

DEEP SIGNAL RECOVERY WITH ONE-BIT QUANTIZATION

Shahin Khobahi^{*†}, Naveed Naimipour[†], Mojtaba Soltanalian[†] and Yonina C. Eldar⁺

[†]Department of Electrical and Computer Engineering, University of Illinois at Chicago, Chicago, USA

⁺Department of Electrical Engineering, Technion, Israel Institute of Technology, Haifa, Israel

ABSTRACT

Machine learning, and more specifically deep learning, have shown remarkable performance in sensing, communications, and inference. In this paper, we consider the application of the deep unfolding technique in the problem of signal reconstruction from its one-bit noisy measurements. Namely, we propose a model-based machine learning method and unfold the iterations of an inference optimization algorithm into the layers of a deep neural network for one-bit signal recovery. The resulting network, which we refer to as *DeepRec*, can efficiently handle the recovery of high-dimensional signals from acquired one-bit noisy measurements. The proposed method results in an improvement in accuracy and computational efficiency with respect to the original framework as shown through numerical analysis.

Index Terms— Deep learning, deep unfolding, MIMO communications, big data, machine learning, neural network, maximum likelihood, one-bit quantization

1. INTRODUCTION

Quantization of signals of interest is an integral part of all modern digital signal processing applications such as sensing, communication, and inference. In an ideal hardware implementation of a quantization system, a high-resolution analog-to-digital converter (ADC) with b -bit resolution and sampling frequency of f_s samples the original analog signal and maps the obtained samples into a discrete state space of size $2^b f_s$. Generally, a large number of bits is required to obtain an accurate digital representation of the analog signal. In such a case, the quantization process has negligible impact on the performance of algorithms which were typically developed on the assumptions of infinite precision samples, and thus, the high-resolution (in terms of amplitude) quantization process can be directly modeled as an additive noise source. However, a crucial obstacle with modern ADCs is that their power consumption, manufacturing cost, and chip area grows exponentially with their resolution b [1–3].

The required high sampling data rate of ADCs used in next generation communications systems is another obstacle that must be tackled in such applications. For instance, the promising millimeter wave (mmWave) multiple-input multiple output (MIMO) communication technology requires a very large bandwidth, and the corresponding sampling rate of the ADCs must increase accordingly. However, manufacturing ADCs with high-resolution (e.g., more than 8 bits) and high sampling rate are extremely costly and may not be available. Moreover, in other applications such as spectral sensing and cognitive radio, which require extremely high sampling rates, the cumulative cost and power consumption of using high-resolution and extremely fast ADCs are typically prohibitive and impractical. Hence, when signals across a wide frequency band are of interest, a fundamental trade-off between sampling rate, amplitude quantization precision, cost, and power consumption is encountered. An immediate solution to such challenges is to use low-resolution, and specifically *one-bit*, ADCs. The use of one-bit signed measurements, and more specifically one-bit ADCs, allows for an extremely high sampling rate at a low cost and low power consumption. From a sampling viewpoint, the most extreme case of quantization is to use only one bit per sample. More precisely, one-bit sampling can be seen as a process through which we repeatedly compare the amplitude of a signal (at each sample) to some reference threshold level and use only one bit to convey whether the signal amplitude resides above or below that threshold. Due to its appealing sampling properties, the problem of recovering a signal from its one-bit measurements has attracted a great deal of interest over the past few years [4–8]. Therefore, it is vital to develop algorithms that can deal with low-resolution samples for different applications.

The fields of machine learning (ML), and more particularly deep learning, are impacting various fields of engineering and have recently attracted a great deal of attention in tackling long-standing signal processing problems. The advent of low-cost specialized powerful computing resources (e.g., GPUs, and more recently TPUs) and the continually increasing amount of massive data generated by the human population and machines, in conjunction with the new optimization and learning methods, have paved the way for deep neural networks (DNNs) and machine learning-based models

* Corresponding author (e-mail: skhoba2@uic.edu). This work was supported in part by U.S. National Science Foundation Grants CCF-1704401 and ECCS-1809225.

to prove their effectiveness in many engineering areas (see, e.g., [9–11] and the references therein).

