Quantized Proximal Averaging Networks for Compressed Image Recovery (Supplementary Material)

Nareddy Kartheek Kumar Reddy, Mani Madhoolika Bulusu, Praveen Kumar Pokala,
Chandra Sekhar Seelamantula
Department of Electrical Engineering, Indian Institute of Science, Bengaluru, 560012.
{nareddyreddy, manibulusu}@iisc.ac.in, praveenkumar.pokala@gmail.com, css@iisc.ac.in

We present additional experimental validation to demonstrate the efficacy of the quantized models on natural image reconstruction and MR image reconstruction from compressed measurements. As mentioned in the main paper, we consider objective metrics, namely, peak signal-to-noise ratio (PSNR) and structural similarity index metric (SSIM) [**?**] for performance quantification and comparison with full-precision models.

We integrate the proposed quantization framework into existing architectures, namely, ISTA-Net and ISTA-Net$^+$ to develop their quantized counterparts. Further, we compare existing architectures and their quantized counterparts with the proposed architectures (the proximal averaging network (PAN), its enhanced version, PAN$^+$) and their quantized counterparts Q-PAN and Q-PAN$^+$, respectively.

1. Experimental Results and Discussion

In the main paper, we proposed using convex combinations of convex and non-convex regularizers (the ℓ_1 norm, minimax concave penalty (MCP) [**?**] and smoothly clipped absolute deviation (SCAD) [**?**] penalty) to promote sparsity of the learnable analysis prior. We experiment with two combinations of regularizers, namely, the (2R) and (3R) models, on the networks under consideration to understand their effect on reconstruction from compressed measurements. The (2R) models use two regularizers, namely, the ℓ_1 norm and MCP with $\alpha_1 = \alpha_2 = \frac{1}{2}$ and the (3R) models employ three regularizers, namely, the ℓ_1 norm, MCP, and SCAD penalty with $\alpha_1 = \alpha_2 = \alpha_3 = \frac{1}{3}$.

In this document, we present a comparative study of the (2R, 3R) variants of PAN, PAN$^+$ with respective K-bit Q-PAN, Q-PAN$^+$, where $K = 1, 2, 3$, on standard datasets – Set11 [**?**], BSD68 [**?**] and brain MR images [**?**].

1.1. Natural Image Reconstruction from Compressed Measurements

We test the effectiveness of proposed techniques and their corresponding quantized versions by reconstructing natural images in the Set11 and BSD68 datasets. The training data consists of 88,912 cropped grayscale image patches, each of size 33×33 and their linearly compressed measurement for given compressive sensing (CS) ratio. We considered 1%, 4%, and 10% compression ratios in our experiments and trained the models for each CS ratio. The training and inference are carried out on an Ubuntu PC with dual Intel® Xeon® Silver 4110 processors and RTX2080Ti GPU. The models are trained for 100 epochs and take approximately nine hours.

The results corresponding to ISTA-Net, ISTA-Net$^+$ models are shown in Tables 1 and 2, PAN models in Tables 3 and 5, and PAN$^+$ models in Tables 4 and 6. From Table 1, we observe that 1-bit Q-ISTA-Net results in a slightly lower reconstruction performance compared with the full-precision ISTA-Net for 10% measurements on Set11 dataset. Nevertheless, the 2-bit variant shows an improvement of 0.4 dB, and the 3-bit variant shows an improvement of 1.17 dB in PSNR over the 1-bit variant Q-ISTA-Net. On the contrary, for the 1% CS ratio considering the Set11 dataset, 1-bit Q-ISTA-Net has shown 0.09 dB PSNR improvement over ISTA-Net. While one would expect the 1-bit quantized model to perform poorly compared to the full-precision model, it is not uncommon that a quantized model outperforms the full precision counterpart. For instance, Zhang *et al.* [**?**] showed that 2-bit weight quantized Alex-Net [**?**] model achieved 3.4% increase in top-1 and 2.5% increase in top-5 accuracy over the full precision Alex-Net on ImageNet [**?**] dataset classification problem. Similar behavior of quantized models can be observed on the Set11 dataset. We hypothesize that quantization acts like a regularizer on the weights, and thus avoids overfitting the network to training data resulting in better generalization capabilities. We shall present results to

