Compressed Sensing Inspired Neural Decoder for Undersampled MRI with Self-Assessment

This is the author's version of the article that has been presented at IEEE BioCAS2021
The editorial version of the paper is available at http://dx.doi.org/10.1109/BioCAS49922.2021.9644958

Compressed Sensing Inspired Neural Decoder for Undersampled MRI with Self-Assessment

Filippo Martinini †, Mauro Mangia *†, Fabio Pareschi ‡†, Riccardo Rovatti *†, Gianluca Setti ‡†
*DEI, †ARCES, University of Bologna, Italy - filippo.martinini@studio.unibo.it, {mauro.mangia2, riccardo.rovatti}@unibo.it
‡DET, Politecnico di Torino, Italy - {fabio.pareschi, gianluca.setti}@polito.it

***Abstract*—An important problem in magnetic resonance imaging (MRI) is the long time lapse required to acquire a fully sampled, high resolution scan. To speed up acquisition, Compressed Sensing (CS) has been used and recently coupled with Neural Networks (NN). In the latter setting, commonly CS has been split into two different problems: *i*) design of the encoder, or selection of the undersampling pattern, and *ii*) design of the decoder. A significant progress was recently introduced by a solution (called LOUPE) where encoding and decoding are simultaneously addressed. Here we propose an improvement of this model, called "regularized-LOUPE" (r-LOUPE), which add *measurement constraint* into the picture, resulting in a $\times 8$ speed-up in the MRI acquisition time. A further benefit of our methodology is that measurement constraint can be leveraged to implement a self-assessment tool able to predict the reconstruction error and to identify possible out-layers.**

I. Introduction

Magnetic Resonance Imaging (MRI) [1] is a non-invasive technology used to perform anatomical investigations. More specifically, MRI acquires 3D frequency volumes that when transformed to the pixel domain allow to easily visualize the inspected part of the body as a series of images (scans).

Unfortunately, MRI scan requires a long time lapse to be completed, leaving a patient in a situation of discomfort inside a large magnet, which partially hinders the wide adoption of this technology. Recently, Compressed Sensing (CS) [2]–[4] has been proposed as a way to accelerate MRI acquisitions, thanks to its capability to reconstruct a sparse signal starting from a limited number of linear measurements of it, and whose adoption has also been recently cleared by FDA [5]. More precisely, since MRI systems perform signal acquisition directly in the frequency domain, such an acceleration is accomplished by using an encoder which imposes an *undersampling* acquisition pattern in the frequency domain, that tells which frequencies to keep or drop. The resulting encoder is then coupled with a CS decoder stage that reconstructs an image well approximating the full resolution MRI scan [4].

To progress along this direction, recent contributions have mainly focused on modifying the decoder, by changing the optimization problem it solves with respect to classical CS where the focus is on reusing the undersampling pattern characterizing the encoder. In many cases, the encoding-decoding problem has been split into two independent problems, where the two blocks have been designed separately. Examples of adapted and data-driven decoders leveraging the statistic of the signals are *sparsifying dictionaries* [6] and total variation penalty [7] while encoder alternatives are, for instance, the random uniform [8], the variable density [9] and equi-spaced Cartesian with skipped lines [10].

The above are examples of a growing trend in the last ten year which relies on the use of Machine Learning (ML) tools in biomedical applications. Those exploiting CS are no exceptions, where ML has been widely used to design CS decoders [11]–[13].

A very successful example of using NN for fast MRI acquisition is reported in [14]–[16], where the authors presented an architecture, called LOUPE, that simultaneously co-designs encoder and decoder. LOUPE finds an optimal undersampling pattern that depends on the reconstruction method and vice versa. Other methods, inspired by LOUPE, that leverage the co-design have been recently proposed, for instance in [17].

Yet, these methods have the important drawback to be unable to preserve the input ground truth frequencies (measurements). Starting from this observation, we show that embedding the ability to preserve the measurements allows to achieve improved performances in terms of MRI acquisition time reduction. As an important byproduct, our approach allows to self-assess the quality of the reconstructed image and to estimate the reconstruction error at inference time. Finally, while the work presented here is based on LOUPE, our contributions are not theoretically limited to it.

The rest of the paper is organized as follows. Before presenting the proposed approach, Section II recaps CS theory and introduces the LOUPE architecture. Section II discusses the self-assessment approaches. Section III describes the adopted dataset and presents achieved results. In Section IV we draw the conclusions.

