $M = 4$ measurements, where 1) EncNet has $J = 3$ layers of widths $e_1 = M = 4$, $e_2 = 5$, and $e_3 = K = 3$; and 2) DecNet has $L = 5$ layers of widths $d_1 = K = 3$, $d_2 = 5$, $d_3 = 6$, $d_4 = 10$, and $d_5 = N = 10$.

In a non-CS setup [53], *compressive autoencoders* were proposed for lossy image compression.

We conclude the section by highlighting that since the inputs of our encoder DNN are the *compressed* measurements, the designated task of our proposed DNN cascade is to copy a *hidden/remote* information source to the decoder output. The main distinction to the above “non-quantized autoencoders” is that **our source compression task calls for optimizing finite representation for the measurement vector, which itself has already undergone the dimensionality reduction stage of CS prior to accessing the encoder.**

Organization: The paper is organized as follows. The system model and the problem definition are presented in Section II. The proposed deep encoder-decoder architecture for QCS is introduced in Section III. Optimization of the proposed method is detailed in Section IV. Simulation results are presented in Section V. Conclusions are drawn in Section VI.

Notations: Boldface capital letters (**A**) denote matrices. Boldface small letters (**a**) denote vectors. Calligraphy letters (\mathcal{A}) denote sets. \mathbb{R}_+ denotes the set of non-negative real numbers. $\mathbf{0}$ is a vector of all entries zero. $\mathbf{1}$ is a vector of all entries one. $(\cdot)^T$ denotes the matrix transpose. \odot denotes the Hadamard product. \circ denotes a composite function. $\tanh(\cdot)$ is the hyperbolic tangent $\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$. $f'(\cdot)$ denotes differentiation of function $f(x)$ with respect to x . $\lceil \cdot \rceil$ denotes rounding up to the nearest integer. $\|\mathbf{a}\|_0$ counts the number of non-zero entries of vector **a**. $\|\cdot\|_1$ and $\|\cdot\|_2$ denote the ℓ_1 -norm and ℓ_2 -norm.

II. SYSTEM MODEL AND PROBLEM DEFINITION

We consider a *remote* signal acquisition setup depicted in Fig. 1. Encoder E (e.g., a low-power wireless sensor)

observes an information source *indirectly* in the form of noisy compressive (dimensionality-reducing) measurements. Encoder E processes the measurements by an *encoder DNN*, *quantizes* the DNN output, and communicates the quantized measurements to decoder D (e.g., a wireless access point). Since our main focus is on source compression, we assume that the communication from encoder E to decoder D is error-free. Decoder D feeds the quantized measurements into a *decoder DNN* to estimate the (remotely observed) source. As will be elaborated later, the method realizes (low-complexity) *vector quantization* (VQ); hence, we dub the proposed deep encoder-decoder architecture for QCS as **DeepVQCS**.

Next, we present the sensing setup and state the considered QCS problem; the DNNs and quantization stage are detailed in Section III.

A. SOURCE SIGNAL AND COMPRESSIVE MEASUREMENTS

Let $\mathbf{x} \in \mathbb{R}^N$ denote a source vector, representing, e.g., a sequence of temperature values at consecutive discrete time instants. The realizations of \mathbf{x} are assumed to be independent and identically distributed across time. We assume that vector² \mathbf{x} is S -sparse, i.e., it has at most S non-zero entries, $\|\mathbf{x}\|_0 = S \leq N$. The *a priori* probabilities of the sparsity patterns are unknown.

Encoder E observes remote source \mathbf{x} (only) indirectly in the form of *noisy compressive measurements* as [7], [8]

$$\mathbf{y} = \Phi \mathbf{x} + \mathbf{n}, \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^M$, $M < N$, is a measurement vector, $\Phi \in \mathbb{R}^{M \times N}$ is a fixed and known measurement matrix, and $\mathbf{n} \in \mathbb{R}^M$ is a

2. For simplicity, we assume that \mathbf{x} itself is sparse in a canonical basis.

noise vector. It is worth emphasizing that encoder \mathbf{E} has no access to information source \mathbf{x} : the encoder device samples and acquires \mathbf{x} merely via the CS, where Φ is dictated by a physical sub-sampling mechanism [41]. Due to the indirect observations, the compression in a QCS setup is referred to as *remote source coding* [14], [15, Sec. 3.5].

B. PROBLEM DEFINITION

The QCS problem for optimizing the DeepVQCS scheme depicted in Fig. 1 is given as follows.

Definition 1 (QCS Problem): Given the compressive measurements (1), a fixed measurement matrix Φ , and total number of quantization levels \bar{I} used to represent (i.e., *compress*) \mathbf{y} , the objective is to jointly optimize encoder \mathbf{E} and decoder \mathbf{D} – the encoder DNN, the quantizers, and the decoder DNN³ – for given DNN configurations to minimize the mean square error (MSE) of the signal reconstruction

$$D(\mathbf{E}, \mathbf{D}) = \mathbb{E}[\|\mathbf{D}(\mathbf{E}(\mathbf{y})) - \mathbf{x}\|_2^2], \quad (2)$$

where the expectation⁴ is with respect to the randomness of source \mathbf{x} and noise \mathbf{n} ; $\mathbf{D}(\mathbf{E}(\mathbf{y}))$ represents an estimate of source vector \mathbf{x} at the output of decoder \mathbf{D} .

Our main interest in tackling the above *source compression* problem in QCS lies in devising the DeepVQCS encoder-decoder architecture to enable a fast and low-complexity encoding-decoding stage, beneficial to real-time applications with large-scale sparse signals. Although the development is not tied to any particular QCS framework, our design is driven by a communication scenario, where encoder \mathbf{E} is a resource-limited device (e.g., a wireless sensor) imposed by substantial limitations on total quantization/communication rate.

III. DEEP ENCODER-DECODER ARCHITECTURE FOR QUANTIZED COMPRESSED SENSING

In this section, we detail the structure and operation of each block of the DeepVQCS architecture. Implementation aspects of DeepVQCS are also discussed.

A. ENCODER

Encoder \mathbf{E} comprises the encoder DNN and a quantizer encoder, described next.

1) ENCODER DNN

As the first stage at encoder \mathbf{E} , the measurements in (1) are fed into an encoder DNN⁵, dubbed **EncNet**. We consider a *feedforward* DNN, or a *multilayer perceptron* [28, Ch. 6], i.e., the connections between the nodes, *neurons*, form no

loops. Moreover, EncNet is *fully connected*,⁶ i.e., each neuron at a layer is connected to all neurons at the next layer. EncNet has J layers, i.e., its *depth* is J . Next, we detail the structure and operation of EncNet.

Let vector $\mathbf{c}_j = [c_{j,1} \cdots c_{j,e_j}]^T \in \mathbb{R}^{e_j}$ be the *weighted input* at layer j , where e_j denotes the number of neurons at layer j , i.e., the *width* of layer j . For each *hidden* layer $j = 2, \dots, J$, the weighted input at layer j has a linear relationship to its preceding layer $j - 1$ as

$$\mathbf{c}_j = \mathbf{F}_j \mathbf{a}_{j-1} + \mathbf{b}_j, \quad j = 2, \dots, J, \quad (3)$$

where $\mathbf{F}_j \in \mathbb{R}^{e_j \times e_{j-1}}$ is the *weight matrix* at layer j , $\mathbf{b}_j \in \mathbb{R}^{e_j}$ is the *bias vector* at layer j , and vector $\mathbf{a}_j \in \mathbb{R}^{e_j}$ is the *output* of layer j , defined as

$$\mathbf{a}_1 = \mathbf{y}, \quad \mathbf{a}_j = \gamma_j(\mathbf{c}_j) = \gamma_j(\mathbf{F}_j \mathbf{a}_{j-1} + \mathbf{b}_j), \quad j = 2, \dots, J, \quad (4)$$

where $\gamma_j(\cdot)$ is an (element-wise) activation function at layer j ; $\mathbf{a}_1 = \mathbf{y}$ implies that the *input layer* of EncNet is formed by measurement vector \mathbf{y} in (1) (i.e., $e_1 = M$).

Let $K \triangleq e_J$ denote the width of the *output layer* $j = J$ of EncNet. Accordingly, EncNet takes $\mathbf{y} \in \mathbb{R}^M$ in (1) as its input and produces $\mathbf{a}_J = [a_{J,1} \cdots a_{J,K}]^T$ of form (4) as an output. We define EncNet as a mapping

$$\Omega_E : \mathbb{R}^M \rightarrow \mathbb{R}^K, \quad \Omega_E(\mathbf{y}) = \mathbf{a}_J. \quad (5)$$

2) QUANTIZER ENCODER

Digital communication of the EncNet outputs, i.e., the *pre-processed* continuous-valued measurements $\mathbf{a}_J = [a_{J,1} \cdots a_{J,K}]^T$ in (5), to decoder \mathbf{D} necessitates *quantization*, performed as the second stage at encoder \mathbf{E} . We consider a low-complexity quantization scheme where each element $a_{J,n}$, $n = 1, \dots, K$, is converted into a discrete representation by a *scalar quantizer* (SQ); we let the SQ to be *non-uniform*, albeit a uniform SQ may be preferred in practice due to its simplicity. Accordingly, the quantization of vector $\mathbf{a}_J \in \mathbb{R}^K$ can be modeled⁷ as K identical SQs (see Fig. 1). A formal definition of the SQ is given next.

Definition 2 (SQ): Let $\mathbf{Q} = \{\mathbf{Q}_E, \mathbf{Q}_D^{-1}\}$ represent an I -level SQ, consisting of quantizer encoder \mathbf{Q}_E (located at encoder \mathbf{E}) and quantizer decoder \mathbf{Q}_D^{-1} (located at decoder \mathbf{D}). Let $\mathcal{R} = \{\mathcal{R}_i\}_{i \in \mathcal{I}}$ be the set of quantization regions, where $\mathcal{I} = \{1, \dots, I\}$ is the set of quantization indices. \mathcal{R} partitions real line with disjoint and exhaustive regions $\mathcal{R}_i = (t_{i-1}, t_i]$, where $t_i \in \mathbb{R}$ is a *threshold* with $t_1 \leq t_2 \leq \dots \leq t_{I-1}$; here, $\mathcal{R}_1 = (-\infty, t_1]$ and $\mathcal{R}_I = (t_{I-1}, \infty)$. Let $\mathcal{G} = \{g_i\}_{i \in \mathcal{I}}$ be the set of discrete *reproduction levels*, where $g_i \in \mathbb{R}$ is the level associated with region \mathcal{R}_i , $i \in \mathcal{I}$. The quantizer encoder is

3. The structure and operation of the encoder DNN, quantizers, and the decoder DNN will become explicit in Section III.

4. Throughout the paper, all expectations $\mathbb{E}[\cdot]$ are taken with respect to the randomness of source \mathbf{x} and noise \mathbf{n} .

5. Regardless of the depth J , we refer to EncNet as a *deep* neural network for brevity; similar convention is used for DecNet.

6. A standard fully connected feedforward DNN is considered due to its universality and suitability for supervised learning on fixed-size input vectors [28, Sec. 11.2]. Designs with more sophisticated DNN architectures are left for future work.

7. Note that portraying K parallel quantizers in Fig. 1 is for illustration purposes; the considered SQ scheme under a single quantization rule can be implemented, e.g., by one serial scalar ADC [54] in a space-efficient manner.

a mapping $\mathbf{Q}_E : \mathbb{R} \rightarrow \mathcal{I}$; for n th EncNet output, $a_{J,n}$, it operates as

$$\mathbf{Q}_E(a_{J,n}) = i_n \in \mathcal{I}, \text{ if } a_{J,n} \in \mathcal{R}_{i_n}, n = 1, \dots, K. \quad (6)$$

The quantizer decoder is a mapping $\mathbf{Q}_D^{-1} : \mathcal{I} \rightarrow \mathcal{G}$; for a received index $i_n \in \mathcal{I}$, it operates as

$$\mathbf{Q}_D^{-1}(i_n) = g_{i_n}, n = 1, \dots, K. \quad (7)$$

In practice, quantization indices $i \in \mathcal{I}^8$ are communicated from encoder \mathbf{E} to decoder \mathbf{D} as *binary* code words. We assume that this communication is lossless; the design of binary code words is outside of the scope of our paper. The total number of quantization levels used for compressing a measurement vector \mathbf{y} is $\bar{I} \triangleq KI$.