The main advantage of the deep learning-based model herein is that it employs several non-linear transformations to obtain an abstract representation of the underlying data. Model-based machine learning frameworks (e.g., probabilistic graphical models) incorporate prior knowledge of the system parameters into the inference process. A recent promising approach in bridging the gap between deep learning-based and model-based methods is the paradigm of *deep unfolding* [12]. Particularly, iterations of a conventional recursive algorithm, such as fast iterative soft thresholding algorithm (FISTA), projected gradient descent, and approximate message passing (AMP), can be used as a baseline to design the architecture of a deep network with trainable parameters specifically customized to the problem of interest. Such a methodology results in an improvement in accuracy, and computational efficiency of the original framework. The deep unfolding method has already shown remarkable performance improvement in a wide range of applications such as MIMO communications [13, 14], multi-channel source separation [15], and sparse inverse problems [16, 17].

In this paper, we consider the general problem of high-dimensional signal recovery from random one-bit measurements. Specifically, we propose an efficient signal recovery framework based on the deep unfolding technique that has the advantage of low-complexity and near-optimal performance compared to traditional methods. Our proposed inference framework has a wide range of applications in the areas of wireless communications, detection and estimation, and sensing.

2. PROBLEM FORMULATION

We begin by considering a general linear signal acquisition and one-bit quantization model with time-varying thresholds, described as follows:

$$\text{Signal Model:} \quad \mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}, \quad (1)$$

$$\text{Quantization Model:} \quad \mathbf{r} \triangleq \text{sign}(\mathbf{y} - \boldsymbol{\tau}), \quad (2)$$

where $\boldsymbol{\tau} = [\tau_1, \dots, \tau_M]^T$ denotes the vector of one-bit quantization thresholds, $\mathbf{y} \in \mathbb{R}^M$ denotes the received signal prior to quantization, $\mathbf{H} \in \mathbb{R}^{M \times N}$ denotes the sensing matrix, $\mathbf{x} \in \mathbb{R}^N$ denotes the multidimensional unknown vector to be recovered, and $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$ denotes the zero-mean Gaussian noise with a known covariance matrix $\mathbf{C} = \text{Diag}(\sigma_1^2, \dots, \sigma_M^2)$. Furthermore, $\text{sign}(\cdot)$ denotes the signum function applied element-wise on the vector argument.

The above model covers a wide range of applications. For instance, the described model (1)-(2) can be used in MIMO communication systems in which \mathbf{H} is the channel matrix, \mathbf{x} is the signal sent by the transmitter, \mathbf{n} is the additive Gaussian noise in the system, and the base station is equipped with

one-bit ADCs, where the goal is to recover the transmitted symbols from \mathbf{r} .

2.1. Maximum Likelihood Estimator Derivation

Given the knowledge of the sensing matrix \mathbf{H} , noise covariance \mathbf{C} , and the corresponding quantization thresholds $\boldsymbol{\tau}$, our goal is to recover the original (likely high-dimensional) signal \mathbf{x} from the one-bit random measurements \mathbf{r} . In such a scenario, each binary observation $\{r_i\}_{i=1}^N$ follows a Bernoulli distribution with parameter p_i , given by:

$$p_i = \text{Prob}\{\mathbf{h}_i^T \mathbf{x} + n_i - \tau_i > 0\} = Q\left(\frac{\tau_i - \mathbf{h}_i^T \mathbf{x}}{\sigma_i}\right), \quad (3)$$

where $Q(x) = 1 - \Phi(x)$ with $\Phi(x)$ representing the cumulative distribution function (cdf) of a standard Gaussian distribution and \mathbf{h}_i^T denotes the i -th row of the matrix \mathbf{H} . In particular, the probability mass function (pmf) of each binary observation can be compactly expressed as:

$$p(r_i) = Q\left(\frac{r_i}{\sigma_i} (\tau_i - \mathbf{h}_i^T \mathbf{x})\right), \quad (4)$$

where $r_i \in \{-1, +1\}$. Therefore, the log-likelihood of the quantized observations \mathbf{r} given the unknown vector \mathbf{x} can be expressed as:

$$\mathcal{L}(\mathbf{x}) = p(\mathbf{r}|\mathbf{x}) = \log \left\{ \prod_{i=1}^N Q\left(\frac{r_i}{\sigma_i} (\tau_i - \mathbf{h}_i^T \mathbf{x})\right) \right\} \quad (5)$$

$$= \sum_{i=1}^N \log \left\{ Q\left(\frac{r_i}{\sigma_i} (\tau_i - \mathbf{h}_i^T \mathbf{x})\right) \right\}, \quad (6)$$

where $\log\{\cdot\}$ denotes the natural logarithm. As a result, the maximum likelihood (ML) estimation of \mathbf{x} can be obtained as

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\text{argmax}} \mathcal{L}(\mathbf{x}). \quad (7)$$

Observe that the maximum likelihood estimator $\hat{\mathbf{x}}$ has to satisfy the following condition:

$$\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}) = \mathbf{0}, \quad (8)$$

where the gradient of the log-likelihood function with respect to the unknown vector \mathbf{x} can be derived as follows:

$$\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}) = \sum_{i=1}^N \left[-\frac{r_i}{\sigma_i} \left(\frac{Q'\left(\frac{r_i}{\sigma_i} (\tau_i - \mathbf{h}_i^T \mathbf{x})\right)}{Q\left(\frac{r_i}{\sigma_i} (\tau_i - \mathbf{h}_i^T \mathbf{x})\right)} \right) \right] \mathbf{h}_i, \quad (9)$$

where $Q'(x) = -\frac{1}{\sqrt{2\pi}} \exp(-x^2/2)$. It can be observed from (9) that the gradient of the log-likelihood function is a linear combination of the rows of the sensing matrix \mathbf{H} . Let $\boldsymbol{\eta}: \mathbb{R}^M \mapsto \mathbb{R}^M$ be a non-linear function defined as follows:

$$\boldsymbol{\eta}(\mathbf{x}) \triangleq \frac{Q'(\mathbf{x})}{Q(\mathbf{x})}, \quad (10)$$

where the functions $Q(\cdot)$, $Q'(\cdot)$, and the division, are applied element-wise on the vector argument \mathbf{x} . In addition, let $\mathbf{\Omega} = \mathbf{Diag}(r_1, \dots, r_M)$ be a diagonal matrix containing the one-bit observations and $\tilde{\mathbf{\Omega}} = \mathbf{\Omega} \mathbf{C}^{-\frac{1}{2}}$ be the semi-whitened version of the *one-bit matrix* $\mathbf{\Omega}$. Then, the gradient of the likelihood function in (9) can be compactly written as follows:

$$\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}) = -\mathbf{H}^T \tilde{\mathbf{\Omega}} \eta \left(\tilde{\mathbf{\Omega}}(\boldsymbol{\tau} - \mathbf{H}\mathbf{x}) \right). \quad (11)$$

Recall that the ML estimator $\hat{\mathbf{x}}$ must satisfy the condition given in (8), i.e.,

$$\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}) = -\mathbf{H}^T \tilde{\mathbf{\Omega}} \eta \left(\tilde{\mathbf{\Omega}}(\boldsymbol{\tau} - \mathbf{H}\mathbf{x}) \right) = \mathbf{0}. \quad (12)$$

Other than certain low-dimensional cases, finding a closed-form expression for $\hat{\mathbf{x}}$ that satisfies (12) is a difficult task [18–20]. Therefore, we resort to iterative methods in order to find the ML estimate, i.e., to solve (7).