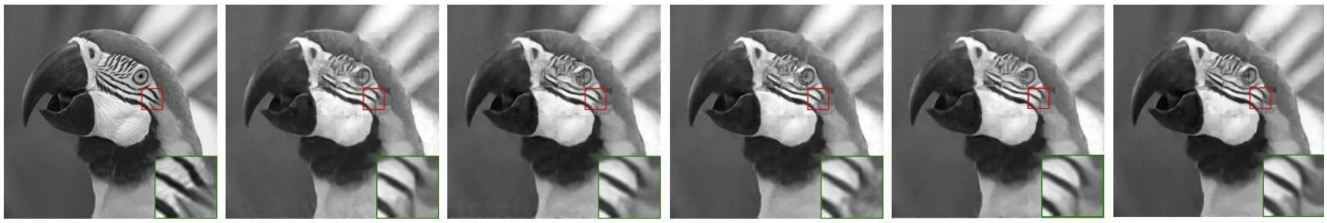

Ground Truth | ISTA-Net, 26.21 dB | PAN (2R), 26.23 dB | ISTA-Net$^+$, 26.37 dB | PAN$^+$ (2R), 26.44 dB | PAN$^+$ (3R), 26.89 dB

Figure 1. Reconstruction results for the *Parrots* image based on 10% compressed measurements. The PSNR values are indicated. It can be observed from the zoomed-in part that PAN$^+$ (3R) offers superior reconstruction of the stripes in comparison with the remaining models.

Table 1. Comparison of quantized variant of ISTA-Net (Q-ISTA-Net) with the full-precision ISTA-Net in terms of the PSNR and SSIM.

CS Ratio	Test Dataset	ISTA-Net (PSNR, SSIM)	1-bit Q-ISTA-Net (PSNR, SSIM)	2-bit Q-ISTA-Net (PSNR, SSIM)	3-bit Q-ISTA-Net (PSNR, SSIM)
10%	Set11	25.80, 0.7961	23.85, 0.6896	24.25, 0.7105	25.02, 0.7460
10%	BSD68	25.02, 0.6953	23.87, 0.6263	24.07, 0.6369	24.75, 0.6729
4%	Set11	21.23, 0.5961	20.43, 0.5520	20.72, 0.5674	20.99, 0.5868
4%	BSD68	22.12, 0.5425	21.61, 0.5191	21.84, 0.5287	21.90, 0.5403
1%	Set11	17.30, 0.4082	17.39, 0.3985	17.40, 0.3538	17.39, 0.4015
1%	BSD68	19.10, 0.4097	19.07, 0.3999	19.01, 0.3466	19.10, 0.4051

show that quantized models outperform full-precision models in the case of ISTA-Net$^+$, PAN$^+$ (2R), and PAN$^+$ (3R) models as well, considering the Set11, BSD68 and brain MRI datasets.

Table 2. Comparison of average PSNR [dB] and SSIM between ISTA-Net$^+$ and $1, 2, 3$ - bit Q-ISTA-Net$^+$. One can observe that there is a degradation of 0.74 dB in PSNR and 0.0364 in SSIM with 1-bit Q-ISTA-Net$^+$ for 10% CS ratio on BSD68. For the same CS ratio, the 2-bit and 3-bit variants show an improvement of 0.43 dB and 0.65 dB, respectively, over the 1-bit variant. For 1% CS ratio on Set11 dataset, 1-bit Q-ISTA-Net$^+$ offers 0.14 dB increase in PSNR over ISTA-Net$^+$.

CS Ratio	Test Dataset	ISTA-Net$^+$ (PSNR, SSIM)	1-bit Q-ISTA-Net$^+$ (PSNR, SSIM)	2-bit Q-ISTA-Net$^+$ (PSNR, SSIM)	3-bit Q-ISTA-Net$^+$ (PSNR, SSIM)
10%	Set11	26.64, 0.8036	24.93, 0.7357	24.37, 0.7282	26.31, 0.7958
10%	BSD68	25.33, 0.6996	24.59, 0.6632	25.02, 0.6857	25.24, 0.6744
4%	Set11	21.31, 0.6240	20.86, 0.5721	20.92, 0.5811	21.51, 0.6125
4%	BSD68	22.17, 0.5569	21.90, 0.5296	21.95, 0.5337	22.34, 0.5531
1%	Set11	17.34, 0.4131	17.48, 0.4055	17.47, 0.4054	17.48, 0.4110
1%	BSD68	19.17, 0.4198	19.19, 0.4169	19.17, 0.4177	19.20, 0.4192