Remark 1: Hardware, energy, and computation restrictions of low-cost encoder devices may limit the quantization resolution, even down to one bit per sample [11]. In such QCS scenarios, the quantization error becomes a dominating factor (despite having a sufficient number of measurements M) in degrading the signal reconstruction accuracy. To mitigate this effect, the presence of a quantizer must be appropriately *integrated* in the design. This is the main focus of our paper – **devising a low-complexity QCS scheme to obtain high rate-distortion performance under severely bit-constrained source encoding**.

3) FULL OPERATION OF ENCODER

Combining the operations of EncNet in (5) and quantizer encoder \mathbf{Q}_E in (6), encoder \mathbf{E} can be expressed as a composite function

$$\mathbf{E} = \bar{\mathbf{Q}}_E \circ \Omega_E : \mathbb{R}^M \rightarrow \mathcal{I}^K, \quad (8)$$

where $\bar{\mathbf{Q}}_E(\mathbf{a}_J)$ represents the aggregate quantization operation of vector \mathbf{a}_J as

$$\bar{\mathbf{Q}}_E(\mathbf{a}_J) = \{\mathbf{Q}_E(a_{J,1}), \dots, \mathbf{Q}_E(a_{J,K})\} = \{i_1, \dots, i_K\}. \quad (9)$$

Finally, the full operation of encoder \mathbf{E} is presented as

$$\{i_1, \dots, i_K\} = \bar{\mathbf{Q}}_E(\gamma_J(\mathbf{F}_J(\cdot \gamma_3(\mathbf{F}_3(\gamma_2(\mathbf{F}_2\mathbf{y} + \mathbf{b}_2)) + \mathbf{b}_3) \cdot \cdot) + \mathbf{b}_J)). \quad (10)$$

Remark 2: While each output element of EncNet, $a_{J,n}$, $n = 1, \dots, K$, is *separately* quantized by \mathbf{Q}_E in (6), encoder \mathbf{E} – the cascade of EncNet and \mathbf{Q}_E – realizes (low-complexity) **vector quantization (VQ)** of measurement vector \mathbf{y} . Namely, according to (8) and (10), encoder \mathbf{E} maps M input elements y_1, \dots, y_M *jointly* into K indices $i_n \in \mathcal{I}$, $n = 1, \dots, K$. “Low-complexity” refers to the fact that quantization of a measurement vector \mathbf{y} requires only a single forward pass through EncNet involving matrix multiplications and activation function operations, followed by a simple SQ stage.

8. Since the SQs are identical, we will omit the neuron index “ n ” from quantization index $i \in \mathcal{I}$ whenever not explicitly needed.

Remark 3: Our proposed VQ-based DNN architecture is motivated by the fact that the optimal compression⁹ in QCS is achieved by constructing a minimum mean square error (MMSE) estimate of source \mathbf{x} at the encoder and compressing the estimate with a VQ [14], [20], [30]. The capability of our DNN-based VQ to realize such optimal compression is mainly dictated by the EncNet configuration: EncNet must surrogate an exponentially complex MMSE estimation stage [57] and a table-lookup stage. We demonstrate in Section V-B5 that DeepVQCS has the ability to realize near-optimal compression. Note that due to the considered QCS architecture, the design of DeepVQCS does not *explicitly* reflect the sparse nature of \mathbf{x} ; rather, the sparsity is an intrinsic element in optimizing the DNN structures so as to form the optimal QCS-aware VQ.

B. DECODER

Decoder \mathbf{D} comprises a quantizer decoder and the decoder DNN, described next.

1) QUANTIZER DECODER

At the first stage of decoder \mathbf{D} , n th quantizer decoder \mathbf{Q}_D^{-1} converts the received index $i_n \in \mathcal{I}$ into reproduction level g_{i_n} according to (7), $n = 1, \dots, K$. The aggregate dequantization operation of index sequence $\{i_1, \dots, i_K\}$ is defined as (see \mathbf{Q}_E in (9))

$$\bar{\mathbf{Q}}_D^{-1}(i_1, \dots, i_K) = [\mathbf{Q}_D^{-1}(i_1) \cdots \mathbf{Q}_D^{-1}(i_K)]^T = [g_{i_1} \cdots g_{i_K}]^T. \quad (11)$$

Combining (9) and (11), we denote the K separate quantization-dequantization operations applied for vector \mathbf{a}_J collectively as $\bar{\mathbf{Q}} = \{\bar{\mathbf{Q}}_E, \bar{\mathbf{Q}}_D^{-1}\}$. Borrowing the nomenclature of DNNs, we refer to $\bar{\mathbf{Q}}$ as the *quantization layer*.

2) DECODER DNN

As the second stage of decoder \mathbf{D} , the output of $\bar{\mathbf{Q}}_D^{-1}$ in (11) is fed into a feedforward fully connected L -layer decoder DNN, dubbed **DecNet**. The structure and operation of DecNet are as follows. Let $\mathbf{z}_l = [z_{l,1} \cdots z_{l,d_l}]^T \in \mathbb{R}^{d_l}$ be the weighted input of layer l , where d_l denotes the width of layer l . The weighted input at layer $l = 2, \dots, L$ is given as

$$\mathbf{z}_l = \mathbf{W}_l \mathbf{p}_{l-1} + \mathbf{r}_l, l = 2, \dots, L, \quad (12)$$

where $\mathbf{W}_l \in \mathbb{R}^{d_l \times d_{l-1}}$ is the weight matrix at layer l , $\mathbf{r}_l \in \mathbb{R}^{d_l}$ is the bias vector at layer l , and $\mathbf{p}_l \in \mathbb{R}^{d_l}$ is the output of layer l , defined as

$$\mathbf{p}_1 = \mathbf{g}, \quad \mathbf{p}_l = \sigma_l(\mathbf{z}_l) = \sigma_l(\mathbf{W}_l \mathbf{p}_{l-1} + \mathbf{r}_l), l = 2, \dots, L, \quad (13)$$

9. This optimal compression strategy is known as “estimate-and-compress”, and has been addressed under QCS in, e.g., [20], [55]. Another strategy is to estimate the support of \mathbf{x} at the encoder and then encode the support pattern losslessly while quantizing the obtained non-zero values; this has been considered in, e.g., [13], [21]. Stringently resource-limited encoding devices may necessitate to quantize the measurements directly, and run a quantization-aware CS reconstruction algorithm at the decoder; this “compress-and-estimate” QCS has been addressed in, e.g., [21], [56].

where $\sigma_l(\cdot)$ is an (element-wise) activation function at layer l ; $\mathbf{p}_1 = \mathbf{g}$ implies that the DecNet input is formed by reproduction levels $\mathbf{g} = [g_{i_1} \cdots g_{i_K}]^T$ obtained via $\bar{\mathbf{Q}}_D^{-1}$ in (11).

Owing to the estimation task, the DecNet output represents an estimate of source vector $\mathbf{x} \in \mathbb{R}^N$; thus, $d_L = N$. DecNet takes K reproduction levels g_{i_n} , $n = 1, \dots, K$, in (11) as its input and produces $\mathbf{p}_L \in \mathbb{R}^N$ of form (13) as an output. We define DecNet as a mapping

$$\Omega_D : \mathbb{R}^K \rightarrow \mathbb{R}^N, \quad \Omega_D(\mathbf{g}) = \mathbf{p}_L. \quad (14)$$

3) FULL OPERATION OF DECODER

Combining the operations of quantizer decoder $\bar{\mathbf{Q}}_D^{-1}$ in (11) and DecNet in (14), decoder \mathbf{D} can be expressed as a composite function

$$\mathbf{D} = \Omega_D \circ \bar{\mathbf{Q}}_D^{-1} : \mathcal{I}^K \rightarrow \mathbb{R}^N. \quad (15)$$

Finally, the full operation of decoder \mathbf{D} is expressed as

$$\mathbf{p}_L = \sigma_L \left(\mathbf{W}_L \left(\cdots \sigma_3 \left(\mathbf{W}_3 \left(\sigma_2 \left(\mathbf{W}_2 \left(\bar{\mathbf{Q}}_D^{-1}(i_1, \dots, i_K) \right) + \mathbf{r}_2 \right) \right) + \mathbf{r}_3 \right) \cdots \right) + \mathbf{r}_L \right). \quad (16)$$

C. IMPLEMENTATION ASPECTS

Several remarks regarding the implementation of the proposed DeepVQCS encoder-decoder architecture are in order. As per (8), EncNet processes measurements \mathbf{y} with real-valued numbers; the same holds true for reproduction levels \mathbf{g} at decoder \mathbf{D} as per (15). These assumptions can be invoked by various design considerations, discussed next.

1) HIGH-RESOLUTION ADC & DIGITAL ENCNET

The encoder device may be equipped with a high-resolution (e.g., 16-bit or 32-bit) ADC, when our model supports a *digital* implementation of EncNet. Encoder \mathbf{E} receives *finely* discretized measurements \mathbf{y} from the ADC, pre-processes \mathbf{y} digitally – this is *approximated* by Ω_E in (5) – and employs *coarse* (e.g., 1–8 bits) quantization of \mathbf{a}_j at (digital) quantizer encoder \mathbf{Q}_E in (6). Here, \mathbf{Q}_E primarily employs source compression. Note that only a moderate-sized EncNet might be viable at a low-power sensor, thereby offloading the sparse signal recovery task primarily to DecNet at a more computationally capable decoding device.

2) ANALOG ENCNET & LOW-RESOLUTION ADC

For low-cost digital devices, a low-resolution ADC precludes the above option. However, owing to the recent advances in neuromorphic computing systems, EncNet can be implemented on an *analog* circuit prior to a (possibly) low-resolution ADC via *memristors*¹⁰ [58], [60]. With exceptions of hardware non-idealities, a memristor implementation

10. Memristor-based mixed hardware–software implementations include a two-layer DNN in [58] and a four-layer fully connected DNN in [59]; fully hardware implementation of a five-layer convolutional DNN was constructed in [60]. Inference performance of memristive DNNs is similar to that of digital DNNs, yet their computational energy efficiency and throughput per unit area are orders of magnitude higher than those of the up-to-date graphical processing units (GPUs) [59].

of EncNet accurately complies with the encoder mapping in (8). In fact, the entire encoder can be implemented via a neuromorphic system: the ADC can be realized by a memristive neural network, providing also flexible training capabilities [61].

As per DecNet, we assume that the receiver is equipped with a high-resolution ADC so that the operations in (15) accurately model a digital DecNet. Another practicality is that digital implementation of a DNN necessitates quantizing the weights and biases.¹¹ Since our main focus is to address the presence of *coarse* quantization from the source compression and communication viewpoint, a specific implementation of EncNet and DecNet as well as the inaccuracies induced by digital DNN operations are outside of the main scope of this article and left for future work.

IV. JOINT OPTIMIZATION OF THE DEEP ENCODER-DECODER

In this section, we elaborate the optimization of the DeepVQCS encoder-decoder scheme via *stochastic gradient descent (SGD)* [28, Ch. 5.9] and *backpropagation* [29]. The optimization problem is formulated in Section IV-A. The technique to overcome the non-differentiability of quantization – *soft-to-hard quantization (SHQ)* – is detailed in Section IV-B. The SGD optimization steps are derived in Section IV-C. Asymptotic quantizer and gradient approximation strategies for the SHQ to facilitate training are proposed in Section IV-D. Quantizer construction is detailed in Section IV-E. A supervised training algorithm is summarized in Section IV-F.