II. Compressed Sensing for MRI

Following the classical CS setting, we represent the generic input instance as a vector $\boldsymbol{a} \in \mathbb{R}^n$, supposing that each input vector is κ-sparse with respect to a basis defined by the columns of an orthonormal matrix $\boldsymbol{D} \in \mathbb{R}^{n \times n}$. In this setting, it is possible to express $\boldsymbol{a}$ as $\boldsymbol{a} = D\boldsymbol{\alpha}$, where $\boldsymbol{\alpha} \in \mathbb{R}^n$ contains only κ non-zero entries, with $\kappa \ll n$. This makes possible the acquisition/compression of $\boldsymbol{a}$ through a very simple operation

$$\boldsymbol{b} = \boldsymbol{A}\boldsymbol{a} = \boldsymbol{A}\boldsymbol{D}\boldsymbol{\alpha} \tag{1}$$

This is the author's version of the article that has been presented at IEEE BioCAS2021
The editorial version of the paper is available at http://dx.doi.org/10.1109/BioCAS49922.2021.9644958

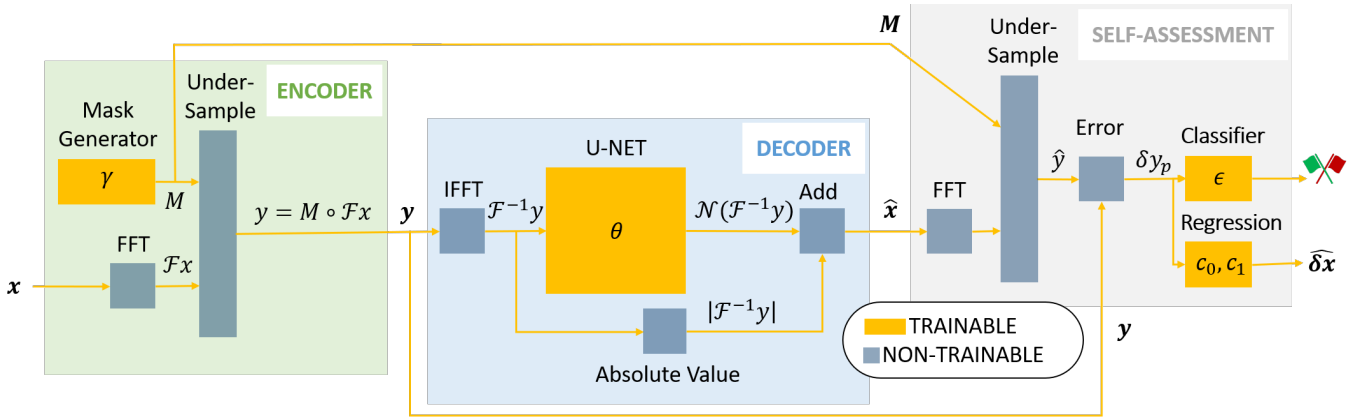

Fig. 1. Block scheme of the Encoder-Decoder couple along with the self-assessment mechanism.

where $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ is the *sensing matrix* and $\boldsymbol{b} \in \mathbb{R}^m$ represents the *measurements vector*, i.e., the encoder output. Note that: *i*) given $m < n$, the signal is being compressed; *ii*) (1) refers to the noiseless compression; in case of a noisy acquisition system, an m-dimensional vector $\boldsymbol{\nu}$ is added to $\boldsymbol{b}$ to model all non-idealities.

In order to solve (1), it is necessary to find the operator that "inverts" $\boldsymbol{B} = \boldsymbol{AD} \in \mathbb{R}^{m\times n}$, the matrix linking the sparse representation with the measurement vector. Since $m < n$, the "inversion" of $\boldsymbol{B}$ is an ill-conditioned problem. Nonetheless, the κ-sparse hypothesis adds enough information on the signal structure to make possible the recovery of $\boldsymbol{\alpha}$, and then $\boldsymbol{a}$, from $\boldsymbol{b}$. One of the most known methods employed for the reconstruction of $\boldsymbol{\alpha}$ is Basis Pursuit (BP) [2], [3], that searches for the sparsest vector $\hat{\boldsymbol{\alpha}}$ that solves the problem:

$$\hat{\boldsymbol{\alpha}} = \arg\min_{\boldsymbol{\alpha}\in\mathbb{R}^n} \|\boldsymbol{\alpha}\|_1 \quad \text{s.t.} \quad \boldsymbol{AD\alpha} = \boldsymbol{b} \tag{2}$$

where the ℓ_1-norm $\|\cdot\|_1$ promotes sparsity and the constraint forces $\hat{\boldsymbol{\alpha}}$ to be a counter-image of $\boldsymbol{b}$ with respect to $\boldsymbol{B}$. To make the recovery effective, standard CS theory requires: i) a number of elements in the measurement vector higher the a suitable lower bound and ii) that the sensing matrix $\boldsymbol{A}$ possesses specific properties such as either being a restricted isometry or being incoherent with respect to $\boldsymbol{D}$ [3].