From the experimental results in Tables. 1, 2, 3, 4, 5, 6, we observe that the PAN$^+$ (2R), PAN$^+$ (3R), and ISTA-Net$^+$ are more robust to quantization compared to the PAN (2R), PAN (3R) and ISTA-Net models. We attribute the robustness to skip connections introduced in the enhanced plus models, which enable easy flow of gradients during backpropagation. Moreover, the 1-bit quantized variants perform better or on par with the full-precision models for 1% compression ratio. Among the quantized models, the 3-bit model performs better than the 1-bit and 2-bit models. Further, 3-bit Q-PAN$^+$ and Q-ISTA-Net$^+$ models have shown performance on par with the full-precision models over all the compression ratios under consideration.

We compare the performance of our proposed model PAN$^+$ (3R) and its 1-bit quantized version against ISTA-Net, ISTA-

Table 3. Comparison of average PSNR [dB] and SSIM with PAN (3R) and its quantized variants $1, 2, 3$ - bit Q-PAN (3R). Observe that the 1-bit Q-PAN (3R) has a degradation of 1.32 dB in PSNR and 0.0737 in SSIM for 4% CS ratio measurements in Set11 data. For 1% CS ratio on Set11 dataset, 3-bit Q-PAN (3R) is on par with PAN (3R).

CS RATIO	TEST DATASET	PAN (3R) (PSNR, SSIM)	1-BIT Q-PAN (3R) (PSNR, SSIM)	2-BIT Q-PAN (3R) (PSNR, SSIM)	3-BIT Q-PAN (3R) (PSNR, SSIM)
10%	SET11	25.91, 0.7848	22.92, 0.6364	23.52, 0.6829	24.35, 0.7117
10%	BSD68	25.13, 0.6852	23.09, 0.5718	23.54, 0.6018	24.09, 0.6426
4%	SET11	21.28, 0.5879	19.96, 0.5142	18.32, 0.3501	20.42, 0.5486
4%	BSD68	22.05, 0.5342	20.95, 0.4614	19.04, 0.3065	21.49, 0.4974
1%	SET11	17.46, 0.4063	15.89, 0.1627	16.39, 0.2119	17.42, 0.3957
1%	BSD68	19.22, 0.4105	17.23, 0.1688	17.85, 0.2108	19.13, 0.3956

Table 4. Comparison of average PSNR [dB] and SSIM with PAN$^+$ (3R) and its quantized variants $1, 2, 3$ - bit Q-PAN$^+$ (3R). We observe that the 1-bit Q-PAN$^+$ (3R) has degradation of 0.79 dB in PSNR and 0.0383 in SSIM, for 10% compression ratio measurements in BSD68 data. 1-bit Q-PAN$^+$ (3R) performs on par with PAN$^+$ (3R) for 1% CS ratio on Set11 dataset.

CS RATIO	TEST DATASET	PAN$^+$ (3R) (PSNR, SSIM)	1-BIT Q-PAN$^+$ (3R) (PSNR, SSIM)	2-BIT Q-PAN$^+$ (3R) (PSNR, SSIM)	3-BIT Q-PAN$^+$ (3R) (PSNR, SSIM)
10%	SET11	26.90, 0.8164	25.15, 0.7481	25.84, 0.7794	26.38, 0.7997
10%	BSD68	25.48, 0.7073	24.69, 0.6690	25.06, 0.6877	25.27, 0.6957
4%	SET11	21.83, 0.6359	20.95, 0.5715	21.20, 0.5940	21.47, 0.6106
4%	BSD68	22.49, 0.5639	21.96, 0.5351	22.11, 0.5445	22.28, 0.5537
1%	SET11	17.43, 0.4126	17.46, 0.4064	17.48, 0.4079	17.49, 0.4112
1%	BSD68	19.21, 0.4195	19.15, 0.4172	19.18, 0.4172	19.21, 0.4191

Table 5. Comparison of average PSNR [dB], SSIM with PAN (2R) and $1, 2, 3$ - bit Q-PAN (2R). 1-bit Q-PAN (2R) leads to a degradation of 0.73 dB in PSNR and 0.0457 in SSIM for 4% compression ratio measurements on Set11 data. For 1% CS ratio, on Set11 dataset, 3-bit Q-PAN (2R) causes only 0.07 dB decrement in PSNR compared with PAN (2R).