A. PROBLEM FORMULATION

We reformulate the QCS problem in Definition 1 to incorporate the defined DeepVQCS system blocks. Let Γ_E and Γ_D be the parameter sets of EncNet and DecNet, respectively, defined as

$$\Gamma_E = \{\mathbf{F}_j, \mathbf{b}_j\}_{j=2}^J, \quad \Gamma_D = \{\mathbf{W}_l, \mathbf{r}_l\}_{l=2}^L. \quad (17)$$

The objective of minimizing MSE distortion D in (2) for a given measurement matrix Φ and quantization resolution $\bar{I} = KI$ is cast as a joint encoder-decoder optimization problem as

$$\begin{aligned} & \{\Gamma_E^*, \Gamma_D^*, \mathbf{t}^*, \mathbf{g}^*\} \\ &= \underset{\Gamma_E, \Gamma_D, \mathbf{t}, \mathbf{g}}{\operatorname{argmin}} D(\Gamma_E, \Gamma_D, \mathbf{t}, \mathbf{g}) \\ &\stackrel{(a)}{=} \underset{\Gamma_E, \Gamma_D, \mathbf{t}, \mathbf{g}}{\operatorname{argmin}} \\ &\quad \times \mathbb{E} \left[\left\| \Omega_D \left(\bar{\mathbf{Q}}_D^{-1} \left(\bar{\mathbf{Q}}_E[\Omega_E(\mathbf{y}; \Gamma_E); \mathbf{t}]; \mathbf{g} \right); \Gamma_D \right) - \mathbf{x} \right\|_2^2 \right], \end{aligned} \quad (18)$$

where (a) follows from the encoder and decoder mappings in (8) and (15), respectively, $\mathbf{t} = [t_1 \cdots t_{I-1}]^T$ is the vector

11. This *DNN quantization* is an active research area; see, e.g., [26], [27], [45]–[47].

of thresholds of \mathbf{Q} and $\mathbf{g} = [g_1 \cdots g_I]^T$ is the vector of reproduction levels of \mathbf{Q} (see Definition 2).

Finding the optimal parameters of EncNet and DecNet along with the optimal quantizer¹² in (18) seems *intractable*. This is due to the complicated and *non-differentiable* nature of quantizer \mathbf{Q} . In particular, the non-differentiability of \mathbf{Q} precludes the use of standard SGD to optimize the DeepVQCS scheme due to the *vanishing gradient problem*¹³ [25]–[27]: since the gradient of a quantization function vanishes almost everywhere, backpropagating “training information” to EncNet is impossible, and thus, EncNet cannot be effectively trained for the considered compression task. Next, we address how to overcome this hindrance at the quantization layer.

B. SOFT-TO-HARD QUANTIZATION

We overcome the incapability of SGD to handle the non-differentiable quantizer by *soft-to-hard quantization (SHQ)* [26], [27], [31], [47] – a *differentiable* approximation of “hard” quantizer \mathbf{Q} . More precisely, in the *offline* training phase, we remove quantization layer $\bar{\mathbf{Q}}$ and replace it by a *virtual* EncNet layer¹⁴ $j = J + 1$ which we call the *SHQ layer*. We adopt SHQ functionality at the SHQ layer to *approximate* the behavior of the quantizer that will be implemented in practice. After training, the virtual SHQ layer is removed and the SQ, \mathbf{Q} , is constructed using the final SHQ parameters. The premise is that once we optimize the DNN parameters – including the SHQ parameters – in a setup without $\bar{\mathbf{Q}}$, the obtained parameters are expected to provide similar performance after the SHQ layer is substituted by $\bar{\mathbf{Q}}$. We emphasize that the SHQ layer is present *only* during the training phase. The detailed description of the SHQ layer is given next.

Since the SHQ surrogates an SQ, the SHQ layer $j = J + 1$ has the width $e_{J+1} = e_J = K$, and it is connected to the EncNet output layer $j = J$ directly, i.e., $\mathbf{c}_{J+1} = \mathbf{a}_J$ (see (3)). According to (4), we have $\mathbf{a}_{J+1} = \gamma_{J+1}(\mathbf{c}_{J+1}) = \gamma_{J+1}(\mathbf{a}_J)$, and we model the activation function $\gamma_{J+1}(\cdot)$ as the SHQ function [31]; thus, n th SHQ output is given as

$$a_{J+1,n} = \gamma_{J+1}(a_{J,n}) = \sum_{i=1}^{I-1} v_i \tanh(h a_{J,n} - s_i), \quad n = 1, \dots, K, \quad (19)$$

where *level* coefficients $\mathbf{v} = [v_1 \cdots v_{I-1}]^T \in \mathbb{R}_+^{I-1}$, *shift* coefficients $\mathbf{s} = [s_1 \cdots s_{I-1}]^T \in \mathbb{R}^{I-1}$, and *steepness*

12. In general, finding optimal thresholds \mathbf{t}^* and reproduction levels \mathbf{g}^* of any quantizer *jointly* is difficult, if not *intractable*. A common approach is to optimize a quantizer via alternating optimization by the Lloyd-Max algorithm [62], [63].

13. Computing the gradient of a loss function with respect to the input of a *hard-thresholding* neuron (e.g., a quantizer) causes the vanishing gradient problem [25]–[27] for backpropagation. *Estimated gradients* have been proposed to overcome the issue; the most common one is a simple *straight-through estimator* [25] which amounts to bypassing the hard-thresholding module.

14. Representing the SHQ as a virtual layer allows us to use the unified DNN terminology established in Section III.




FIGURE 2. Illustration of the SHQ function in (19) for $I = 4$, level coefficients $\mathbf{v} = [0.15 \ 0.4 \ 0.45]^T$, shift coefficients $\mathbf{s} = h[-0.2 \ 0 \ 2/3]^T$, and steepness coefficient $h = [5, 12, 40]$. Adaptive adjustment of h plays a key role in optimizing the DeepVQCS scheme. A four-level non-uniform SQ with reproduction levels $\{-1, -0.7, 0.1, 1\}$ is depicted for comparison.

coefficient $h \in \mathbb{R}_+$ are *tunable* parameters. Recall that I represents the number of quantization levels of an actual SQ (see Section III-A2). The SHQ function is illustrated in Fig. 2.

The SHQ function in (19) approximates a *non-uniform* SQ as a weighted sum of shifted v_i -weighted hyperbolic tangents, with the input argument being scaled by h . This steepness coefficient h controls the *asymptotic* continuous-to-discrete mapping: the higher the value of h , the steeper the slope of $\tanh(h a_{J,n} - s_i)$ for a small input $a_{J,n}$. Thus, $\tanh(h a_{J,n} - s_i)$ saturates quickly to 1 (−1) for a small positive (negative) $a_{J,n}$, i.e., $\gamma_{J+1}(\cdot)$ operates like an I -level quantizer. Level coefficients \mathbf{v} play the role of reproduction levels $\mathbf{g} = [g_1 \cdots g_I]^T$ of quantizer \mathbf{Q} ; shift coefficients \mathbf{s} are analogous to thresholds $\mathbf{t} = [t_1 \cdots t_{I-1}]^T$ (see Definition 2). Construction of \mathbf{Q} from the SHQ parameters is detailed in Section IV-E.

Remark 4: Owing to the differentiability of the SHQ function in (19), one can optimize level coefficients \mathbf{v} and shift coefficients \mathbf{s} along with the other DNN parameters in (17) within a single end-to-end SGD loop. This amounts to optimizing the quantization regions and reproduction levels of a *non-uniform* SQ that will be implemented in the system.

Remark 5: Steepness coefficient h is a *hyperparameter* and it is thus excluded from the SGD optimization. Since the vanishing gradient problem is still present at the SHQ layer (manifested in Fig. 2 for $h = 40$), *gradual increase* of h plays a crucial role in our training procedure, as elaborated in Section IV-D.

C. STOCHASTIC GRADIENT DESCENT OPTIMIZATION

In this section, we first formulate the training objective and then, use backpropagation to derive the SGD expressions needed in training the DeepVQCS scheme.

1) TRAINING OBJECTIVE

We formulate the training objective by modifying (18) to incorporate the SHQ layer in (19) while accounting for the

removal of the quantizer. Let $\mathcal{D}_{\text{tr}} = \{\mathbf{x}_{\text{tr}}^k, \mathbf{y}_{\text{tr}}^k\}_{k=1}^{N_{\text{tr}}}$ be a *training data set* of N_{tr} source and measurement vectors sampled from their joint distribution. The *training cost function* is defined as

$$\begin{aligned} C(\Gamma_E, \Gamma_D, \mathbf{v}, \mathbf{s}) &= (1/N_{\text{tr}}) \sum_{k=1}^{N_{\text{tr}}} \left\| \Omega_D(\gamma_{J+1}(\Omega_E(\mathbf{y}_{\text{tr}}^k; \Gamma_E); \mathbf{v}, \mathbf{s}); \Gamma_D) - \mathbf{x}_{\text{tr}}^k \right\|_2^2 \\ &= (1/N_{\text{tr}}) \sum_{k=1}^{N_{\text{tr}}} \left\| \mathbf{p}_L^k - \mathbf{x}_{\text{tr}}^k \right\|_2^2, \end{aligned} \quad (20)$$

where $\mathbf{p}_L^k \in \mathbb{R}^N$ is the DecNet output associated with k th training sample; here, we have $\Omega_D(\mathbf{a}_{J+1})$ (instead of $\Omega_D(\mathbf{g})$ as per (14)) to account for the fact that n th SHQ output (19) is directly connected to the n th input of DecNet, i.e., $p_{1,n} = a_{J+1,n}$, $n = 1, \dots, K$.

The training objective for the DeepVQCS scheme is to find parameter sets $\{\Gamma_E, \Gamma_D\}$ in (17) and SHQ parameters $\{\mathbf{v}, \mathbf{s}\}$ in (19) that minimize the training cost in (20) by solving the joint encoder-decoder optimization problem

$$\{\Gamma_E^*, \Gamma_D^*, \mathbf{v}^*, \mathbf{s}^*\} = \underset{\Gamma_E, \Gamma_D, \mathbf{v}, \mathbf{s}}{\text{argmin}} C(\Gamma_E, \Gamma_D, \mathbf{v}, \mathbf{s}). \quad (21)$$

The problem (21) is solved by the SGD optimization, for which the required gradient updates are derived by the backpropagation. These are detailed in the next subsections.