In the MRI context, the input signal is an image modelled by $\boldsymbol{x} \in \mathbb{R}^{N\times N}$ that is acquired through its 2-dimensional Fourier transform $\mathcal{F}(\boldsymbol{x})$. Undersampling is modelled by considering a binary mask $\boldsymbol{M} \in \{0,1\}^{N\times N}$ where only $m = rN^2$ elements are non-null, with $0 < r < 1$, and by saying that the actual measurements are the elements of

$$\boldsymbol{y} = \boldsymbol{M} \circ \mathcal{F}(\boldsymbol{x}) \tag{3}$$

where, $\circ$ represents the Hadamard element-wise product. Note that, in general, $\boldsymbol{y}$ is $\mathbb{C}^{N\times N}$. In practice only the elements of $\boldsymbol{y}$ that are not set to zero by masking are needed, thus yielding the saving in acquisition time. As a result, a physical system acquiring only the signal components associated to the ones in $\boldsymbol{M}$ achieves a speed-up factor up to $R = 1/r$.

The decoder receives measurements $\boldsymbol{y}$ corresponding to the undersampling patter and reconstruct the MRI scan by solving (2) where some sparse representations have been considered [18].

A. LOUPE

The LOUPE architecture is conceptually an *autoencoder*, i.e., it has an encoding $\text{Enc}_{\boldsymbol{\gamma}}$ and a decoding stage $\text{Dec}_{\boldsymbol{\theta}}$, depending on parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\theta}$, respectively. Hence, the output image can be expressed as $\hat{\boldsymbol{x}} = \text{Dec}_{\boldsymbol{\theta}}(\text{Enc}_{\boldsymbol{\gamma}}(\boldsymbol{x}))$, and the parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\theta}$ need to be trained to make $\hat{\boldsymbol{x}}$ as close as possible to $\boldsymbol{x}$ despite the fact that intermediate encoded signal has smaller dimensionality with respect to the original one.

In the encoder, parameters $\boldsymbol{\gamma}$ control the generation of the random binary mask $\boldsymbol{M}(\boldsymbol{\gamma}) \in \{0,1\}^{N\times N}$ so that

$$\boldsymbol{y} = \text{Enc}_{\boldsymbol{\gamma}}(\boldsymbol{x}) = \boldsymbol{M}(\boldsymbol{\gamma}) \circ \mathcal{F}(\boldsymbol{x})$$

.

Furthermore, the mask generation mechanism is the same described in [15] and is sketched in Section III.

The decoder starts from a first approximation of the original image $\mathcal{F}^{-1}(\boldsymbol{y})$ and then applies a correction $\mathcal{N}_{\boldsymbol{\theta}}(\mathcal{F}^{-1}(\boldsymbol{y}))$ calculated by a deep neural network $\mathcal{N}_{\boldsymbol{\theta}}$ that is a slightly modified version of the U-NET network [19] with an additional residual layer, and whose parameters $\boldsymbol{\theta}$ are trained to produce

$$\hat{\boldsymbol{x}} = \left|\mathcal{F}^{-1}(\boldsymbol{y})\right| + \mathcal{N}_{\boldsymbol{\theta}}\left(\mathcal{F}^{-1}(\boldsymbol{y})\right)$$

as the best possible estimation of $\boldsymbol{x}$. Note that the modulus $|\cdot|$ is needed since the inverse of a subsampled 2-dimensional Fourier transform is not necessarily a matrix of real numbers.

Figure 1 gives a visual representation of the auto-encoding architecture (LOUPE architecture) as well as of the self-assessment stage that is one of the novelties of this contribution.

In its original conception [14]–[16], the parameters of this architecture are trained by defining the error

$$\delta\boldsymbol{x}_l = \|\text{Dec}_{\boldsymbol{\theta}}(\text{Enc}_{\boldsymbol{\gamma}}(\boldsymbol{x})) - \boldsymbol{x}\|_l$$

This is the author's version of the article that has been presented at IEEE BioCAS2021
The editorial version of the paper is available at http://dx.doi.org/10.1109/BioCAS49922.2021.9644958

where $\|\cdot\|_l$ is the ℓ_l norm and using the loss function

$$L(\boldsymbol{\gamma}, \boldsymbol{\theta}) = \mathbf{E}_{\boldsymbol{x}}[\delta \boldsymbol{x}_1] \tag{4}$$

where $\mathbf{E}_{\boldsymbol{x}}[\cdot]$ stands for expectation over all possible $\boldsymbol{x}$.