In this paper, the well-known gradient ascent method is employed to iteratively solve (7). Namely, given an initial point $\mathbf{x}^{(0)}$, the update equation at each iteration is given by:

$$\begin{aligned} \mathbf{x}^{(k+1)} &= \mathbf{x}^{(k)} + \delta^{(k)} \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}) \\ &= \mathbf{x}^{(k)} - \delta^{(k)} \mathbf{H}^T \tilde{\mathbf{\Omega}} \eta \left(\tilde{\mathbf{\Omega}}(\boldsymbol{\tau} - \mathbf{H}\mathbf{x}^{(k)}) \right), \end{aligned} \quad (13)$$

where $\delta^{(k)}$ is the step size at the k -th iteration. The obtained maximum likelihood estimator derived from the signal model, and the corresponding optimization steps, can be unfolded into a multi-layer deep neural network, which improves the accuracy and computational efficiency of the original framework.

In the next section, we *unfold* the above iterations into the layers of a deep neural network where each layer denotes one iteration of the above optimization method. Interestingly, we fix the complexity budget of the inference framework (via fixing the number of layers), and apply the gradient descent method to yield the most accurate estimation of the parameter in at most K iterations.

3. SIGNAL RECOVERY VIA DEEP UNFOLDING

Conventionally, first-order optimization methods, such as gradient descent algorithms, have slow convergence rate, and thus take a large number of iterations to converge to a solution. Herein, we are interested in finding a good solution under the condition that the complexity of the inference algorithm is fixed. This is important since, via unfolding the optimization algorithm, we fix the computational complexity of the inference model (a DNN with K layers in such a case) and optimize the parameters of the network to find the best possible estimator with a fixed-complexity constraint. Below, we introduce *DeepRec*, our deep learning based signal recovery framework which is designed based on the iterations of

the form (14), to find the maximum likelihood estimation of the unknown parameter.

—**The DeepRec Architecture.** The construction of DeepRec involves the unfolding of $k = 1, \dots, K$, iterations each of which are of the form (14), as the layers of a deep neural network. Particularly, each step of the gradient descent method depends on the previous signal estimate $\mathbf{x}^{(k)}$, the step size $\delta^{(k)}$, the scaled one-bit matrix $\tilde{\mathbf{\Omega}}$, the sensing matrix \mathbf{H} , and the threshold vector $\boldsymbol{\tau}$. In addition, the form of the gradient vector (11) makes it convenient and insightful to unfold the iterations onto the layers of a DNN in that each iteration of the gradient descent method is a linear combination of the system parameters followed by a non-linearity. The k -th layer of DeepRec can be characterized via the following operations and variables:

$$\mathbf{z}^{(k)} = \mathbf{W}_{1k} \tilde{\mathbf{\Omega}} \boldsymbol{\tau} - \mathbf{W}_{2k} \tilde{\mathbf{\Omega}} \mathbf{H} \mathbf{x}^{(k)} + \mathbf{b}_{1k}, \quad (15)$$

$$\mathbf{p}^{(k)} = \eta \left(\mathbf{z}^{(k)} \right), \quad (16)$$

$$\mathbf{t}^{(k)} = \mathbf{H}^T \tilde{\mathbf{\Omega}} \mathbf{p}^{(k)}, \quad (17)$$

$$\mathbf{x}^{(k+1)} = f \left(\mathbf{W}_{3k} \begin{bmatrix} \mathbf{x}^{(k)} \\ \mathbf{t}^{(k)} \end{bmatrix} + \mathbf{b}_{2k} \right), \quad (18)$$

where $\mathbf{x}^{(1)} = \mathbf{0}$, $f(\cdot)$ denotes a non-linear activation function where in this work we consider $f(x) \triangleq \text{ReLU}(x) = \max\{0, x\}$, and the goal is to optimize the DNN parameters, described as follows:

$$\Xi = \{\mathbf{W}_{1k}, \mathbf{W}_{2k}, \mathbf{W}_{3k}, \mathbf{b}_{1k}, \mathbf{b}_{2k}\}_{k=1}^K. \quad (19)$$