2) COMPUTATION OF GRADIENTS VIA BACKPROPAGATION

To apply the SGD for problem (21), we need to compute the gradients of cost function¹⁵ C in (20) with respect to each DNN parameter set in (17) and SHQ parameters in (19). The crux of the backpropagation is to use the chain rule of the partial derivatives of C to interrelate associated gradients at two consecutive layers, enabling efficient computations. Let $\xi_j = [\xi_{j,1} \dots \xi_{j,e_j}]^T \in \mathbb{R}^{e_j}$ be the gradient of C with respect to the weighted input of layer j of EncNet, i.e., \mathbf{c}_j in (3), as

$$\xi_j = \nabla_{\mathbf{c}_j} C = \left[\frac{\partial C}{\partial c_{j,1}} \dots \frac{\partial C}{\partial c_{j,e_j}} \right]^T, \quad j = 2, \dots, J+1, \quad (22)$$

where $\frac{\partial C}{\partial c_{j,n}}$ denotes the partial derivative of C with respect to $c_{j,n}$. Similarly, we define $\delta_l = [\delta_{l,1} \dots \delta_{l,d_l}]^T \in \mathbb{R}^{d_l}$ as the gradient of C with respect to the weighted input of layer l of DecNet, i.e., \mathbf{z}_l in (12), as

$$\delta_l = \nabla_{\mathbf{z}_l} C = \left[\frac{\partial C}{\partial z_{l,1}} \dots \frac{\partial C}{\partial z_{l,d_l}} \right]^T, \quad l = 1, \dots, L. \quad (23)$$

Next, we present the gradients in (22) and (23) for each layer by traversing the DeepVQCS layers in the reverse order. For the DecNet output, the gradient (23) is $\delta_L = \nabla_{\mathbf{z}_L} C = 2(\mathbf{p}_L - \mathbf{x})$. For DecNet layers

$l = 1, \dots, L-1$, using the well-established backpropagation equations [29], [28, Alg. 6.4], gradients δ_l and δ_{l+1} are interrelated as

$$\delta_l = \mathbf{W}_{l+1}^T \delta_{l+1} \odot \sigma'_l(\mathbf{z}_l), \quad l = 1, \dots, L-1. \quad (24)$$

Focus now on the SHQ layer $j = J+1$. Due to the absence of the quantizer, its adjacent deeper layer is the first layer of DecNet, interconnected as $\mathbf{p}_1 = \mathbf{a}_{J+1}$. Therefore, ξ_{J+1} in (22) is expressed as a function of δ_1 in (23) as

$$\xi_{J+1} = \delta_1 \odot \gamma'_{J+1}(\mathbf{a}_J). \quad (25)$$

Finally, for EncNet layers $j = 2, \dots, J$, gradients ξ_j and ξ_{j+1} are interrelated as (see (24))

$$\xi_j = \mathbf{F}_{j+1}^T \xi_{j+1} \odot \gamma'_j(\mathbf{c}_j), \quad j = 2, \dots, J. \quad (26)$$

Next, we present the gradient of C with respect to each DNN parameter in (17) and SHQ parameters in (19) as a function of derived quantities $\{\xi_j\}_{j=2}^{J+1}$ and $\{\delta_l\}_{l=1}^L$. For weight matrices and bias vectors, the gradients are given as [29], [28, Alg. 6.4]

$$\begin{aligned} \nabla_{\mathbf{F}_j} C &= \xi_j \mathbf{a}_{j-1}^T, \quad \nabla_{\mathbf{b}_j} C = \xi_j, \quad j = 2, \dots, J \\ \nabla_{\mathbf{W}_l} C &= \delta_l \mathbf{p}_{l-1}^T, \quad \nabla_{\mathbf{r}_l} C = \delta_l, \quad l = 2, \dots, L. \end{aligned} \quad (27)$$

For the SHQ layer (19), the gradient for level coefficients $\nabla_{\mathbf{v}} C = \left[\frac{\partial C}{\partial v_1} \dots \frac{\partial C}{\partial v_{I-1}} \right]^T$ and the gradient for shift coefficients $\nabla_{\mathbf{s}} C = \left[\frac{\partial C}{\partial s_1} \dots \frac{\partial C}{\partial s_{I-1}} \right]^T$ are given as follows. The partial derivative of C with respect to level coefficient v_i is derived in the Appendix and is given as

$$\frac{\partial C}{\partial v_i} = \sum_{n=1}^K \delta_{1,n} \tanh(ha_{J,n} - s_i), \quad i = 1, \dots, I-1. \quad (28)$$

The partial derivative of C with respect to shift coefficient s_i is given as (see Appendix)

$$\frac{\partial C}{\partial s_i} = \sum_{n=1}^K \delta_{1,n} v_i \tanh'_{s_i}(ha_{J,n} - s_i), \quad i = 1, \dots, I-1, \quad (29)$$

where $\tanh'_{s_i}(\cdot)$ denotes the derivative of $\tanh(ha_{J,n} - s_i)$ with respect to s_i , given as $-4/(\exp\{ha_{J,n} - s_i\} + \exp\{-ha_{J,n} + s_i\})^2$.

Remark 6: The expression in (29) reveals that for large h , $\tanh'_{s_i}(ha_{J,n} - s_i)$, and consequently, $\frac{\partial C}{\partial s_i}$ are close to zero almost everywhere. Thus, optimization of \mathbf{s} may be difficult in practice.

3) MINI-BATCH SGD UPDATES

Above, we derived all required gradient expressions to apply SGD for each DNN parameter set in (17) and SHQ parameters in (19) to train the DeepVQCS scheme. As a practical means, we employ the *mini-batch SGD* [28, Ch. 5.9] as follows. Let $\mathcal{B}^{(t)}$ be a *mini-batch* at iteration t , which consists of $B \leq N_{\text{tr}}$ samples $\{\mathbf{x}^k, \mathbf{y}^k\}$ from training set \mathcal{D}_{tr} . Taking weight

15. To lighten the derivations, we drop the dependency of $C(\Gamma_E, \Gamma_D, \mathbf{v}, \mathbf{s})$ on its arguments.

matrix \mathbf{F}_j as an example, the mini-batch SGD updates are of the form:

$$\begin{aligned}\mathbf{F}_j^{(t+1)} &= \mathbf{F}_j^{(t)} - \Lambda_{\mathbf{F}_j}^{(t)} \odot \Delta_{\mathbf{F}_j}^{(t)}, \quad j = 2, \dots, J \\ &\stackrel{(a)}{=} \mathbf{F}_j^{(t)} - \Lambda_{\mathbf{F}_j}^{(t)} \odot (1/B) \sum_{k=1}^B \xi_j^{k,(t)} \mathbf{a}_{j-1}^{k,(t)T}, \quad (30)\end{aligned}$$

where superscript $t = 1, 2, \dots$ denotes the (SGD) iteration, $\Lambda_{\mathbf{F}_j}^{(t)} \in \mathbb{R}^{e_j \times e_{j-1}}$ is the step size matrix at iteration t , $\Delta_{\mathbf{F}_j}^{(t)} \in \mathbb{R}^{e_j \times e_{j-1}}$ is the stochastic gradient for weight matrix \mathbf{F}_j computed over mini-batch $\mathcal{B}^{(t)}$ at iteration t , $(\cdot)^{k,(t)}$ represents a quantity computed for k th sample at iteration t , and equality (a) follows from (27). Thus, $\Delta_{\mathbf{F}_j}^{(t)}$ approximates $\nabla_{\mathbf{F}_j} C$ in (27). The mini-batch SGD updates for the other DNN parameters can be derived similarly.

D. QUANTIZER AND GRADIENT APPROXIMATION AT THE SHQ LAYER

At the SHQ layer, steepness coefficient h in (19) trade offs between the smoothness of the EncNet-DecNet interface and the resemblance of an I -level quantizer. Clearly, a very large value of h brings the vanishing gradient problem (see Fig. 2 for $h = 40$) for (25), i.e., no training information flows from DecNet to EncNet, inhibiting the achievable performance. On the other hand, a small value of h creates a smooth transition between the EncNet output and DecNet input (see Fig. 2 for $h = 5$), passing an intact gradient flow to EncNet. However, the shortcoming is that an *over-relaxed* soft quantizer does not authentically represent the actual “hard” quantizer \mathbf{Q} , detrimental to the rate-distortion performance.

The aforementioned trade-off motivates to *gradually increase* the presence of quantization during training. To this end, we propose two strategies that are employed to facilitate the training of DeepVQCS: 1) *asymptotic quantizer approximation*¹⁶ that adjusts steepness coefficient h , and 2) *gradient approximation* that (re)adjusts the gradient pass through the SHQ layer. The crux is to asymptotically increase the degree of a continuous-to-(near)-discrete mapping; initially, an ample gradient flow trains EncNet for a *coarse* approximate quantizer, whereas in the course of iterations, the SHQ layer becomes an accurate replica of an I -level quantizer and *fine-tunes* DeepVQCS for the given quantization resolution. These two strategies are detailed next.

1) ASYMPTOTIC QUANTIZER APPROXIMATION

Steepness coefficient h in (19) is updated as

$$h^{(t)} = \min(h^{\text{init}} + \alpha^{(t)}, h^{\text{max}}), \quad (31)$$

where $\alpha^{(t)}$ is a step size, and parameters h^{init} and h^{max} set the initial and maximum value of h , respectively. A small

16. This is akin to *annealing*, a well-established strategy in quantization; see, e.g., the VQ design in [64]. In DNNs, annealing has been applied, e.g., for DNN quantization in [27] and for DNN model/data compression using the softmax operator in [26].

$h^{(t)}$ approximates an identity function (see Fig. 2 for $h = 5$), whereas increasing $h^{(t)}$ slowly to a large value approaches an I -level quantizer (see Fig. 2 for $h = 40$). The simulation results in Section V-B2 show this to be an efficient strategy to ameliorate training.

2) GRADIENT APPROXIMATION

Besides (31), we propose a gradual soft-to-hard transition for the backpropagating gradient through the SHQ layer. Recall that by (25), the gradient of C with respect to SHQ input \mathbf{a}_j is $\xi_{j+1} = \delta_1 \odot \gamma'_{j+1}(\mathbf{a}_j)$. We propose a gradient approximation policy that uses an *adjustable* weighted combination¹⁷ of the true gradient and the saturation-aware *straight-through estimator* (STE)¹⁸ [46]. Thus, at SGD iteration t , we have for k th sample:

$$\begin{aligned}\xi_{j+1}^{k,(t)} &= (1 - \beta^{(t)}) \left[\delta_1^{k,(t)} \odot \mathbf{1} \left\{ |\mathbf{a}_j^{k,(t)}| \leq \sum_{i=1}^{I-1} v_i \right\} \right] \\ &\quad + \beta^{(t)} \left[\delta_1^{k,(t)} \odot \gamma'_{j+1}(\mathbf{a}_j^{k,(t)}) \right], \quad (32)\end{aligned}$$

where $\beta^{(t)} = [0, 1]$ is a step size and binary vector $\mathbf{1}\{\cdot\} \in \mathbb{B}^K$ is an element-wise indicator function: its n th element is zero if the magnitude of SHQ input $a_{j,n}^{k,(t)}$ exceeds the SHQ output range $a_{j+1,n} \in [-\sum_{i=1}^{I-1} v_i, \sum_{i=1}^{I-1} v_i]$, $n = 1, \dots, K$ (see (33)), i.e., it nullifies the n th gradient entry. For a small $\beta^{(t)}$, (32) at early iterations tends to the STE as $\xi_{j+1}^{k,(t)} \approx \delta_1^{k,(t)}$. One intuition at moderate values of $\beta^{(t)}$ is that the STE term of (32) keeps passing “noisy gradient” for a coarse training of EncNet, overriding the fact that the true gradient term of (32) has small values around the progressively emerging flat regions of the SHQ. Finally, $\beta^{(t)} \rightarrow 1$ ensures that the true gradient is used towards the end of training, which, along with large h , refines DecNet for quantization resolution I . In our conducted numerical experiments, combination of (31) and (32) with appropriate learning schedules for $\alpha^{(t)}$ and $\beta^{(t)}$ yielded the most robust training behavior.

E. QUANTIZER CONSTRUCTION

Once the DeepVQCS scheme has been trained, the SHQ layer $j = J + 1$ will be removed and quantizer $\mathbf{Q} = \{\mathbf{Q}_E, \mathbf{Q}_D^{-1}\}$ is implemented in the system. Since we devoted the trainable SHQ layer to approximate an I -level quantizer (advocated by the policies of Section IV-D), we use directly the learned SHQ parameters to construct quantizer \mathbf{Q} as follows.

After training, the SHQ outputs $a_{j+1,n}$ concentrate around I discrete values which are dictated by level coefficients $\{v_i\}_{i=1}^I$ (in Fig. 2, these are $\{-1, -0.7, 0.1, 1\}$); we place reproduction levels $\{g_i\}_{i=1}^I$ of \mathbf{Q} to coincide with these saturation values. Thresholds $\{t_i\}_{i=1}^{I-1}$ of \mathbf{Q} are set as $t_i = s_i/h$

17. A weighted combination for gradually increasing the impact of quantization in backpropagation is used by, e.g., the “alpha-blending” method [27] developed to optimize low-precision representations of a DNN model.