CS RATIO	TEST DATASET	PAN (2R) (PSNR, SSIM)	1-BIT Q-PAN (2R) (PSNR, SSIM)	2-BIT Q-PAN (2R) (PSNR, SSIM)	3-BIT Q-PAN (2R) (PSNR, SSIM)
10%	SET11	26.26, 0.7951	22.78, 0.6294	23.25, 0.6624	24.53, 0.7253
10%	BSD68	25.13, 0.6947	22.59, 0.5327	23.23, 0.5811	24.25, 0.6465
4%	SET11	20.86, 0.5858	20.13, 0.5401	20.14, 0.5480	20.10, 0.5352
4%	BSD68	21.85, 0.5340	21.20, 0.4836	21.21, 0.4962	21.16, 0.4837
1%	SET11	17.52, 0.4108	16.73, 0.2537	16.45, 0.2274	17.45, 0.3974
1%	BSD68	19.15, 0.4097	18.12, 0.2437	17.91, 0.2273	18.96, 0.3888

Net$^+$ for CS ratios 1%, 4% and 10% on Set11 dataset. The average PSNR values are plotted in Figure 2. An interesting observation from Figure 2 and Tables 2, 4, 6 is that the quantized variants of enhanced plus models perform on par with the full-precision models even for the case of extreme compression of 1%. We repeat the experiment on BSD68 dataset (c.f. Figure 3) and document the findings.

Table 6. Comparison of average PSNR [dB], SSIM with PAN$^+$ (2R) and $1, 2, 3$ - bit Q-PAN$^+$ (2R) on the Set11 and BSD68 datasets. We observe that the 1-bit Q-PAN$^+$ (2R) has degradation of 0.69 dB in PSNR and 0.0320 in SSIM for 10% CS ratio measurements in BSD68 data when compared with PAN$^+$ (2R). But for 1% CS ratio on the same dataset, the model shows 0.08 dB improvement in PSNR over PAN$^+$ (2R).

CS RATIO	TEST DATASET	PAN$^+$ (2R) (PSNR, SSIM)	1-BIT Q-PAN$^+$ (2R) (PSNR, SSIM)	2-BIT Q-PAN$^+$ (2R) (PSNR, SSIM)	3-BIT Q-PAN$^+$ (2R) (PSNR, SSIM)
10%	SET11	26.53, 0.8017	25.05, 0.7436	25.82, 0.7796	26.37, 0.7982
10%	BSD68	25.30, 0.6972	24.61, 0.6652	25.05, 0.6844	25.24, 0.6959
4%	SET11	21.72, 0.6302	20.90, 0.5762	21.19, 0.5923	21.32, 0.6045
4%	BSD68	22.45, 0.5610	21.91, 0.5351	22.10, 0.5459	22.23, 0.5455
1%	SET11	17.34, 0.4143	17.48, 0.4055	17.52, 0.4111	17.56, 0.4119
1%	BSD68	19.08, 0.4190	19.16, 0.4166	19.20, 0.4189	19.22, 0.4196