B. Loss regularization by measurement constraint

As previously sketched, LOUPE is similar to CS as it tries to recover the original instance from the measurement vector knowing that the information content in the original domain is redundant. i.e., that the sensing process preserves enough information to make decoding effective. What is different from classic CS is the set of images in which the decoder looks for the reconstruction: LOUPE does not force $\boldsymbol{y}$ to be obtained by a linear projection of $\hat{\boldsymbol{x}}$. Nevertheless, if recovery succeeds, then, considering (3), $\hat{\boldsymbol{x}}$ should be close enough to $\boldsymbol{x}$ to make

$$\delta \boldsymbol{y}_p = \left\| \boldsymbol{y} - \boldsymbol{M} \circ \mathcal{F}\left(\hat{\boldsymbol{x}}\right) \right\|_p \tag{5}$$

vanishing. Hence, smaller $\delta \boldsymbol{y}_p$ are certainly preferable to larger ones. Such a prior can be introduced in the network training by adding a regularization term and thus considering the alternative loss

$$L'_p(\boldsymbol{\gamma}, \boldsymbol{\theta}) = \mathbf{E}_{\boldsymbol{x}} \left[\phi \delta \boldsymbol{x}_1 + (1-\phi) \delta \boldsymbol{y}_p \right] \tag{6}$$

where ϕ is a tunable weight that balances the two contributions. With this approach, it is possible to emulate the constraint in (2). We call a model trained using L'_p a regularized-LOUPE (r-LOUPE$_p$) and consider $p \in \{1, 2\}$.

C. Adding self-assessment capabilities

Though (5) should vanish for properly recovered signals, when the reconstruction fails then $\delta \boldsymbol{y}_p$ may be expected to be substantially larger than 0. Hence, $\delta \boldsymbol{y}_p$, which can be computed starting from the knowledge of the actual measurements $\boldsymbol{y}$ only, has a magnitude that can be seen as a proxy of the decoder performance.

This can be exploited in a self-assessment stage that follows reconstruction and whose aim is to give the user additional information on the quality of the output image.

As shown in Fig. 1, once $\boldsymbol{x}$ is estimated, we first compute $\delta \boldsymbol{y}_p$. Then, we may either match $\delta \boldsymbol{y}_p$ against a threshold ϵ to raise a warning if $\delta \boldsymbol{y}_p > \epsilon$ hints at a bad reconstruction, or use it to compute an estimation $\widehat{\delta \boldsymbol{x}}_l = \widehat{\delta \boldsymbol{x}}_l(\delta \boldsymbol{y}_p)$ of the output quality based on the value of $\delta \boldsymbol{y}_p$. For simplicity's sake we will concentrate on affine estimations $\widehat{\delta \boldsymbol{x}}_l(\delta \boldsymbol{y}_p) = c_0 \delta \boldsymbol{y}_p + c_1$.

III. Numerical evidence

A. dataset

Our models have been trained and tested using the freely available data-set "Brain MRI segmentation"[1] that contains brain MRI images.

The database consists of 2D brain scans acquired from 110 patients included in The Cancer Genome Atlas lower-grade glioma collection[2]. The number of slices varies from 20 to 88 depending on the patient, and we exclude some of the least informative images so to obtain a training set consisting of 2753 scans (80%) and a test set containing 688 scans (20%). The original image dimensions are 256×256, we resize all the scans to obtain 128×128 images. We also re-scale each scan to have pixel values in the range [0,1].