18. While simple, STE has empirically been shown to be a viable means in training [25].

by invoking the fact that $\tanh(ha_{J,n} - s_i) = 0$, if $a_{J,n} = s_i/h$: each threshold coincides with a (nearly) vertical step occurring at an input value $a_{J,n} = s_i/h$ (in Fig. 2, these points are $\{-0.2, 0, 2/3\}$). Formally, assuming without loss of generality that $v_1 \leq v_2 \leq \dots \leq v_{I-1}$ and $s_1 \leq s_2 \leq \dots \leq s_{I-1}$, the quantizer \mathbf{Q} is constructed as

$$g_i = \begin{cases} -\sum_{i'=1}^{I-1} v_{i'}, & i = 1 \\ \sum_{i'=1}^{I-1} v_{i'} - 2\sum_{i'=i}^{I-1} v_{i'}, & 2 \leq i \leq I-1 \\ \sum_{i'=1}^{I-1} v_{i'}, & i = I, \end{cases} \quad (33)$$

$$t_i = s_i/h, \quad i = 1, \dots, I-1.$$

Note that the mismatch between the SHQ and quantizer \mathbf{Q} constructed according to (33) vanishes when h is sufficiently large (see Fig. 2), and thus, the additional distortion incurred by implementing the actual quantizer in the system becomes minimal.

F. SUPERVISED LEARNING ALGORITHM

A practical mini-batch SGD algorithm to train the DeepVQCS scheme in a supervised fashion is summarized in Algorithm 1. At each iteration t , training involves a *forward pass* and a *backward pass*, summarized in Algorithm 2 and Algorithm 3, respectively. At iteration t , the rate-distortion performance of the current DeepVQCS scheme with $\{\Gamma_E^{(t)}, \mathbf{t}^{(t)}, \mathbf{g}^{(t)}, \Gamma_D^{(t)}\}$ can be evaluated using a validation set $\mathcal{D}_{va} = \{\mathbf{x}_{va}^k, \mathbf{y}_{va}^k\}_{k=1}^{N_{va}}$ (or test set \mathcal{D}_{te}) as (see (18))

$$\tilde{D}_{va} = (1/N_{va}) \sum_{k=1}^{N_{va}} \|\mathbf{p}_L^{k,(t)} - \mathbf{x}_{va}^k\|_2^2 = (1/N_{va}) \sum_{k=1}^{N_{va}} \left\| \Omega_D \left[\tilde{\mathbf{Q}}_D^{-1} \left(\tilde{\mathbf{Q}}_E \left[\Omega_E \left(\mathbf{y}_{va}^k; \Gamma_E^{(t)} \right); \mathbf{t}^{(t)} \right]; \mathbf{g}^{(t)} \right); \Gamma_D^{(t)} \right] - \mathbf{x}_{va}^k \right\|_2^2. \quad (34)$$

The time and computation cost of the training phase can be high, typical to supervised learning. Thus, the entire training of the encoder and decoder is to be performed *offline* at a computationally capable entity, e.g., a general-purpose computer. Once trained, the DeepVQCS scheme communicates a measurement vector \mathbf{y} using only a *single forward pass* in Algorithm 2 (with Step 7 replaced by quantizer $\tilde{\mathbf{Q}}$). As this involves only matrix multiplications and activation function operations, the proposed DeepVQCS scheme has a fast, low-complexity encoding-decoding stage, enabling to process time-sensitive large-scale data. To assess the computational complexity and latency of the proposed method, the algorithm running time of the online phase is evaluated in Section V-B6.

V. SIMULATION RESULTS

Simulation results are presented to assess the rate-distortion performance and algorithm time complexity of the proposed DeepVQCS scheme summarized in Algorithm 1. The DeepVQCS scheme as well as the considered baseline methods were implemented in MATLAB.

Algorithm 1 DeepVQCS Training via SGD

- 1: **Input:** 1) Measurement matrix Φ ; 2) quantization levels I ; 3) DNN configurations $J, L, K, \{\epsilon_j, \gamma_j\}_{j=1}^J$, and $\{d_l, \sigma_l\}_{l=1}^L$; 4) data sets $\mathcal{D}_{tr} = \{\mathbf{x}_{tr}^k, \mathbf{y}_{tr}^k\}_{k=1}^{N_{tr}}$, $\mathcal{D}_{va} = \{\mathbf{x}_{va}^k, \mathbf{y}_{va}^k\}_{k=1}^{N_{va}}$, and $\mathcal{D}_{te} = \{\mathbf{x}_{te}^k, \mathbf{y}_{te}^k\}_{k=1}^{N_{te}}$
- 2: Set SGD iteration index as $t = 1$
- 3: **while** stopping criteria are not met **do** ▷ Training
- 4: Generate a mini-batch $\mathcal{B}^{(t)}$ from \mathcal{D}_{tr}
- 5: Run Algorithm 2 with mini-batch $\mathcal{B}^{(t)}$ ▷ Forward pass
- 6: Run Algorithm 3 with mini-batch $\mathcal{B}^{(t)}$ ▷ Backward pass
- 7: a) Construct $\mathbf{Q} = \{\mathbf{Q}_E, \mathbf{Q}_D^{-1}\}$ in (33), b) run Algorithm 2 for \mathcal{D}_{va} by replacing Step 7 with $\tilde{\mathbf{Q}}$, and c) evaluate \tilde{D}_{va} in (34) ▷ Validation
- 8: Set $t = t + 1$
- 9: **end while**
- 10: Evaluate \tilde{D}_{te} in (34) using test set \mathcal{D}_{te} ▷ Testing
- 11: **Output:** DeepVQCS encoder-decoder architecture with estimated performance \tilde{D}_{te}

Algorithm 2 Forward Pass at Iteration t

- 1: **Input:** Mini-batch $\mathcal{B}^{(t)}$
- 2: **for** mini-batch sample $k = 1, \dots, B$ **do**
- 3: EncNet input: $\mathbf{a}_1^{k,(t)} = \mathbf{y}^k$ ▷ Encoder E
- 4: **for** EncNet layer $j = 2, \dots, J$ **do** ▷ EncNet
- 5: $\mathbf{c}_j^{k,(t)} = \mathbf{F}_j^{(t)} \mathbf{a}_{j-1}^{k,(t)} + \mathbf{b}_j^{(t)}$, $\mathbf{a}_j^{k,(t)} = \gamma_j(\mathbf{c}_j^{k,(t)})$
- 6: **end for**
- 7: $\mathbf{a}_{J+1,n}^{k,(t)} = \sum_{i=1}^{I-1} v_i^{(t)} \tanh(h^{(t)} a_{J,n}^{k,(t)} - s_i^{(t)})$, $n = 1, \dots, K$ ▷ SHQ layer
- 8: DecNet input: $\mathbf{p}_1^{k,(t)} = \mathbf{a}_{J+1}^{k,(t)}$ ▷ Decoder D
- 9: **for** DecNet layer $l = 2, \dots, L$ **do** ▷ DecNet
- 10: $\mathbf{z}_l^{k,(t)} = \mathbf{W}_l^{(t)} \mathbf{p}_{l-1}^{k,(t)} + \mathbf{r}_l^{(t)}$, $\mathbf{p}_l^{k,(t)} = \sigma_l(\mathbf{z}_l^{k,(t)})$
- 11: **end for**
- 12: **end for**
- 13: **Output:** EncNet: $\{\mathbf{a}_j^{k,(t)}, \mathbf{c}_j^{k,(t)}\}_{j=1}^J$; SHQ: $\mathbf{a}_{J+1}^{k,(t)}$; DecNet: $\{\mathbf{p}_l^{k,(t)}, \mathbf{z}_l^{k,(t)}\}_{l=1}^L$

A. SIMULATION SETUP

The simulation setup for the experiments is set as follows, unless otherwise stated.

1) SIGNAL MODEL

For the CS setup in (1), we consider that 1) each non-zero entry of \mathbf{x} is Gaussian $\mathcal{N}(0, 1)$, 2) the sparsity patterns are uniformly distributed, 3) each measurement noise entry is Gaussian $\mathcal{N}(0, \sigma_n^2)$ with $\sigma_n^2 = 10^{-4}$, and 4) Φ is generated by taking the first M rows of an $N \times N$ discrete cosine transform matrix and normalizing the columns as $\|\cdot\|_2 = 1$.

2) DEEPVQCS

EncNet has $J = 3$ layers with $e_2 = 5K$. DecNet has $L = 5$ layers with $d_2 = d_3 = d_4 = 4N$. The SHQ layer width is $K = M$. Activation functions $\{\gamma_j\}_{j=2}^J$ and $\{\sigma_l\}_{l=2}^{L-1}$ are $\tanh(\cdot)$; γ_1 , σ_1 , and σ_L are identity functions. Each entry of weight matrix $\{\mathbf{F}_j\}_{j=2}^J$ (\mathbf{W}_l) is initialized by the Xavier initialization as $\mathcal{N}(0, 1/\epsilon_{j-1})$ [65]. The bias vectors are initialized as zero vectors. The level coefficients are initialized as $\mathbf{v} = \frac{0.8}{I-1} \mathbf{1}$.

Algorithm 3 Backward Pass at Iteration t

```

1: Input: 1) Mini-batch  $\mathcal{B}^{(t)}$ ; 2) EncNet:  $\{\mathbf{a}_j^{k,(t)}, \mathbf{c}_j^{k,(t)}\}_{j=1}^J$ ; 3)
   SHQ:  $\mathbf{a}_{j+1}^{k,(t)}$ ; 4) DecNet:  $\{\mathbf{p}_l^{k,(t)}, \mathbf{z}_l^{k,(t)}\}_{l=1}^L$ 
2: for DecNet layer  $l = L, \dots, 1$  do ▷ Decoder D
3:   if  $l = L$  then ▷ DecNet
4:      $\delta_L^{k,(t)} = 2(\mathbf{p}_L^{k,(t)} - \mathbf{x}^{(t)}), \forall k = 1, \dots, B$ 
5:   else
6:      $\delta_l^{k,(t)} = \mathbf{W}_{l+1}^{(t)T} \delta_{l+1}^{k,(t)} \odot \sigma'_l(\mathbf{z}_l^{k,(t)}), \forall k = 1, \dots, B$ 
7:   end if
8:   if  $l > 1$  then
9:      $\mathbf{W}_l^{(t+1)} = \mathbf{W}_l^{(t)} - \Lambda_{\mathbf{W}_l}^{(t)} \odot \Delta_{\mathbf{W}_l}^{(t)}, \mathbf{r}_l^{(t+1)} = \mathbf{r}_l^{(t)} - \lambda_{\mathbf{r}_l}^{(t)} \odot \Delta_{\mathbf{r}_l}^{(t)}$ 
10:  end if
11: end for
12:  $\mathbf{v}^{(t+1)} = \mathbf{v}^{(t)} - \lambda_{\mathbf{v}}^{(t)} \odot \Delta_{\mathbf{v}}^{(t)}$  ▷ Encoder E
    $\mathbf{s}^{(t+1)} = \mathbf{s}^{(t)} - \lambda_{\mathbf{s}}^{(t)} \odot \Delta_{\mathbf{s}}^{(t)}$ 
13:  $\xi_{j+1}^{k,(t)} = (1 - \beta^{(t)})[\delta_1^{k,(t)} \odot \mathbf{1}\{|\mathbf{a}_j^{k,(t)}| \leq \sum_{i=1}^{I-1} v_i\}] +$ 
    $\beta^{(t)}[\delta_1^{k,(t)} \odot \gamma'_{j+1}(\mathbf{a}_j^{k,(t)})], \forall k = 1, \dots, B$  ▷ SHQ layer
14: for EncNet layer  $j = J, \dots, 2$  do ▷ EncNet
15:    $\xi_j^{k,(t)} = \mathbf{F}_{j+1}^{(t)T} \xi_{j+1}^{k,(t)} \odot \gamma'_j(\mathbf{c}_j^{k,(t)}), \forall k = 1, \dots, B$ 
16:    $\mathbf{F}_j^{(t+1)} = \mathbf{F}_j^{(t)} - \Lambda_{\mathbf{F}_j}^{(t)} \odot \Delta_{\mathbf{F}_j}^{(t)}, \mathbf{b}_j^{(t+1)} = \mathbf{b}_j^{(t)} - \lambda_{\mathbf{b}_j}^{(t)} \odot \Delta_{\mathbf{b}_j}^{(t)}$ 
17: end for
18: Output: EncNet:  $\{\mathbf{F}_j^{(t+1)}, \mathbf{b}_j^{(t+1)}\}_{j=2}^J$ ;
   SHQ:  $\{\mathbf{v}^{(t+1)}, \mathbf{s}^{(t+1)}\}$ ; DecNet:  $\{\mathbf{W}_l^{(t+1)}, \mathbf{r}_l^{(t+1)}\}_{l=2}^L$ 