The proposed DeepRec architecture with L layers can be interpreted as a class of estimator functions $\Psi_{\Xi}(\mathbf{r}, \mathbf{H}, \boldsymbol{\tau})$ parametrized by Ξ to estimate the unknown parameter \mathbf{x} given the system parameters. In order to find the best estimator function $\Psi_{\Xi}(\mathbf{r}, \mathbf{H}, \boldsymbol{\tau})$ associated with our problem, we conduct a learning process via minimizing a loss function $\mathcal{R}(\mathbf{x}; \Psi_{\Xi}(\mathbf{r}, \mathbf{H}, \boldsymbol{\tau}))$, i.e.,

$$\min_{\Xi} \mathcal{R}(\mathbf{x}; \Psi_{\Xi}(\mathbf{r}, \mathbf{H}, \boldsymbol{\tau})). \quad (20)$$

In this work, we employ the following least squares (LS) loss function:

$$\mathcal{R}(\mathbf{x}; \Psi_{\Xi}(\mathbf{r}, \mathbf{H}, \boldsymbol{\tau})) = \|\mathbf{x} - \Psi_{\Xi}(\mathbf{r}, \mathbf{H}, \boldsymbol{\tau})\|_2^2, \quad (21)$$

where during the training phase, we synthetically generate the system parameters $\Theta = \{\mathbf{x}, \mathbf{r}, \mathbf{H}, \boldsymbol{\tau}\}$ according to their statistical model.

4. NUMERICAL RESULTS

We now demonstrate the performance of the proposed DeepRec framework for the problem of one-bit signal recovery. The proposed framework was implemented using the TensorFlow library [21], with the ADAM stochastic optimizer [22]

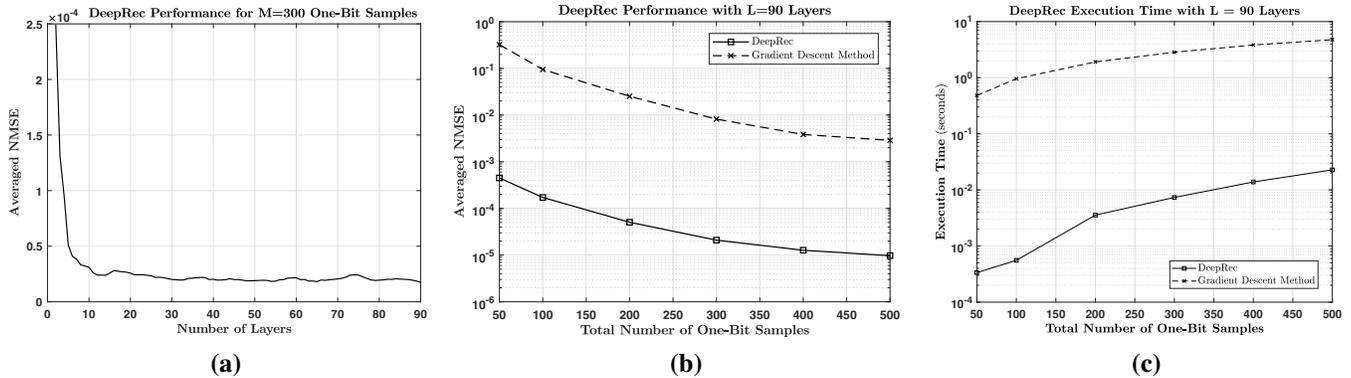


Fig. 1. The performance of DeepRec: **(a)** demonstrates the NMSE performance of the DeepRec network for different numbers of layers K . **(b)** shows the performance of the proposed DeepRec architecture and the original gradient descent method of (14) in terms of averaged NMSE for different numbers of one-bit samples M . **(c)** shows a comparison of the computational cost between the gradient descent method and the proposed DeepRec network for different numbers of one-bit samples M .

and an exponential decaying step size. In the learning process of the network, we employed the batch training method with a batch size of 500 at each epoch and we performed the training for 2000 epochs. In all of the simulations, we assumed $N = 3$, i.e., $\mathbf{x} \in \mathbb{R}^3$, and we used the normalized mean square error (NMSE) defined as $NMSE = \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 / \|\mathbf{x}\|_2^2$, for the performance metric.