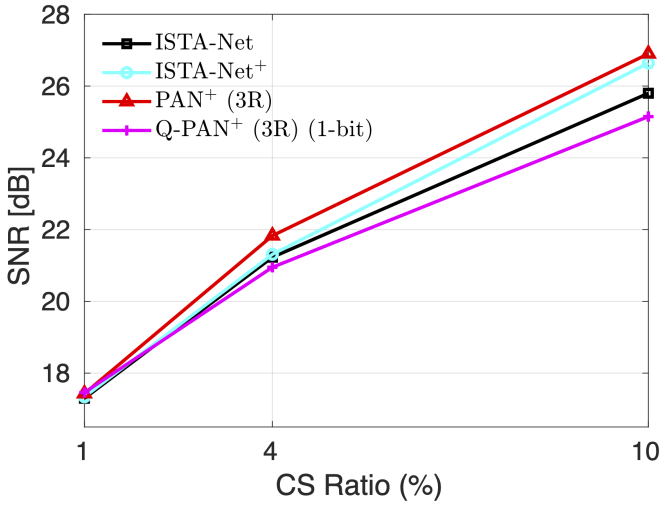

Figure 2. Set11 dataset - image reconstruction: Differentiation of Q-PAN$^+$ with full-precision PAN$^+$ (Ours), ISTA-Net$^+$, ISTA-Net in terms of PSNR performance over natural image reconstruction. The 1-bit Q-PAN$^+$ performs on par with the full-precision models at high levels of compression, i.e. CS ratio 1% where the number of observations is very low.

1.2. MR Image Reconstruction from Compressed Measurements

We now aim to validate our techniques on real-world data for which we consider the magnetic resonance image reconstruction problem. The PAN, PAN$^+$ models and their quantized counterparts Q-PAN, Q-PAN$^+$ are tested on MRI dataset [**?**] for compressed image reconstruction. We compare the Q-ISTA-Net and Q-ISTA-Net$^+$ models also to analyze the effect of quantization.

The training data for brain MRI comprises 800 images of size 256×256. We train the network models for compression ratios of 20%, 30%, 40%, and 50% . The training and inference are carried out on *Google Colab Pro* [**?**] GPUs. The models are trained for 200 epochs and take approximately 10 hours.

From the experimental results shown in Tables. 7, 8, 9, 10, 11, 12, we observe that the PAN$^+$ (2R), PAN$^+$ (3R), and ISTA-Net$^+$ are more robust to quantization, compared with the PAN (2R), PAN (3R) and ISTA-Net models. This could be attributed to the skip connections. A similar observation was made in the case of Set11 and BSD68 datasets as well. The validation on real-world brain MR data further strengthens the argument. Moreover, the 3-bit quantized variants have shown improved performance over the full-precision models across all CS ratios in the PAN$^+$ (2R), PAN$^+$ (3R), and ISTA-Net$^+$ models. This is an exciting finding, where full-precision models, which are prone to over-fit, can be replaced by better

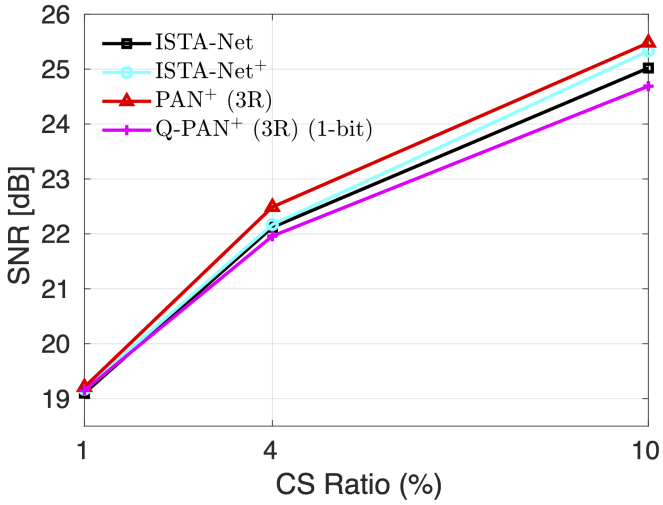

Figure 3. Image reconstruction performance on BSD68 dataset. We observe that at 1% CS ratio, the full-precision models and quantized model have nearly identical performance. We can conclude that, quantized models have reconstruction capabilities on par with the full-precision models at higher compression.

Table 7. Comparison of average PSNR [dB], SSIM for different compression ratios with ISTA-Net and its quantized variants $(1, 2, 3$-bit) Q-ISTA-Net on MRI datasets. We observe that 1-bit Q-ISTA-Net degrades by 1.81 dB in PSNR and by 0.0124 in SSIM for 50% compression ratio measurements. For the same number of observations one can observe that $2, 3$-bit Q-ISTA-Nets show an improvement of 1.05 dB and 1.81 dB in PSNR over the 1-bit Q-ISTA-Net.