```

For $I = 2$, the shift coefficients are fixed to $\mathbf{s} = \mathbf{0}$; for $I > 2$, the shifts are adjusted¹⁹ as $\mathbf{s}^{(t)} = h^{(t)}[-0.8 : \frac{1.6}{I-2} : 0.8]^\top$. The mini-batch size is $B = 100$, and the data set sizes are $N_{\text{tr}} = 5 \times 10^5$ and $N_{\text{va}} = N_{\text{te}} = 3 \times 10^5$. For (31) and (32), we use linear step size schedules as $\alpha^{(t)} = \alpha t$ and $\beta^{(t)} = \min(\beta t, 1)$ with $h^{\text{init}} = 5$, $\alpha = 10^{-5}$, $h^{\text{max}} = 300$, and $\beta = 10^{-7}$. The step size for each DNN parameter is set by the Adam optimizer [66, Alg. 1] and diminishing learning schedule as $(\mathbf{F}_j$ as an example) $\Lambda_{\mathbf{F}_j}^{(t)} = \max\{\eta_{\mathbf{F}}^{\min}, \eta_{\mathbf{F}}/\sqrt{t}\} \mathcal{A}(\Delta_{\mathbf{F}_j}^{(1)}, \dots, \Delta_{\mathbf{F}_j}^{(t)}; \beta_1, \beta_2, \epsilon)$, where parameters $\eta_{\mathbf{F}}$ and $\eta_{\mathbf{F}}^{\min}$ adjust the initial and minimum step size, respectively; Adam $\mathcal{A}(\cdot)$ is run with the “default” parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. For weight matrices and bias vectors, we use $\eta_{(\cdot)} = 10^{-2}$ and $\eta_{(\cdot)}^{\min} = 10^{-4}$; for the level coefficients, we use $\eta_{\mathbf{v}} = 5 \times 10^{-5}$ and $\eta_{\mathbf{v}}^{\min} = 5 \times 10^{-7}$.

Given a signal setup (N , M , and S), the DeepVQCS scheme is (only) empirically tuned in that the chosen learning parameters and DNN configurations (the widths, depths, activation functions etc.) remain *fixed* across the quantization rates. The SGD iterations are repeated until \tilde{D}_{va} does not significantly decrease or the maximum number of iterations 10^7 is reached.

3) BASELINE QCS METHODS

- A *compress-and-estimate* (CE) QCS method [20], [21] where 1) the encoder quantizes measurements \mathbf{y} in (1)

19. As pointed out in Remark 6, optimizing $\{s_i\}_{i=1}^{I-1}$ becomes challenging for large h . For the conducted experiments, we found that increasing s_i proportional to $h^{(t)}$ to preserve the ratio s_i/h and thus, to ensure well-separated SHQ regions (see Fig. 2 and (33)) resulted in the best performance.

oblivious to \mathbf{x} , and 2) the decoder estimates \mathbf{x} from quantized measurements $\tilde{\mathbf{y}} \in \mathbb{R}^M$ through a quadratically constrained polynomial-complexity basis pursuit (BP) problem²⁰ $\min_{\mathbf{x} \in \mathbb{R}^N} \|\mathbf{x}\|_1$ s.t. $\|\tilde{\mathbf{y}} - \Phi \mathbf{x}\|_2 \leq \mu_{\text{qc}}$. Three variants are considered: 1) CE-USQ-L1 with uniform SQ (USQ), 2) CE-SQ-L1 with an SQ that is optimized to minimize the quantization distortion via the Lloyd algorithm [63], and 3) CE-VQ-L1 with a Lloyd-optimized VQ. We use $\mu_{\text{qc}} = \sqrt{\sigma_{\mathbf{n}}}(1 + 1/I)$, which is verified in Section V-B1.

- A low-complexity USQ-based CE method, CE-USQ-OMP, that estimates \mathbf{x} via (greedy) orthogonal matching pursuit (OMP) [70] with *known sparsity* S .
- A DNN-based CE method, CE-DecNet, where 1) the encoder uses SQ, and 2) the decoder estimates \mathbf{x} via the decoder DNN, DecNet; we train CE-DecNet similarly as DeepVQCS in Algorithm 1 but without EncNet. This “SQ+DNN” architecture resembles that of “BW-NQ-DNN” [24]; however, a major difference is that “BW-NQ-DNN” optimizes Φ , which is not applicable in our remote sensing setup.
- An *estimate-and-compress* (EC) QCS method, EC-VQ [20], [21], where 1) the encoder forms an MMSE estimate of \mathbf{x} from \mathbf{y} which is an exponentially complex task [57], and 2) quantizes the resulting estimate with a Lloyd-optimized VQ. The EC strategy is known to be the *optimal* compression strategy for remote source coding [14], [30], while suffering from its high complexity.
- The *remote rate-distortion function* (RDF) of \mathbf{x} , generated by the modified Blahut-Arimoto algorithm in [20, Alg. 1]; this is an information-theoretic *lower bound* to any QCS method.

4) PERFORMANCE METRICS

Reconstruction accuracy is measured as the normalized MSE (NMSE) as $10 \log_{10}(\mathbb{E}[\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2] / \mathbb{E}[\|\mathbf{x}\|_2^2])$ (dB), where $\hat{\mathbf{x}} \in \mathbb{R}^N$ represents a source estimate. The rate is measured as $R = R_{\text{tot}}/N$ (bits), where R_{tot} is the total number of bits a QCS method uses to compress an encoder input \mathbf{y} . For the ease of exposition, we consider that DeepVQCS employs independent coding of K indices $\{i_1, \dots, i_K\}$ and thus, spends $R_{\text{tot}} = K \lceil \log_2 I \rceil$ bits.

B. SIMULATION RESULTS

1) COMPARISON TO BASELINES

Fig. 3 depicts the rate-distortion performance of the DeepVQCS scheme against several baseline QCS methods for $N = 20$, $M = 10$, and $S = 2$. The proposed DeepVQCS scheme significantly outperforms the baseline QCS methods,

20. The BP problem is solved via the ℓ_1 -MAGIC package [67] using “llqc_logbarrier.m” with stopping parameter 10^{-3} . The problem is equivalent to the well-known basis pursuit denoising (BPDN) [68] $\min_{\mathbf{x} \in \mathbb{R}^N} \mu_{\text{bp}} \|\mathbf{x}\|_1 + \|\tilde{\mathbf{y}} - \Phi \mathbf{x}\|_2^2$ for certain parameters μ_{qc} and μ_{bp} [69, Proposition 3.2.].




FIGURE 3. Rate-distortion performance of the proposed DeepVQCS method versus baseline QCS methods for $N = 20$, $M = 10$, and $S = 2$. DeepVQCS outperforms the considered baseline methods. The figure illustrates the importance of appropriately handling the vanishing gradient problem at the quantization layer.

which are ranked in the ascending order of performance as CE-USQ-OMP, CE-USQ-L1, CE-SQ-L1, and CE-DecNet. The CE-USQ-L1 and CE-SQ-L1 methods nearly coincide, indicating that SQ optimization provides negligible gain. Thus, we use CE-USQ-L1 instead of CE-SQ-L1 in sequel.

Second, we verified the choice $\mu_{qc} = \sqrt{\sigma_n}(1 + 1/I)$ for CE-USQ-L1 as follows. For each test sample $\{\mathbf{x}_{te}^k, \mathbf{y}_{te}^k\}$, $k = 1, \dots, 5 \times 10^5$, we ran the BP decoder for 56 different values $\mu_{qc} = \{10^{-5}, 10^{-4.9}, \dots, 10^{0.4}, 10^{0.5}\}$, and read off the minimum MSE – using the knowledge of \mathbf{x} – among the candidate solutions. As Fig. 3 shows, this *genie-aided* CE-USQ-L1* variant provides only small improvement, corroborating a valid choice of μ_{qc} for CE-USQ-L1.

Third, the CE-DecNet method outperforms the standard CE methods, substantiating the high potential of a DNN to replace a polynomial-complexity decoder in a QCS setup. However, because CE-DecNet confines to use SQ, the gap to the proposed VQ-based DeepVQCS scheme is immense: CE-DecNet achieves its minimum NMSE of around -16.7 dB for $R = 4.0$ bits, whereas DeepVQCS achieves the same NMSE with more than 2.5 times fewer bits, $R = 1.55$. The efficacy of VQ in the DeepVQCS scheme is evident in that the slope of the decay of NMSE is unrivalled; also, for the considered range of R , saturation is not yet encountered.

2) GRADIENT PASS STRATEGIES

Fig. 3 also illustrates the impact of different gradient pass strategies at the SHQ layer for the DeepVQCS scheme. Modifying (31) and (32), we consider four DeepVQCS variants: 1) DeepVQCS-STE with the saturation-aware STE [46] and no gradual increase of h with $h^{\text{init}} = h^{\text{max}} = 400$ and $\beta = 0$; 2) DeepVQCS-Q with using only the asymptotic quantizer approximation (31) with a modified step size schedule $\alpha^{(i)} = 0.05 \lceil t/100 \rceil$ with $h^{\text{init}} = 5$, $h^{\text{max}} = 300$, and $\beta = 1$; 3) DeepVQCS-QG with using both the quantizer and

gradient approximation (31) and (32) with “fast” step size schedules $h^{\text{init}} = 5$, $\alpha = 10^{-4}$, $h^{\text{max}} = 300$, and $\beta = 10^{-6}$; and 4) our standard DeepVQCS setting employing both (31) and (32) with “slow” step size schedules $h^{\text{init}} = 5$, $\alpha = 10^{-5}$, $h^{\text{max}} = 300$, and $\beta = 10^{-7}$.

Fig. 3 shows the benefits of the proposed strategies for the SHQ layer in (31) and (32) in that they provide the best rate-distortion performance. Using only the gradual increase of h as per (31) is a viable option, albeit DeepVQCS-Q encountered slightly unstable behavior at high rates. The STE cannot provide authentic training information to EncNet; the shortcoming of DeepVQCS-STE is logically more pronounced for low rates. Similarly, too rapid increase of quantizer presence inhibits the performance for DeepVQCS-QG. While quantitative comparison is not present, we found throughout our experiments that using the combination of (31) and (32) provided the most robust convergence with least sensitive choices of the learning parameters.

3) SCALABILITY TO DIFFERENT SIGNAL LENGTHS

Fig. 4 shows the rate-distortion performance of the DeepVQCS method versus the CE-USQ-L1 method in four different signal settings. The proposed DeepVQCS method scales well to setups of different signal lengths, outperforming CE-USQ-L1 in all setups; the gap reduces when N , M , and S increase.

4) NUMBER OF QUANTIZERS

Fig. 4(b) shows the performance of the DeepVQCS scheme for different widths of the EncNet output $K = \{10, 15, 20\}$, illustrating the trade-off for having either a few high-resolution SQs or several low-resolution SQs. The figure shows that the proposed DeepVQCS scheme is flexible in terms of K ; same rate-distortion performance can be achieved via multiple quantizer configurations. This can be beneficial for practical implementations having different hardware/operational constraints on the quantization stage.