The training was performed based on the data generated via the following model. Each element of the vector \mathbf{x} is assumed to be i.i.d and uniformly distributed, i.e., $x \sim \mathcal{U}(\delta_l^x, \delta_u^x)$. The sensing matrix is assumed to be fixed and follow a Normal distribution, where we consider $\mathbf{H} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The quantization thresholds were also generated from a uniform distribution, $\tau \sim \mathcal{U}(\delta_l^\tau, \delta_u^\tau)$, where the lower and upper bound of the distribution is chosen in a fashion that at least covers the domain of \mathbf{x} . The noise is assumed to be independent from one sample to another and follows a Normal distribution, where the variance of each corresponding noise element is different, e.g., the noise covariance $\mathbf{C} = \mathbf{Diag}(\sigma_1^2, \dots, \sigma_M^2)$, with $\sigma_i^2 \sim \mathcal{U}(\delta_1^n, \delta_M^n)$. Note that we trained the network over a wide range of noise powers in order to make the DeepRec network more robust to noise.

Fig. 1(a) demonstrates the performance of the DeepRec network for different numbers of layers K . It can be observed that the averaged NMSE decreases dramatically as the number of layers increases. Such a result is also expected as each layer corresponds to one iteration of original optimization algorithm. Thus, as the number of layers increases, the output of the network will converge to a better estimation as well.

Fig. 1(b) demonstrates the performance of the proposed DeepRec architecture and the original Gradient Descent method of (14) in terms of averaged NMSE for different numbers of one-bit samples M . In this simulation, we implemented the DeepRec network with $K = 90$ layers. It can be

clearly seen from Fig. 1(b) that the proposed deep recovery architecture (DeepRec) significantly outperforms the original optimization method in terms of accuracy and provides a considerably better estimation than that of the gradient descent method for the same number of iterations/layers. As a fair comparison, we also assumed a fixed-step size of $\delta = 0.01$ for the gradient descent method.

Fig. 1(c) shows a comparison of the computational cost (machine runtime) between the gradient descent method and the proposed DeepRec network for different numbers of one-bit samples M . It can be seen that our proposed method (DeepRec) has a significantly lower computational cost than that of the original optimization algorithm for our problem. Hence, making the DeepRec a good candidate for real-time signal processing or big data applications (the results were obtained on a standard PC with a quad-core 2.30GHz CPU and 4 GB memory).

5. CONCLUSION

We have considered the application of model-based machine learning, and specifically the deep unfolding technique, in the problem of recovering a high-dimensional signal from its one-bit quantized noisy measurements via random thresholding. We proposed a novel deep architecture, which we refer to as *DeepRec*, that was able to accurately perform the task of one-bit signal recovery. Our numerical results show that the proposed DeepRec network significantly improves the performance of traditional optimization methods both in terms of accuracy and efficiency.

