Signal Processing 228 (2025) 109746

Contents lists available at ScienceDirect

Signal Processing

journal homepage: www.elsevier.com/locate/sigpro

Incremental Undersampling MRI Acquisition With Neural Self Assessment

Filippo Martinini [a,*], Mauro Mangia [a], Alex Marchioni [a], Gianluca Setti [b], Riccardo Rovatti [a]

[a] University of Bologna, Via Toffano 2/2, Bologna, 40125, Italy
[b] King Abdullah University of Science and Technology, Thuwal, 23954, Saudi Arabia

ARTICLE INFO

Keywords:
Accelerated MRI
Adaptive acquisition
Deep neural network
Incremental acquisition
Quality assessment

ABSTRACT

Accelerated MRI acquisition is widely adopted and basically consists in undersampling the current slice at the cost of a quality degradation. What samples to skip is determined by an encoder, while the quality loss is partially compensated by the use of a decoder. The hypothesis behind accelerated MRI acquisition is that to higher acceleration factors always correspond lower reconstruction qualities with an undersampling pattern that is usually fixed at design time, neglecting adaptability on the slice acquired at inference time. This paper proposes a novel accelerated MRI acquisition method that enables single-slice adaptation by dividing the acquisition into incremental batches and estimating the reconstruction quality at the end of each batch. The acquisition terminates as soon as the target quality is reached. We demonstrate the efficacy of our novel method using a state-of-the-art neural model capable of jointly optimizing the encoder and decoder. To estimate the current quality of the slice we reconstruct and propose a neural quality predictor. We demonstrate the advantages of our novel acquisition method compared to classic acquisition for two different datasets and for both line-constrained and unconstrained Cartesian sampling strategies (theoretically implementable via 2D and 3D imaging respectively).

1. Introduction

Magnetic Resonance Imaging (MRI) is widely adopted in the clinical world, to investigate the state of health of critical regions such as the brain, the heart and the musculoskeletal system.

The MRI output is typically a 3D volume composed of 2D slices, although 3D imaging is also possible [1]. These slices represent sections of the region under analysis. Each slice is formed by acquiring samples in what is known as the *k-space*, which is a matrix in the spatial frequency domain. Each point in k-space corresponds to a specific spatial frequency, and its value is related to the amplitude and phase of the signal generated by the processing nuclear relaxation signals in response to applied magnetic fields. The process of filling k-space involves applying a sequence of gradients, with the frequency-encoding and phase-encoding gradients determining the trajectory of acquisition. Human-readable high-resolution images are then generated by transforming each slice into the spatial domain using the Inverse Fourier Transform. The number of slices per volume and the slice dimensions (usually expressed in pixels) are variable, depending on the clinical requirements and the resources of the MRI machinery.

While acquiring the full k-space ensures maximum resolution, it is often necessary to reduce acquisition time by undersampling. This approach does lead to some loss of resolution. However, the literature shows plenty of studies to retain as much image quality as possible, effectively balancing the trade-off between speed and resolution.

Researchers who pioneered this research demonstrated that even if a portion of k-space is not acquired (to save time), it is possible to infer it, and obtain high-quality images. Eventually Compressed Sensing (CS) [2,3] was introduced into the MRI acquisition scheme [4,5] to improve results. With CS-based skipping and inferring strategies, researchers achieved faster acquisition with negligible quality losses. This idea brought to the formalization of a general two-step implementation framework, where first an encoder undersamples the image in the k-space. Second, a decoder compensates the effect of the undersampling, thus producing images with reasonable high quality [6]. Compressed Sensing (CS) revolutionized MRI by transforming the image reconstruction process into an optimization problem, typically solved by minimizing hand-crafted regularization terms [7,8].