5) RATE-DISTORTION LIMITS

To compare the DeepVQCS method against the rate-distortion limits of QCS, we consider the experiment in [20, Fig. 5(c)] with $N = 7$, $M = 4$, $S = 1$, and $\sigma_n^2 = 0.01$. We use $J = 4$, $e_2 = e_3 = 5K$, $L = 4$, $d_2 = d_3 = d_4 = 5N$, $\alpha = 10^{-3}$, and $\beta = 10^{-6}$. We consider a special structure for EncNet: its output layer is formed by K 1-bit SQs. Thus, $I = 2$ and the SHQ function in (19) is a single (weighted) $\tanh(\cdot)$. Note that although the signal setup is small, it allows to elucidate the fundamental compression capabilities of DeepVQCS.

Fig. 5 shows the rate-distortion performance of the DeepVQCS method against the baselines and theoretical limits of QCS. The SQ-based CE-DecNet method slightly outperforms CE-USQ-L1*. The figure corroborates the ability of the proposed DeepVQCS scheme to realize near-optimal compression: by employing VQ through the cascade of EncNet and SQ, DeepVQCS takes a significant leap




FIGURE 4. Rate-distortion performance of the proposed DeepVQCS method versus the baseline CE-USQ-L1 method for (a) $N = 20$, $M = 10$, and $S = 2$; (b) $N = 30$, $M = 15$, and $S = 3$; (c) $N = 60$, $M = 30$, and $S = 6$; (d) $N = 80$, $M = 40$, and $S = 8$. In (b), the width of EncNet output is varied as $K = \{10, 15, 20\}$; for the other plots, we have $K = M$. DeepVQCS scales well and outperforms the baseline in all signal setups.

TABLE 1. Algorithm Running Time Comparison in the Online Phase: Encoding/Decoding/Total Time of CE-USQ-OMP (First row) and CE-USQ-L1 (Second row) Normalized With Respect to Those of the DeepVQCS Scheme

$N = 20, M = 10$, and $S = 2$			$N = 80, M = 40$, and $S = 8$			$N = 160, M = 80$, and $S = 16$		
$R = 1$	$R = 2.5$	$R = 4$	$R = 1$	$R = 2.5$	$R = 4$	$R = 1$	$R = 2.5$	$R = 4$
0.4/2.5/1.8	0.5/2.6/1.8	0.7/2.6/1.6	0.2/3.4/2.6	0.2/3.4/2.5	0.4/3.5/1.9	0.1/7.6/6.6	0.2/7.6/6.3	0.5/7.4/5.3
0.4/57/38	0.5/58/36	0.7/58/26	0.2/53/40	0.2/54/39	0.4/54/27	0.1/38/33	0.2/37/31	0.5/37/26

from CE-DecNet and performs close to the tabular-search VQ-based CE-VQ-L1 and EC-VQ methods. Recall that EC-VQ involves exponentially complex MMSE estimation at the encoder, which becomes prohibitive for large-scale signals. In fact, the best anticipated performance for DeepVQCS is to match with EC-VQ as they both apply VQ of a single vector \mathbf{y} at a time; note that the remote RDF is portrayed as the information-theoretic limit of QCS that can be achieved only via (excessively complex) VQ of multiple vector inputs \mathbf{y} [20].

6) ALGORITHM RUNNING TIME

To assess the computational complexity and latency for communicating a measurement vector \mathbf{y} , Table 1 compares

the algorithm running time²¹ in the online phase for three different signal setups. The decoding and total running times of the proposed DeepVQCS scheme are around 35 – 60 and 25 – 40 times lower than those of the polynomial-complexity CE-USQ-L1 method, respectively. The DeepVQCS scheme is faster than the greedy CE-USQ-OMP method (which runs only S loops). Additional pre-processing via EncNet inevitably increases the encoding time of DeepVQCS. Note, however, that the encoding time of each algorithm is small in proportion to the decoding time, i.e., the decoding time dominates the total latency incurred in the encoding-decoding process.

²¹ Algorithm running time was evaluated using “tic” function in MATLAB.




FIGURE 5. Rate-distortion performance of the proposed DeepVQCS method versus practical QCS methods and the information-theoretic limit of QCS ("Remote RDF") for $N = 7$, $M = 4$, and $S = 1$. The performance of DeepVQCS is close to that of EC-VQ which represents the optimal compression achievable in QCS by fixed-rate VQ of a measurement vector \mathbf{y} .

To summarize the findings from the conducted experiments, the proposed VQ-based DeepVQCS method obtains superior rate-distortion performance with orders of magnitude lower algorithm running time as compared to the conventional QCS methods, rendering DeepVQCS a potential method for finite-rate communication of sparse signals with resource-limited encoding devices.

VI. CONCLUSION

We proposed the DeepVQCS architecture, consisting of the encoder DNN, quantizer, and decoder DNN, for low-complexity acquisition of sparse sources through vector quantized noisy compressive measurements. A supervised SGD learning algorithm and techniques to overcome the non-differentiability of quantization were proposed for training the DeepVQCS scheme. Simulation results showed the superior rate-distortion performance and algorithm complexity of the proposed DeepVQCS scheme compared to standard QCS methods. These are desirable features to make DeepVQCS as a potential candidate for rate-limited communication of sparse signals through QCS under limited encoding and decoding capabilities.

The present study opens several research avenues. First, it would be interesting to test the proposed method using real-world sparse signals to obtain insights about the performance in a practical scenario. As potential future work, different encoder/decoder DNN types could be considered. The system could be extended to incorporate lossy communication channels, calling for a design of a channel-aware QCS method to counteract the erroneous transmissions of the code words. An extension to a distributed QCS setting with multiple encoding devices has its relevance to model, e.g., a practical wireless sensor network.

APPENDIX

The partial derivative of cost function C in (20) with respect to level coefficient v_i , $i = 1, \dots, I - 1$, is derived as

$$\begin{aligned} \frac{\partial C}{\partial v_i} &= \sum_{n=1}^K \sum_{n'=1}^K \frac{\partial C}{\partial z_{1,n'}} \frac{\partial z_{1,n'}}{\partial a_{J+1,n}} \frac{\partial a_{J+1,n}}{\partial v_i} \\ &\stackrel{(a)}{=} \sum_{n=1}^K \sum_{n'=1}^K \frac{\partial C}{\partial z_{1,n'}} \frac{\partial z_{1,n'}}{\partial a_{J+1,n}} \tanh(ha_{J,n} - s_i) \\ &\stackrel{(b)}{=} \sum_{n=1}^K \delta_{1,n} \tanh(ha_{J,n} - s_i), \end{aligned} \quad (35)$$

where equality (a) follows from the differentiation of the SHQ function with respect to v_i as $\frac{\partial a_{J+1,n}}{\partial v_i} = \frac{\partial \gamma_{J+1}(a_{J,n})}{\partial v_i} = \frac{\partial (\sum_{i=1}^{I-1} v_i \tanh(ha_{J,n} - s_i))}{\partial v_i} = \tanh(ha_{J,n} - s_i)$; equality (b) follows i) because by $\mathbf{z}_1 = \mathbf{a}_{J+1}$, we have $\frac{\partial z_{1,n'}}{\partial a_{J+1,n}} = \frac{\partial a_{J+1,n'}}{\partial a_{J+1,n}} = 1$, if $n = n'$, and otherwise 0, and ii) from substitution $\delta_{1,n} = \frac{\partial C}{\partial z_{1,n}}$ from (23).

The partial derivative of C with respect to shift coefficient s_i , $i = 1, \dots, I - 1$, is derived as

$$\begin{aligned} \frac{\partial C}{\partial s_i} &= \sum_{n=1}^K \sum_{n'=1}^K \frac{\partial C}{\partial z_{1,n'}} \frac{\partial z_{1,n'}}{\partial a_{J+1,n}} \frac{\partial a_{J+1,n}}{\partial s_i} \\ &\stackrel{(a)}{=} \sum_{n=1}^K \sum_{n'=1}^K \frac{\partial C}{\partial z_{1,n'}} \frac{\partial z_{1,n'}}{\partial a_{J+1,n}} v_i \tanh'_{s_i}(ha_{J,n} - s_i) \\ &\stackrel{(b)}{=} \sum_{n=1}^K \delta_{1,n} v_i \tanh'_{s_i}(ha_{J,n} - s_i), \end{aligned} \quad (36)$$

where equality (a) follows from $\frac{\partial a_{J+1,n}}{\partial s_i} = \frac{\partial (\sum_{i=1}^{I-1} v_i \tanh(ha_{J,n} - s_i))}{\partial s_i} = v_i \tanh'_{s_i}(ha_{J,n} - s_i)$, where $\tanh'_{s_i}(\cdot)$ denotes the derivative of $\tanh(ha_{J,n} - s_i)$ with respect to s_i , given as $-4/(\exp\{ha_{J,n} - s_i\} + \exp\{-ha_{J,n} + s_i\})^2$; equality (b) follows similarly as step (b) in (35).

REFERENCES

- [1] M. Leinonen and M. Codreanu, "Quantized compressed sensing via deep neural networks," in *Proc. 6G Wireless Summit*, Levi, Finland, Mar. 2020, pp. 1–5.
- [2] M. Leinonen, M. Codreanu, and G. B. Giannakis, "Compressed sensing with applications in wireless networks," *Found. Trends Signal Process.*, vol. 13, nos. 1–2, pp. 1–282, Nov. 2019. [Online]. Available: <http://dx.doi.org/10.1561/20000000107>
- [3] G. Quer, R. Masiero, G. Pilonetto, M. Rossi, and M. Zorzi, "Sensing, compression, and recovery for WSNs: Sparse signal modeling and monitoring framework," *IEEE Trans. Wireless Commun.*, vol. 11, no. 10, pp. 3447–3461, Oct. 2012.
- [4] D. Malioutov, M. Çetin, and A. S. Willsky, "A sparse signal reconstruction perspective for source localization with sensor arrays," *IEEE Trans. Signal Process.*, vol. 53, no. 8, pp. 3010–3022, Aug. 2005.
- [5] J. A. Bazerque and G. B. Giannakis, "Distributed spectrum sensing for cognitive radio networks by exploiting sparsity," *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1847–1862, Mar. 2010.