6. REFERENCES

- [1] Bin Le, Thomas W Rondeau, Jeffrey H Reed, and Charles W Bostian, "Analog-to-digital converters,"

- IEEE Signal Processing Magazine*, vol. 22, no. 6, pp. 69–77, 2005.
- [2] Nir Shlezinger, Yonina C Eldar, and Miguel RD Rodrigues, “Hardware-limited task-based quantization,” *arXiv preprint arXiv:1807.08305*, 2018.
- [3] Alon Kipnis, Yonina C Eldar, and Andrea J Goldsmith, “Fundamental distortion limits of analog-to-digital compression,” *IEEE Transactions on Information Theory*, vol. 64, no. 9, pp. 6013–6033, 2018.
- [4] Christopher Gianelli, Luzhou Xu, Jian Li, and Petre Stoica, “One-bit compressive sampling with time-varying thresholds: Maximum likelihood and the cramer-rao bound,” in *2016 50th Asilomar Conference on Signals, Systems and Computers*. IEEE, 2016, pp. 399–403.
- [5] Naveed Naimipour and Mojtaba Soltanalian, “Graph clustering using one-bit comparison data,” in *2018 IEEE Asilomar Conference on Signals, Systems and Computers*. IEEE, 2018.
- [6] Yongzhi Li, Cheng Tao, Gonzalo Seco-Granados, Amine Mezghani, A Lee Swindlehurst, and Liu Liu, “Channel estimation and performance analysis of one-bit massive MIMO systems,” *IEEE Trans. Signal Process.*, vol. 65, no. 15, pp. 4075–4089, 2017.
- [7] Shahin Khobahi and Mojtaba Soltanalian, “Signal recovery from 1-bit quantized noisy samples via adaptive thresholding,” in *2018 IEEE Asilomar Conference on Signals, Systems and Computers*. IEEE, 2018.
- [8] Fangqing Liu, Heng Zhu, Jian Li, Pu Wang, and Philip V Orlik, “Massive MIMO channel estimation using signed measurements with antenna-varying thresholds,” in *2018 IEEE Statistical Signal Processing Workshop (SSP)*. IEEE, 2018, pp. 188–192.
- [9] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436, 2015.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.
- [11] Li Deng, Dong Yu, et al., “Deep learning: methods and applications,” *Foundations and Trends® in Signal Processing*, vol. 7, no. 3–4, pp. 197–387, 2014.
- [12] John R Hershey, Jonathan Le Roux, and Felix Weninger, “Deep unfolding: Model-based inspiration of novel deep architectures,” *arXiv preprint arXiv:1409.2574*, 2014.
- [13] Hengtao He, Chao-Kai Wen, Shi Jin, and Geoffrey Ye Li, “A model-driven deep learning network for MIMO detection,” *arXiv preprint arXiv:1809.09336*, 2018.
- [14] Neev Samuel, Tzvi Diskin, and Ami Wiesel, “Learning to detect,” *arXiv preprint arXiv:1805.07631*, 2018.
- [15] Scott Wisdom, John Hershey, Jonathan Le Roux, and Shinji Watanabe, “Deep unfolding for multichannel source separation,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 121–125.
- [16] Karol Gregor and Yann LeCun, “Learning fast approximations of sparse coding,” in *Proceedings of the 27th International Conference on Machine Learning*. Omnipress, 2010, pp. 399–406.
- [17] Mark Borgerding and Philip Schniter, “Onsager-corrected deep learning for sparse linear inverse problems,” in *2016 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*. IEEE, 2016, pp. 227–231.
- [18] Michel T Ivrlac and Josef A Nossek, “On MIMO channel estimation with single-bit signal-quantization,” in *ITG Smart Antenna Workshop*, 2007.
- [19] Amine Mezghani, Felix Antreich, and Josef A Nossek, “Multiple parameter estimation with quantized channel output,” in *2010 International ITG Workshop on Smart Antennas (WSA)*. IEEE, 2010, pp. 143–150.
- [20] Jianhua Mo, Philip Schniter, Nuria González Prelcic, and Robert W Heath, “Channel estimation in millimeter wave MIMO systems with one-bit quantization,” in *2014 48th Asilomar Conference on Signals, Systems and Computers*. IEEE, 2014, pp. 957–961.
- [21] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al., “Tensorflow: a system for large-scale machine learning,” in *OSDI*, 2016, vol. 16, pp. 265–283.
- [22] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.