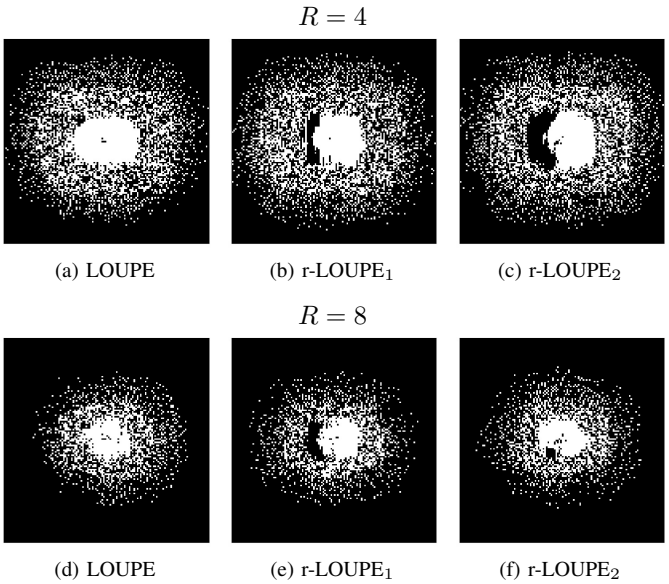

(a) LOUPE (b) r-LOUPE$_1$ (c) r-LOUPE$_2$

(d) LOUPE (e) r-LOUPE$_1$ (f) r-LOUPE$_2$

Fig. 2. Matrices containing the probability of each entry of $\boldsymbol{M}$ to be one. White means one while black stands for zero. Most values are very close to these boundaries. Top plots are for $R = 4$, bottom plots are for $R = 8$.

B. results

The main figure of merit we adopt for the assessment of the reconstruction quality is the Peak Signal-to-Noise Ratio (PSNR) defied as:

$$\mathrm{PSNR} = 10 \log_{10} \left(\frac{\max\{\boldsymbol{x}\}}{\delta \boldsymbol{x}_2} \right) = -10 \log_{10} \left(\delta \boldsymbol{x}_2 \right) \tag{7}$$

since in our case $\max\{\boldsymbol{x}\} = 1$. We considered speed-up values $R \in \{4, 8\}$, i.e., 25% and 12.5% of ones in $\boldsymbol{M}$. Furthermore, regularization weights characterizing r-LOUPE$_p$ are such that $0.99 \leq \phi < 1$.

The undersampling patterns $\boldsymbol{M}$ adopted by the three different approaches are the output of trained masks generator blocks [15]. In particular, random instances $\boldsymbol{M}$ are generated such that the probability of each entry to be 1 is equal to the corresponding entry of a $N \times N$ trained matrix. Fig. 2 reports such probability matrices for LOUPE, r-LOUPE$_1$ and r-LOUPE$_2$ in case of $R = 4$ (top plots) and $R = 8$ (bottom plots).

Table I shows averages PSNR along with the fraction of cases for which the PSNR achieved by r-LOUPE$_1$ or r-LOUPE$_2$ exceeds the PSNR of LOUPE. Both r-LOUPE$_1$ and r-LOUPE$_2$ outperform LOUPE in terms of average PSNR in at least 97.7% of the cases. Furthermore, Fig. 3 reports

[1] https://www.kaggle.com/mateuszbuda/lgg-mri-segmentation

[2] https://www.cancer.gov/tcga

This is the author's version of the article that has been presented at IEEE BioCAS2021
The editorial version of the paper is available at http://dx.doi.org/10.1109/BioCAS49922.2021.9644958

TABLE I
AVERAGE PSNR AND FRACTION OF MRI IMAGES FOR WHICH R-LOUPE$_p$ ACHIEVES HIGHER PSNR WITH RESPECT TO LOUPE

Model	Aver. PSNR [dB]		% of success	
	$R=4$	$R=8$	$R=4$	$R=8$
LOUPE	38.65	32.67	-	-
r-LOUPE$_1$	39.26	34.15	98.7	97.7
r-LOUPE$_2$	39.86	34.23	99.9	98.4

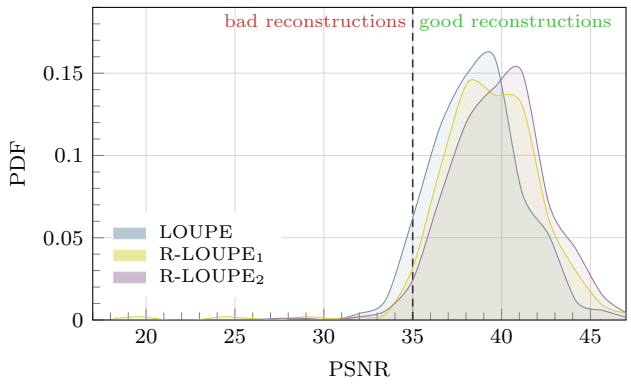

(a)

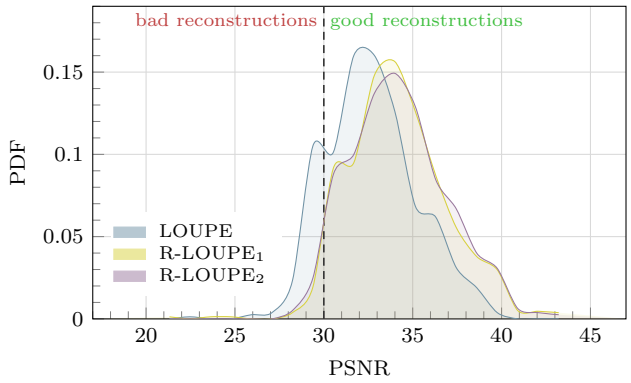

(b)