CS RATIO	ISTA-NET (PSNR, SSIM)	1-BIT Q-ISTA-NET (PSNR, SSIM)	2-BIT Q-ISTA-NET (PSNR, SSIM)	3-BIT Q-ISTA-NET (PSNR, SSIM)
20%	38.30, 0.9483	36.11, 0..9106	37.92, 0.9383	38.41, 0.9431
30%	40.52, 0.9638	39.02, 0.9468	39.22, 0.9501	40.52, 0.9594
40%	42.12, 0.9727	40.55, 0.9610	40.01, 0.9614	42.13, 0.9693
50%	43.60, 0.9795	41.79, 0.9671	42.84, 0.9739	43.60, 0.9767

generalized, memory and computation efficient quantized models.

We carried out a comparative study between ISTA-Net, ISTA-Net$^+$ and proposed models. We evaluated the performance of the models on the MRI dataset. The corresponding results are shown in Figure 5 for visual inspection.

Figure 5 shows that the 1-bit Q-PAN$^+$(3R) performs on par with the full-precision ISTA-Net across various CS ratios. The performance of PAN$^+$(3R) is consistently superior to that of ISTA-Net$^+$.

Table 8. Comparison of average PSNR [dB] and SSIM on Brain MR images for different compression ratios with ISTA-Net$^+$ and its quantized variants $(1, 2, 3$ - bit) Q-ISTA-Net$^+$. We observe that 1-bit Q-ISTA-Net$^+$ degrades by 0.7 dB in PSNR and by 0.0064 in SSIM for 40% compression ratio measurements. For the same number of observations, 2-bit Q-ISTA-Net$^+$ offers 0.07 dB and 3-bit variant offers 0.31 dB improvement in PSNR over ISTA-Net$^+$.

CS RATIO	ISTA-NET$^+$ (PSNR, SSIM)	1-BIT Q-ISTA-NET$^+$ (PSNR, SSIM)	2-BIT Q-ISTA-NET$^+$ (PSNR, SSIM)	3-BIT Q-ISTA-NET$^+$ (PSNR, SSIM)
20%	38.73, 0.9490	37.92, 0.9378	38.69, 0.9446	38.91, 0.9467
30%	40.89, 0.9637	40.31, 0.9567	40.42, 0.9578	41.13, 0.9627
40%	42.52, 0.9729	41.82, 0.9665	42.59, 0.9707	42.83, 0.9719
50%	44.09, 0.9799	43.74, 0.9765	44.08, 0.9776	44.07, 0.9774

Table 9. Comparison of average PSNR [dB] and SSIM on Brain MR images for different compression ratios with PAN (3R) and $(1, 2, 3$ - bit) Q-PAN (3R). The 1-bit Q-PAN (3R) has degradation of PSNR by 3.72 dB and SSIM by 0.0612 for 30% compression ratio measurements. For the same number of observations one can observe that 2-bit Q-PAN (3R) has an improvement of 1.17 dB and 3-bit variant has improvement 2.1 dB of PSNR over 1-bit Q-PAN (3R).

CS RATIO	PAN (3R) (PSNR, SSIM)	1-BIT Q-PAN (3R) (PSNR, SSIM)	2-BIT Q-PAN (3R) (PSNR, SSIM)	3-BIT Q-PAN (3R) (PSNR, SSIM)
20%	38.00, 0.9433	34.76, 0.8840	34.19, 0.8684	35.96, 0.9518
30%	39.48, 0.9530	35.76, 0.8918	36.93, 0.9257	37.86, 0.9323
40%	41.57, 0.9671	36.25, 0.9148	38.42, 0.9381	39.54, 0.9514
50%	43.20, 0.9757	38.24, 0.9416	39.49, 0.9521	40.96, 0.9662

Table 10. Comparison of average PSNR [dB] and SSIM on Brain MR images with PAN$^+$ (3R) and its quantized variants $(1, 2, 3$ - bit) Q-PAN$^+$ (3R). At 40% CS ratio measurements, we observe that while the 1-bit Q-PAN$^+$ (3R) has degradation of PSNR by 0.3 dB and SSIM by 0.0030, the 2-bit Q-PAN$^+$ (3R) gives the same PSNR as the full precision model, PAN$^+$ (3R) and the 3-bit Q-PAN$^+$ (3R) offers an improvement of 0.3 dB over PAN$^+$ (3R).