Since the breakthrough of Deep Learning (DL) and Deep Neural Network (DNN), the field veered in the data-driven direction. The first emerging DNN-based models, specific for fast MRI acquisition, were based on the so-called UNET [9,10], that gave way to prolific research branches. For example, [11] leverages multi-contrast MRI images in the training process; [12] justifies the use of novel attention layers inside the model; [13,14] explore the use of Generative Adversarial Neural

* Corresponding author.
E-mail address: filippo.martinini@unibo.it (F. Martinini).

https://doi.org/10.1016/j.sigpro.2024.109746
Received 28 March 2024; Received in revised form 29 September 2024; Accepted 15 October 2024
Available online 23 October 2024

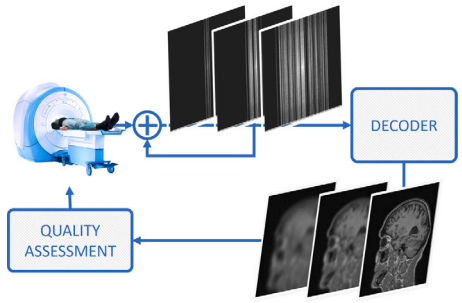

Fig. 1. Incremental Acquisition scheme. At every iteration the MRI machine acquires a batch of samples, that are accumulated with the previously acquired samples in the k-space. A decoder block recreates the anatomical image which quality is assessed by a specific block. If a target quality is reached the acquisition stops, otherwise another iteration is launched.

Networks (GAN); [15,16] specialize on DL-based denoising applied to MRI images; [17] interweaves graph learning and DNN-based image reconstruction. However, DNNs lack full explainability, which is a critical aspect in the medical field, and efforts to address this are ongoing, as seen in [18].

The above methods introduce novelty by upgrading the DNN-based decoder, but they leave the undersampling task to simple encoders, e.g., uniform random sample skipping [19], or variable density [20] and Line-Constrained sampling with skipped lines [21] (indicated here as *Line-Constrained*). Different sampling schemes also affect the artifacts acquisition differently. Radial sampling is generally more robust to motion, while Line-Constrained sampling is more prone to artifacts like blurring. Faster acquisitions reduce the time during which motion can occur. The encoder importance has been widely overlooked, and only recently novel DNN-based methods pioneered the joint encoder–decoder optimization reaching superior performances. For example, [22,23] design trainable continuous matrices that encode the probability of every sample in the k-space to be sampled. Recently, methods such as [24–26] further adapted [23], showing interesting results. In particular, [24] explores various loss functions to improve the joint dependency between the encoder and decoder, [25] refines gradient estimation strategies and replaces the UNET-like decoder with an unrolled architecture, and [26] extends the joint learning strategy to incorporate radial-constrained sampling patterns. Alternative trainable encoder architectures have also been proposed. For example, [27] relies on the Softmax function to select frequencies; [28] performs a non-uniform Fourier transform with continuously defined sampling locations; [29] builds a reinforcement learning problem. Recently, the ability of DL-based methods to reconstruct images acquired with ¼ of the standard time has been determined to be of diagnostic equivalence with the conventionally acquired images for knee MRI [30].

Typically, at least in the ML framework, the fast acquisition problem is tackled by optimizing the acquisition/reconstruction over a dataset made of several volumes, each comprising several slices. In general, central slices are complex and abound with details, while external slices are plainer and less sophisticated. It follows that it is counter-productive to give the same span of attention (acquisition-speaking) to all the slices within the same volume. Nevertheless, this is what the classic acquisition does. Intuitively, it discards the single-slice optimization, that consequently precludes adaptability and quality control.

Inspired by [31], we design a novel acquisition strategy by dividing the classic slice acquisition into batches, each implementing three steps: (*i*) sensing, (*ii*) reconstruction, and (*iii*) quality estimation. Based on the estimated quality the sequence continues or terminates. This procedure is depicted in Fig. 1. Our strategy adapts to the information content of the slice and gains control over the slice reconstruction quality, hence overcoming a great limitation of the classic acquisition. Because of its properties, we dub our method the *incremental adaptive MRI acquisition*

This work is not the first, in the MRI context, to build slice-adaptive k-space acquisition patterns (also referred to as *masks* or *binary masks*) [32–34]. In [32], authors create slice-specific sampling patterns based on a first acquisition batch of a low-frequency portion of the k-space. Similarly, [33,34] adapt the k-space acquisition iteratively across a predefined number of K iterations. Each iteration is based on the previous iteration's results. While [33] uses a single unrolled model across all iterations, [34] employs separate networks for sampling and reconstruction. Interestingly, our method could be integrated into the previous to remove the need for a predefined number of iterations K. Indeed, our method dynamically determines when to stop the acquisition process based on the target quality, allowing K to vary.