Fig. 3. Estimated probability density functions of PSNR for LOUPE, r-LOUPE$_1$ and r-LOUPE$_2$. In case of (a) $R=4$ and PSNR values lower than 35 dB are labeled as *bad reconstruction*. In case of (b) $R=8$ and the PSNR threshold is 30 dB.

the PSNR estimated probability density functions (PDF) for $R=4$ (Fig. 4b) and $R=8$ (Fig. 4c). To give a visual representation of the reconstructed scans we report in Fig. 4 an example of the original scan and the corresponding r-LOUPE$_2$ reconstructed images with speed-up factors $R=\{4,8\}$.

To evaluate the self-assessment capability, we define lower bounds for PSNR such that reconstructions with PSNR values larger than those thresholds are labeled as *good reconstructions*. Adopted thresholds are 35dB and 30dB respectively for R equal to 4 and 8. The proposed classifier tries to mark a reconstruction as good or bad by comparing $\delta \boldsymbol{y}_p$ with a threshold ϵ. Depending on the value of ϵ we could have a certain fraction of good reconstruction misclassified as bad

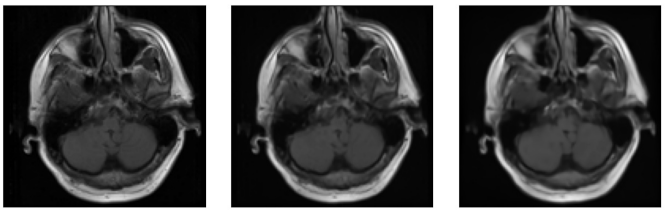

(a) original scan (b) $R=4$ (c) $R=8$

Fig. 4. Original MRI scan (a) andh reconstructed images with r-LOUPE$_2$ with: (b) $R=4$ and PSNR=37.5 dB, (c) $R=8$ and PSNR=31.6 dB.

TABLE II
AUC OF THE ROC FOR ALL CONSIDERED MODELS WITH $\delta \boldsymbol{y}_1$ OR $\delta \boldsymbol{y}_2$ AND WITH $R=4$ OR $R=8$

Model	AUC using $\delta \boldsymbol{y}_1$		AUC using $\delta \boldsymbol{y}_2$	
	$R=4$	$R=8$	$R=4$	$R=8$
LOUPE	0.971	0.934	0.969	0.955
r-LOUPE$_1$	0.992	0.978	0.991	0.979
r-LOUPE$_2$	0.992	0.983	0.993	0.986

(false positive rate, FPR) and a fraction of bad reconstructions correctly classified as bad (true positive rate, TPR) where both FPR and TPR are in $[0,1]$. The whole classifier performance is given by the so-called Receiver Operating Characteristic (ROC), i.e., the set of pairs of FPR, TPR associated to different ϵ values. The ideal classifier has FPR$=0$, TPR$=1$ for any possible ϵ corresponding to a ROC curve degenerating to the point (0,1). As a result, a possible figure of merit is the Area Under the ROC curve (AUC). The closer the AUC to one, the better the classifier.

Table II reports AUC values for LOUPE, r-LOUPE$_1$ and r-LOUPE$_2$ highlighting how the introduced regularization in the training phase also increases the classifier capability. It also shows a negligible difference in terms of AUC between the choice of $\delta \boldsymbol{y}_1$ or $\delta \boldsymbol{y}_2$. Moreover, r-LOUPE$_2$ behaves better if compared with r-LOUPE$_1$. As such, we focus only on r-LOUPE$_2$ for which ROC curve shown in Fig. 5.

As anticipated in the previous section, self-assessment can

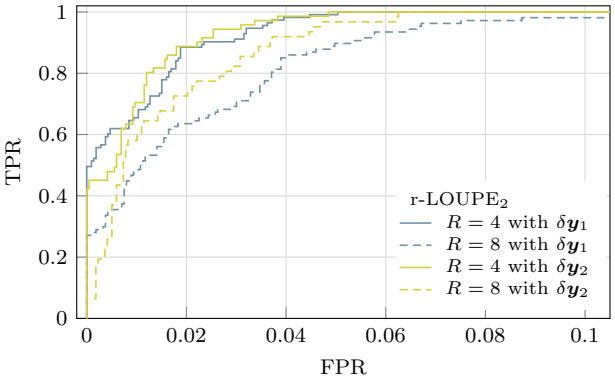

Fig. 5. ROC curve for r-LOUPE$_2$ with $R=\{4,8\}$ and both $\delta \boldsymbol{y}_1$ and $\delta \boldsymbol{y}_2$.