CS RATIO	PAN$^+$ (3R) (PSNR, SSIM)	1-BIT Q-PAN$^+$ (3R) (PSNR, SSIM)	2-BIT Q-PAN$^+$ (3R) (PSNR, SSIM)	3-BIT Q-PAN$^+$ (3R) (PSNR, SSIM)
20%	38.90, 0.9505	38.17, 0.9400	36.96, 0.9257	39.11, 0.9486
30%	41.02, 0.9640	40.36, 0.9579	40.98, 0.9614	41.31, 0.9636
40%	42.71, 0.9729	42.41, 0.9699	42.71, 0.9715	43.01, 0.9728
50%	44.29, 0.9800	43.84, 0.9769	44.21, 0.9782	44.50, 0.9794

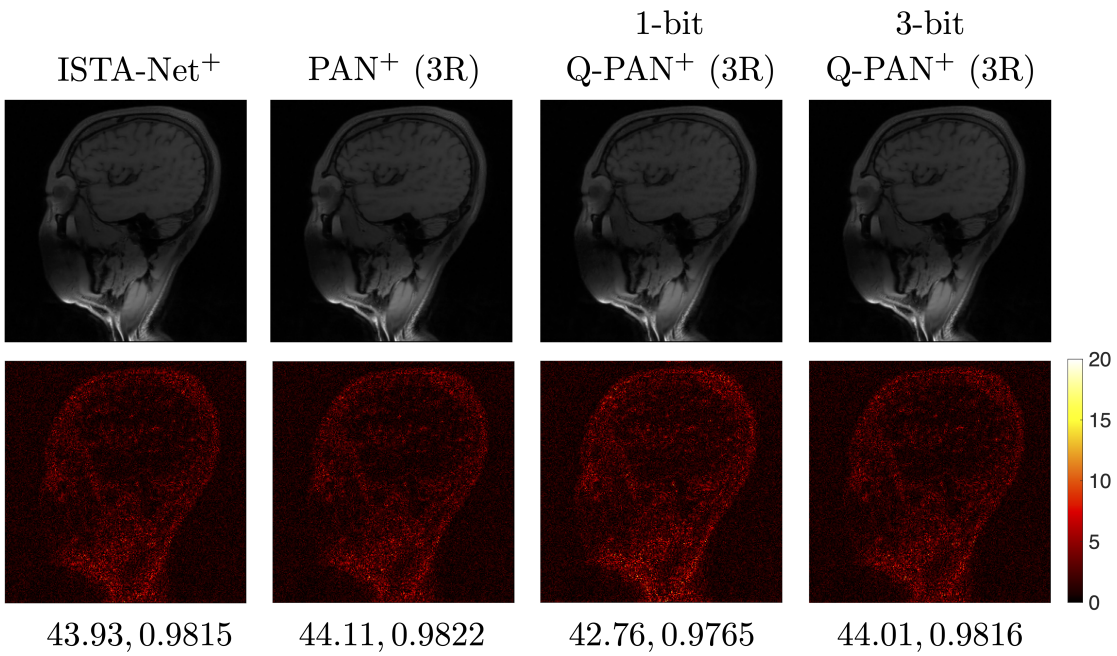

Figure 4. Brain MR image reconstruction for CS ratio 30% by ISTA-Net$^+$ and the proposed models. The bottom row visualizes the absolute difference between the ground truth and the reconstructed images. The 3-bit Q-PAN$^+$ (3R) has shown performance on par with the full-precision ISTA-Net$^+$ model. The numbers indicates the PSNR, SSIM values.

Table 11. Comparison of average PSNR [dB], SSIM for different CS ratios using PAN (2R) and its quantized variants ($1, 2, 3$ - bit) Q-PAN (2R) on the MRI dataset. In this comparative study, we observe that 1-bit Q-PAN (2R) suffered performance decrement of 2.98 dB in PSNR and 0.0258 in SSIM on average across different CS ratios. In the case of 1-bit Q-PAN (2R), the dip in PSNR is 2.46 dB, while 3-bit Q-PAN (2R) has 1.66 dB of PSNR drop in comparison with PAN (2R).