Also, we are not first trying to associate a score to the MRI reconstructed image. [35] quantifies uncertainty by resorting to the Stein's Unbiased Risk Estimator (SURE) measure. SURE studies how adding Gaussian white noise to the encoded image affects the reconstruction by running a Monte Carlo simulation. SURE quantifies uncertainty by computing second-order statistics over the reconstructed set of images. The authors show how SURE can be adopted as an indicator of the uncertainty of reconstruction. Interestingly, they also highlight how SURE is not a measure for the quality of the reconstruction. For example, authors demonstrate that training a model with an adversarial loss returns less reliable results (lower SURE) than training the same model with a standard loss, even if adversarial loss tends to improve performances in terms of reconstruction quality.

In general, the implementation of the quality estimator falls in the vast realm of No-Reference Image Quality Assessment (NR-IQA), which objective is to find the quality of a corrupted image without knowing the reference image [36,37]. Nevertheless, differently from the NR-IQA setting, some prior domain knowledge is available, even if limited. NR-IQA applied to MRI acquisition mostly focus on the design of novel metrics that better approximate the human perceptual quality [38].

In this work, we relied on the well-consolidated Mean Absolute Error (MAE), Peak Signal-to-Noise Ratio (PSNR), and Structural Similarity (SSIM) for quality assessment. Nevertheless, we are confident that alternative figures of merit, such as the one in [35,39], the ones detecting motion-related artifacts [40], or the ones from pure NR-IQA [41], can possibly be integrated with our scheme and benefit the overall acquisition.

This study explores unconventional MRI acquisition strategies with a focus on simplified scenarios, i.e., simplified MRI physics, to highlight conceptual clarity over technical complexities. Our objective is to explore future advancements in MRI acquisition, fostering inspiration, rather than proposing a singular fully-optimized model.

The rest of the paper is organized as follows. In Section 2, an overview of the MRI acquisition is presented, with a focus on related work in Section 2.1, and the presentation of our novel incremental acquisition in Section 3. Numerical results are shown in Section 4, finally in Section 5 we draw the conclusions.

2. Undersampled MRI acquisition

An MRI machine acquires *k-space samples*, which are points in the k-space: a spatial frequency domain represented as a Cartesian grid, denoted as $\boldsymbol{y} \in \mathbb{C}^{d \times d}$. These $d \times d$ samples represent the *spatial frequency coefficients* of the image. The acquired data is then transformed into an image through the Inverse Fourier Transform, i.e., $\boldsymbol{x} = \mathcal{F}^{-1}(\boldsymbol{y})$. In this context, k-space sample refers to an individual value of the k-space representing the frequency domain information.

MRI machine follows k-space trajectories, i.e., path traced in the spatial-frequency domain during data collection as determined by the applied gradients. To speed-up the acquisition of an image it is possible to *undersample*. Such strategies follow binary masks that specifies what k-space samples to acquire and what to discard [42].