This is the author's version of the article that has been presented at IEEE BioCAS2021
The editorial version of the paper is available at http://dx.doi.org/10.1109/BioCAS49922.2021.9644958

TABLE III
AVERAGE ABSOLUTE AND RELATIVE ERRORS IN ESTIMATING $\delta\boldsymbol{x}_l$ FROM $\delta\boldsymbol{y}_p$

l	p	$\lvert\delta\boldsymbol{x}_l - \widehat{\delta\boldsymbol{x}_l}\rvert$		$\lvert\delta\boldsymbol{x}_l - \widehat{\delta\boldsymbol{x}_l}\rvert / \lvert\delta\boldsymbol{x}_l\rvert$	
		$R=4$	$R=8$	$R=4$	$R=8$
1	1	4.58e-04	8.19e-04	0.0752	0.0776
1	2	4.67e-05	1.01e-04	0.530	0.350
2	1	6.39e-04	1.06e-03	0.104	0.103
2	2	3.50e-05	8.61e-05	0.434	0.266

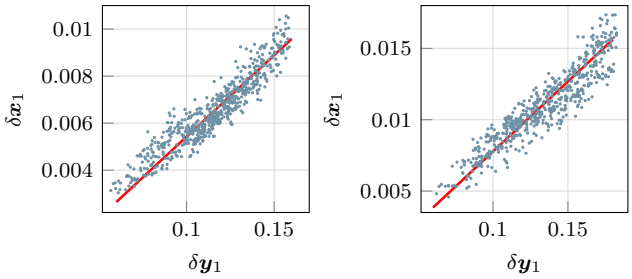

(a) $R = 4$ (b) $R = 8$

Fig. 6. Scatter plots reporting couples of $\delta\boldsymbol{x}_1$ and $\delta\boldsymbol{x}_2$ associated to r-LOUPE$_2$. Plots also report lines $\widehat{\delta\boldsymbol{x}}_1(\delta\boldsymbol{y}_1) = c_0\delta\boldsymbol{y}_1 + c_1$, used to predict $\delta\boldsymbol{x}_1$. We found: (a) $R = 4$ with $c_0 = 0.06937$ and $c_1 = -0.001515$; (b) $R = 8$ with $c_0 = 0.09815$ and $c_1 = -0.002042$.

be also exploited to predict $\delta\boldsymbol{x}_l$ by looking at $\delta\boldsymbol{y}_p$ with $l, p \in \{1, 2\}$. Considering r-LOUPE$_2$, as a first and simple choice we focus on a linear predictor estimating $\delta\boldsymbol{x}_l$ with $\widehat{\delta\boldsymbol{x}_l} = c_0\delta\boldsymbol{y}_p + c_1$. To reduce the effect of outliers, for any given l and p, we compute the coefficients c_0 and c_1, by a Theil-Sen linear regression [20], [21] on the training set. Performances on the test set in terms of prediction error are reported in Table III. An example of a visual representation of the estimation of $\delta\boldsymbol{x}_1$ given $\delta\boldsymbol{y}_1$ is reported in Fig. 6 for both $R = 4$ and $R = 8$. A relative error always smaller than 8% in the prediction of $\delta\boldsymbol{x}_1$ shows that the $\delta\boldsymbol{y}$ proxy can be quite effective.

IV. CONCLUSION

The method we present is based on the LOUPE architecture for the acquisition of subsampled MRI. We show that the quality of the reconstructed images can be improved by embedding the measurements constraint via a regularization term. On the average, such an improvement results in a PSNR increase of up to 1.6dB for speed-up factors of 4 and 8.

Moreover, we show that the error between the re-acquisitions of the reconstructed image and the actual acquisitions is a good proxy of the quality of the reconstructed image measured against the unknown ground truth. This gives raise to two self-assessment tools: the first enables to discriminate, at inference time, whether a reconstructed scan meets a given quality standard; the second estimates the image reconstruction error. We demonstrate how these two methods can be efficiently applied to have an additional degree of robustness to out-layers and, in general, to predict how well an image has been reconstructed.

ACKNOWLEDGMENTS

This work was supported in part by the Italian Ministry for Education, University and Research (MIUR) under the program "Dipartimenti di Eccellenza" (2018-2022) and in part by the SmartData Center of Politecnico di Torino.