CS Ratio	PAN (2R) (PSNR, SSIM)	1-bit Q-PAN (2R) (PSNR, SSIM)	2-bit Q-PAN (2R) (PSNR, SSIM)	3-bit Q-PAN (2R) (PSNR, SSIM)
20%	38.55, 0.9478	35.56, 0.9071	34.47, 0.8730	35.28, 0.8878
30%	40.38, 0.9604	37.17, 0.9305	38.22, 0.9405	39.79, 0.9538
40%	42.12, 0.9703	38.88, 0.9509	39.63, 0.9486	40.89, 0.9626
50%	43.31, 0.9763	40.83, 0.9629	42.20, 0.9700	41.74, 0.9702

Table 12. Comparison of average PSNR [dB] and SSIM with PAN$^+$ (2R) and $(1, 2, 3$ - bit) Q-PAN$^+$ (2R) on the MRI dataset. The 3-bit Q-PAN$^+$ (2R) outperformed the PAN$^+$ (2R) for all CS ratios in terms of PSNR. Whereas the 2-bit Q-PAN$^+$ (2R) has shown on par performance with PAN$^+$ (2R). The PSNR performance degradation with 1-bit Q-PAN$^+$ (2R) compared to PAN$^+$ (2R) is minimal, and at most 0.63 dB.

CS RATIO	PAN$^+$ (2R) (PSNR, SSIM)	1-BIT Q-PAN$^+$ (2R) (PSNR, SSIM)	2-BIT Q-PAN$^+$ (2R) (PSNR, SSIM)	3-BIT Q-PAN$^+$ (2R) (PSNR, SSIM)
20%	38.82, 0.9492	38.19, 0.9396	38.55, 0.9438	39.13, 0.9486
30%	41.03, 0.9647	40.46, 0.9582	40.76, 0.9599	41.32, 0.9640
40%	42.70, 0.9728	42.23, 0.9687	42.60, 0.9709	43.04, 0.9730
50%	44.12, 0.9790	43.93, 0.9771	44.29, 0.9786	44.41, 0.9793

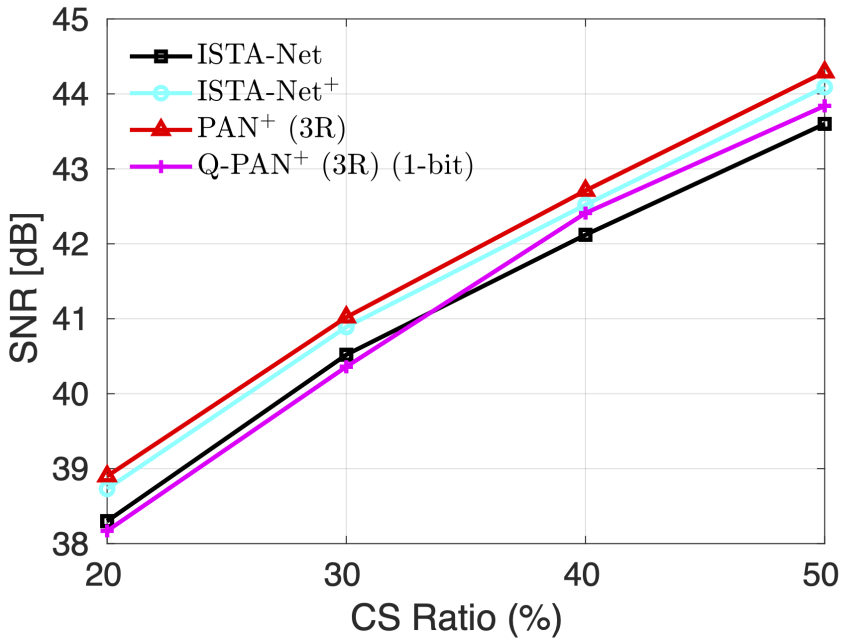

Figure 5. MR Image Reconstruction: Performance comparison of Q-PAN$^+$ with the full-precision models, namely, PAN (Ours), ISTA-Net, and ISTA-Net$^+$ in terms of PSNR as a function of CS ratio. One can infer that Q-PAN$^+$ with 1-bit quantization achieves image reconstruction performance on par with that of the full-precision models.