- [6] X. Wang, G. Li, and P. K. Varshney, "Detection of sparse signals in sensor networks via locally most powerful tests," *IEEE Signal Process. Lett.*, vol. 25, no. 9, pp. 1418–1422, Sep. 2018.
- [7] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [8] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [9] J. Haupt and R. Nowak, "Signal reconstruction from noisy random projections," *IEEE Trans. Inf. Theory*, vol. 52, no. 9, pp. 4036–4048, Sep. 2006.
- [10] M. Duarte, M. Wakin, D. Baron, and R. Baraniuk, "Universal distributed sensing via random projections," in *Proc. IEEE 5th Int. Symp. Inf. Process. Sens. Netw.*, Nashville, TN, USA, 2006, pp. 177–185.
- [11] P. T. Boufounos and R. G. Baraniuk, "1-bit compressive sensing," in *Proc. Conf. Inf. Sci. Syst.*, Princeton, NJ, USA, Mar. 2008, pp. 16–21.
- [12] X. Cheng, D. Ciunzio, and P. S. Rossi, "Multibit decentralized detection through fusing smart and dumb sensors based on Rao test," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 56, no. 2, pp. 1391–1405, Apr. 2020.
- [13] V. Goyal, A. Fletcher, and S. Rangan, "Compressive sampling and lossy compression," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 48–56, Mar. 2008.
- [14] R. Dobrushin and B. Tsybakov, "Information transmission with additional noise," *IRE Trans. Inf. Theory*, vol. 8, no. 5, pp. 293–304, Sep. 1962.
- [15] T. Berger, *Rate-Distortion Theory: A Mathematical Basis for Data Compression* (Prentice-Hall Series in Information and System Sciences). Englewood Cliffs, NJ, USA: Prentice Hall, 1971.
- [16] J. Sun and V. Goyal, "Optimal quantization of random measurements in compressed sensing," in *Proc. IEEE Int. Symp. Inf. Theory*, Seoul, South Korea, Jun./Jul. 2009, pp. 6–10.
- [17] A. Zymnis, S. Boyd, and E. J. Candès, "Compressed sensing with quantized measurements," *IEEE Signal Process. Lett.*, vol. 17, no. 2, pp. 149–152, Feb. 2010.
- [18] A. Shirazinia, S. Chatterjee, and M. Skoglund, "Joint source-channel vector quantization for compressed sensing," *IEEE Trans. Signal Process.*, vol. 62, no. 14, pp. 3667–3681, Jul. 2014.
- [19] M. Leinonen, M. Codreanu, and M. Juntti, "Distributed distortion-rate optimized compressed sensing in wireless sensor networks," *IEEE Trans. Commun.*, vol. 66, no. 4, pp. 1609–1623, Apr. 2018.
- [20] M. Leinonen, M. Codreanu, M. Juntti, and G. Kramer, "Rate-distortion performance of lossy compressed sensing of sparse sources," *IEEE Trans. Commun.*, vol. 66, no. 10, pp. 4498–4512, Oct. 2018.
- [21] M. Leinonen, M. Codreanu, and M. Juntti, "Practical compression methods for quantized compressed sensing," in *Proc. IEEE INFOCOM Conf. Comput. Commun. Workshop*, Paris, France, Apr./May 2019, pp. 756–761.
- [22] H. M. Shi, M. Case, X. Gu, S. Tu, and D. Needell, "Methods for quantized compressed sensing," in *Proc. Inf. Theory Appl. Workshop*, La Jolla, CA, USA, Jan./Feb. 2016, pp. 1–9.
- [23] A. Mousavi, A. B. Patel, and R. G. Baraniuk, "A deep learning approach to structured signal recovery," in *Proc. Allerton Conf. Commun. Control Comput.*, Monticello, IL, USA, Sep./Oct. 2015, pp. 1336–1343.
- [24] B. Sun, H. Feng, K. Chen, and X. Zhu, "A deep learning framework of quantized compressed sensing for wireless neural recording," *IEEE Access*, vol. 4, pp. 5169–5178, 2016.
- [25] Y. Bengio, N. Léonard, and A. C. Courville, "Estimating or propagating gradients through stochastic neurons for conditional computation," 2013. [Online]. Available: <https://arxiv.org/abs/1308.3432>.
- [26] E. Agustsson, *et al.*, "Soft-to-hard vector quantization for end-to-end learning compressible representations," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, Long Beach, CA, USA, Dec. 2017, pp. 1141–1151.
- [27] Z.-G. Liu and M. Mattina, "Learning low-precision neural networks without straight-through estimator (STE)," in *Proc. 28th Int. Joint Conf. Artif. Intell.*, Macao, China, May 2019, pp. 3066–3072. [Online]. Available: <https://arxiv.org/abs/1903.01061>.
- [28] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016. [Online]. Available: <http://www.deeplearningbook.org>
- [29] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, no. 6088, pp. 533–536, Oct. 1986.
- [30] J. Wolf and J. Ziv, "Transmission of noisy information to a noisy receiver with minimum distortion," *IEEE Trans. Inf. Theory*, vol. 16, no. 4, pp. 406–411, Jul. 1970.
- [31] N. Shlezinger and Y. C. Eldar, "Deep task-based quantization," Aug. 2019. [Online]. Available: <https://arxiv.org/abs/1908.06845>.
- [32] K. Kulkarni, S. Lohit, P. Turaga, R. Kerviche, and A. Ashok, "ReconNet: Non-iterative reconstruction of images from compressively sensed measurements," in *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recogn.*, Las Vegas, NV, USA, Jun./Jul. 2016, pp. 449–458.
- [33] A. Mousavi, G. Dasarathy, and R. G. Baraniuk, "DeepCodec: Adaptive sensing and recovery via deep convolutional neural networks," 2017. [Online]. Available: <https://arxiv.org/abs/1707.03386>.
- [34] A. Mousavi and R. G. Baraniuk, "Learning to invert: Signal recovery via deep convolutional networks," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, New Orleans, LA, USA, Mar. 2017, pp. 2272–2276.
- [35] Z. Zhang, Y. Wu, C. Gan, and Q. Zhu, "The optimally designed autoencoder network for compressed sensing," *EURASIP J. Image Video Process.*, vol. 2019, no. 56, pp. 1–12, Apr. 2019.
- [36] Y. Yang, J. Sun, H. Li, and Z. Xu, "ADMM-CSNet: A deep learning approach for image compressive sensing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 3, pp. 521–538, Mar. 2020.
- [37] B. Sun and H. Feng, "Efficient compressed sensing for wireless neural recording: A deep learning approach," *IEEE Signal Process. Lett.*, vol. 24, no. 6, pp. 863–867, Jun. 2017.
- [38] Z. Wang, Q. Ling, and T. S. Huang, "Learning deep ℓ_0 encoders," in *Proc. AAAI Conf. Artif. Intell.*, Phoenix, AZ, USA, Feb. 2016, pp. 2194–2200.
- [39] A. Bora, A. Jalal, E. Price, and A. G. Dimakis, "Compressed sensing using generative models," in *Proc. Int. Conf. Mach. Learn.*, Sydney, NSW, Australia, Aug. 2017, pp. 537–546.
- [40] H. Palangi, R. Ward, and L. Deng, "Distributed compressed sensing: A deep learning approach," *IEEE Trans. Signal Process.*, vol. 64, no. 17, pp. 4504–4518, Sep. 2016.
- [41] M. Mishali and Y. Eldar, "Sub-Nyquist sampling," *IEEE Signal Process. Mag.*, vol. 28, no. 6, pp. 98–124, Nov. 2011.
- [42] W. Cui, F. Jiang, X. Gao, S. Zhang, and D. Zhao, "An efficient deep quantized compressed sensing coding framework of natural images," in *Proc. 26th ACM Int. Conf. Multimedia*, Seoul, South Korea, Oct. 2018, pp. 1777–1785.
- [43] R. K. Mahabadi, J. Lin, and V. Cevher, "A learning-based framework for quantized compressed sensing," *IEEE Signal Process. Lett.*, vol. 26, no. 6, pp. 883–887, Jun. 2019.
- [44] M. Shohat, G. Tsintsadze, N. Shlezinger, and Y. C. Eldar, "Deep quantization for MIMO channel estimation," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, Brighton, U.K., May 2019, pp. 3912–3916.
- [45] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding," in *Proc. Int. Conf. Learn. Represent.*, San Juan, Puerto Rico, May 2016, pp. 1–14. [Online]. Available: <https://arxiv.org/abs/1510.00149>.
- [46] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, "Quantized neural networks: Training neural networks with low precision weights and activations," *J. Mach. Learn. Res.*, vol. 18, no. 1, pp. 1–30, Jan. 2017. [Online]. Available: <https://arxiv.org/abs/1609.07061>.
- [47] R. Gong *et al.*, "Differentiable soft quantization: Bridging full-precision and low-bit neural networks," in *Proc. Int. Conf. Comp. Vision*, Seoul, South Korea, Oct./Nov. 2019, pp. 4851–4860. [Online]. Available: <https://arxiv.org/abs/1908.05033>.
- [48] H. Bourlard and Y. Kamp, "Auto-association by multilayer perceptrons and singular value decomposition," *Biol. Cybern.*, vol. 59, nos. 4–5, pp. 291–294, Sep. 1988.
- [49] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, "Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion," *J. Mach. Learn. Res.*, vol. 11, pp. 3371–3408, Dec. 2010.
- [50] S. Wu *et al.*, "Learning a compressed sensing measurement matrix via gradient unrolling," in *Proc. Int. Conf. Mach. Learn.*, Long Beach, CA, USA, Jun. 2019, pp. 6828–6939.

- [51] A. Grover and S. Ermon, "Uncertainty autoencoders: Learning compressed representations via variational information maximization," in *Proc. Mach. Learn. Res.*, vol. 89, Apr. 2019, pp. 2514–2524.
- [52] S. Li, W. Zhang, Y. Cui, H. V. Cheng, and W. Yu, "Joint design of measurement matrix and sparse support recovery method via deep auto-encoder," *IEEE Signal Process. Lett.*, vol. 26, no. 12, pp. 1778–1782, Dec. 2019.
- [53] L. Theis, W. Shi, A. Cunningham, and F. Huszár, "Lossy image compression with compressive autoencoders," in *Proc. Int. Conf. Learn. Represent.*, Toulon, France, Apr. 2017, pp. 1–19.
- [54] N. Shlezinger, Y. C. Eldar, and M. R. D. Rodrigues, "Hardware-limited task-based quantization," *IEEE Trans. Signal Process.*, vol. 67, no. 20, pp. 5223–5238, Oct. 2019.
- [55] A. Kipnis, G. Reeves, Y. Eldar, and A. Goldsmith, "Compressed sensing under optimal quantization," in *Proc. IEEE Int. Symp. Inf. Theory*, Aachen, Germany, Jun. 2017, pp. 2148–2152.
- [56] A. Kipnis, G. Reeves, and Y. C. Eldar, "Single letter formulas for quantized compressed sensing with Gaussian codebooks," in *Proc. IEEE Int. Symp. Inf. Theory*, Vail, CO, USA, Jun. 2018, pp. 71–75.
- [57] M. Elad and I. Yavneh, "A plurality of sparse representations is better than the sparsest one alone," *IEEE Trans. Inf. Theory*, vol. 55, no. 10, pp. 4701–4714, Oct. 2009.
- [58] C. Li *et al.*, "Efficient and self-adaptive in-situ learning in multilayer memristor neural networks," *Nat. Commun.*, vol. 9, no. 2385, pp. 1–8, Jun. 2018.
- [59] S. Ambrogio *et al.*, "Equivalent-accuracy accelerated neural-network training using analogue memory," *Nature*, no. 558, pp. 60–67, Jun. 2018.
- [60] P. Yao *et al.*, "Fully hardware-implemented memristor convolutional neural network," *Nature*, no. 577, pp. 641–646, Jan. 2020.
- [61] L. Danial, N. Wainstein, S. Kraus, and S. Kvatinisky, "Breaking through the speed-power-accuracy tradeoff in ADCs using a memristive neuromorphic architecture," *IEEE Trans. Emerg. Topics Comput. Intell.*, vol. 2, no. 5, pp. 396–409, Oct. 2018.
- [62] J. Max, "Quantizing for minimum distortion," *IRE Trans. Inf. Theory*, vol. 6, no. 1, pp. 7–12, Mar. 1960.
- [63] S. Lloyd, "Least squares quantization in PCM," *IEEE Trans. Inf. Theory*, vol. 28, no. 2, pp. 129–137, Mar. 1982.
- [64] K. Rose, E. Gurewitz, and G. C. Fox, "Vector quantization by deterministic annealing," *IEEE Trans. Inf. Theory*, vol. 38, no. 4, pp. 1249–1257, Jul. 1992.
- [65] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. Int. Conf. Artif. Intell. Stat.*, Chia Laguna Resort, Italy, May 2010, pp. 249–256.
- [66] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014. [Online]. Available: <https://arxiv.org/abs/1412.6980>.
- [67] E. J. Candès and J. Romberg, (Oct. 2005). *l_1 -MAGIC: Recovery of Sparse Signals Via Convex Programming*. [Online]. Available: <http://users.ece.gatech.edu/~justin/l1magic/>
- [68] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.
- [69] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing* (Applied and Numerical Harmonic Analysis). New York, NY, USA: Birkhäuser, 2013.
- [70] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. 27th Asilomar Conf. Signals Syst. Comput.*, Pacific Grove, CA, USA, Nov. 1993, pp. 40–44.