REFERENCES

[1] J. L. Prince and J. Links, *Medical Imaging Signals and Systems, 2nd edition*. Pearson, 2015.
[2] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 4 2006. [Online]. Available: http://ieeexplore.ieee.org/document/1614066/
[3] E. Candes and T. Tao, "Decoding by Linear Programming," *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203–4215, 12 2005. [Online]. Available: http://ieeexplore.ieee.org/document/1542412/
[4] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 72–82, 2008.
[5] FDA, "Fda clears compressed sensing mri acceleration technology from siemens healthineers," 2 2017.
[6] S. Ravishankar and Y. Bresler, "MR Image Reconstruction From Highly Undersampled k-Space Data by Dictionary Learning," *IEEE Transactions on Medical Imaging*, vol. 30, no. 5, pp. 1028–1041, 5 2011. [Online]. Available: http://ieeexplore.ieee.org/document/5617283/
[7] Shiqian Ma, Wotao Yin, Yin Zhang, and A. Chakraborty, "An efficient algorithm for compressed MR imaging using total variation and wavelets," in *2008 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 6 2008, pp. 1–8. [Online]. Available: http://ieeexplore.ieee.org/document/4587391/
[8] U. Gamper, P. Boesiger, and S. Kozerke, "Compressed sensing in dynamic MRI," *Magnetic Resonance in Medicine*, vol. 59, no. 2, pp. 365–373, 2 2008. [Online]. Available: http://doi.wiley.com/10.1002/mrm.21477
[9] Zhongmin Wang and G. Arce, "Variable Density Compressed Image Sampling," *IEEE Transactions on Image Processing*, vol. 19, no. 1, pp. 264–270, 1 2010. [Online]. Available: http://ieeexplore.ieee.org/document/5256257/
[10] J. P. Haldar, D. Hernando, and Zhi-Pei Liang, "Compressed-Sensing MRI With Random Encoding," *IEEE Transactions on Medical Imaging*, vol. 30, no. 4, pp. 893–903, 4 2011. [Online]. Available: http://ieeexplore.ieee.org/document/5599301/
[11] M. Mangia, L. Prono, A. Marchioni, F. Pareschi, R. Rovatti, and G. Setti, "Deep Neural Oracles for Short-window Optimized Compressed Sensing of Biosignals," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 14, no. 3, pp. 545–557, 6 2020. [Online]. Available: https://ieeexplore.ieee.org/document/9044777/
[12] L. Prono, M. Mangia, A. Marchioni, F. Pareschi, R. Rovatti, and G. Setti, "Deep Neural Oracle with Support Identification in the Compressed Domain," in *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 10, no. 4. Institute of Electrical and Electronics Engineers Inc., 12 2020, pp. 458–468.
[13] 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. [Online]. Available: http://ieeexplore.ieee.org/document/7560597/
[14] C. D. Bahadir, A. V. Dalca, and M. R. Sabuncu, "Learning-Based Optimization of the Under-Sampling Pattern in MRI," 2019, pp. 780–792. [Online]. Available: http://link.springer.com/10.1007/978-3-030-20351-1_61
[15] C. D. Bahadir, A. Q. Wang, A. V. Dalca, and M. R. Sabuncu, "Deep-Learning-Based Optimization of the Under-Sampling Pattern in MRI," *IEEE Transactions on Computational Imaging*, vol. 6, pp. 1139–1152, 2020.
[16] J. Zhang, H. Zhang, A. Wang, Q. Zhang, M. Sabuncu, P. Spincemaille, T. D. Nguyen, and Y. Wang, "Extending LOUPE for K-Space Under-Sampling Pattern Optimization in Multi-coil MRI," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 12450 LNCS. Springer Science and Business Media Deutschland GmbH, 2020, pp. 91–101.

[17] H. K. Aggarwal and M. Jacob, "J-Modl: Joint model-based deep learning for optimized sampling and reconstruction," *IEEE Journal on Selected Topics in Signal Processing*, vol. 14, no. 6, pp. 1151–1162, 2020.

[18] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI: A look at how CS can improve on current imaging techniques," pp. 72–82, 2008.

[19] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 9351, pp. 234–241, 2015.

[20] H. Theil, "A Rank-Invariant Method of Linear and Polynomial Regression Analysis," in *Proceedings of the Royal Netherlands Academy of Sciences 53 (1950) Part I: 386-392, Part II: 521-525, Part III: 1397-1412*. Springer, Dordrecht, 1950. [Online]. Available: https://link.springer.com/chapter/10.1007/978-94-011-2546-8_20

[21] P. K. Sen, "Estimates of the Regression Coefficient Based on Kendall's Tau," *Journal of the American Statistical Association*, vol. 63, no. 324, pp. 1379–1389, 1968.