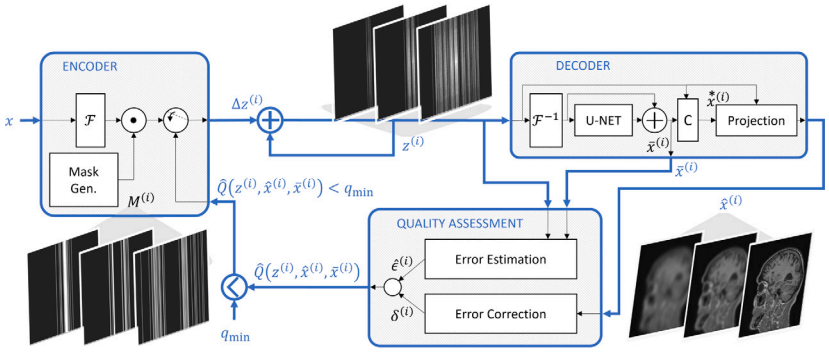

Fig. 2. Scheme of the incremental acquisition, where a fully-sampled image $\boldsymbol{x}$ is the input of an encoder block, that, at every iteration (i), acquires a portion of its k-space $\Delta \boldsymbol{z}^{(i)}$. All the samples are accumulated into a matrix $\boldsymbol{z}^{(i)}$. Such matrix is then fed to a decoder block that creates $\hat{\boldsymbol{x}}$, an image designed to be as close as possible to $\boldsymbol{x}$. The reconstruction is fed to the quality assessment block, that estimates the reconstruction quality $\hat{Q}(\boldsymbol{z}^{(i)}, \hat{\boldsymbol{x}}^{(i)}, \overline{\boldsymbol{x}}^{(i)})$. Intermediate results, such as $\overline{\boldsymbol{x}}^{(i)}$ and $\boldsymbol{z}^{(i)}$, can be exploited by the quality assessment block to improve the estimation. Finally, $\hat{Q}(\boldsymbol{z}^{(i)}, \hat{\boldsymbol{x}}^{(i)}, \overline{\boldsymbol{x}}^{(i)})$ is compared with a target quality $q_{\min}$. As long as $\hat{Q}(\boldsymbol{z}^{(i)}, \hat{\boldsymbol{x}}^{(i)}, \overline{\boldsymbol{x}}^{(i)}) < q_{\min}$, another iteration $(i+1)$ is launched.

When dealing with undersampling patterns, each sample of the k-space is usually assumed to be acquired independently of the others and in the same amount of time [11–14,19–30,43,44]. This assumption simplifies the analysis by allowing us to isolate the contribution of each k-space sample to the image reconstruction. However, it is important to recognize that this is a significant idealization. In practice, acquisition times can vary due to the chosen acquisition strategy and physical constraints. Lifting this assumption may also affect the response of the imaging system, e.g., the response to a point source (that takes the name of Point Spread Function), that is often used to describe image blurring.

In this work, we consider two types of undersampling patterns: (*i*) unconstrained sampling, which allows for a more flexible selection of k-space samples, and (*ii*) Line-Constrained sampling, where k-space is sampled along straight lines. The adopted constraint in the Line-Constrained sampling is an example of specific requirements of the imaging system. Conversely, the Unconstrained sampling does not restrict the selection to linear paths, offering more freedom in how k-space is sampled.

To formalize the undersampling procedure, image acquisition can be modeled as an encoder $\mathrm{Enc}(\cdot)$ providing $\boldsymbol{z}$ as an undersampled version of $\boldsymbol{y}$ where only $m < n = d^2$ samples of the k-space are sampled. The encoder is responsible for determining which k-space samples to acquire and which to skip, effectively reducing the total number of data points the machine had to collect. More in detail, $\mathrm{Enc}(\cdot)$ applies a binary mask $\boldsymbol{M} \in \{0,1\}^{d\times d}$ with m ones and $n-m$ zeros such that

$$\boldsymbol{z} = \mathrm{Enc}(\boldsymbol{x}) = \boldsymbol{M} \circ \boldsymbol{y} = \boldsymbol{M} \circ \mathcal{F}(\boldsymbol{x}) \tag{1}$$

where $\circ$ indicates the Hadamard element-wise product between matrices. We call $r = \frac{m}{n}$ the acquisition rate and $s = \frac{1}{r} = \frac{n}{m}$ the undersampling ratio. Note that we do not model noise in the acquisition process and we limit the analysis to the ideal case in (1), as in [22,23].

The recovery of $\boldsymbol{x}$ from $\boldsymbol{z}$ bases its effectiveness on the fact that medical images often show patterns in spatial arrangement that translate to redundancy in the k-space.

However, computing only $\mathcal{F}^{-1}(\boldsymbol{z})$ produces a first-guess reconstructed image that is affected by aliasing artifacts. Therefore, a decoder $\mathrm{Dec}(\cdot)$ exploiting the natural redundancy of the images is necessary to produce an estimate $\hat{\boldsymbol{x}} = \mathrm{Dec}(\boldsymbol{z}) \approx \boldsymbol{x}$ that is free of unwanted aliasing effect. The decoder processes the undersampled k-space to reconstruct the missing information compensating for the skipped samples.

An option to design the Enc-Dec pair is Compressed Sensing (CS), which leverages the fact that image redundancy implies sparsity in some representation [45]. Following the CS framework, $\boldsymbol{M}$ can be randomly drawn and, if the random pattern m is sufficiently large, an estimate of $\boldsymbol{x}$ is obtained by looking at the sparsest image matching (1) [46,47].

Recent data-driven approaches exploit different priors that are directly learnt from the data, i.e., $\hat{\boldsymbol{x}}$ is not imposed to be sparse but forced to be *similar* to the images observed during the training. Hence, the decoder involves a Deep Neural Network (DNN) that compensates for the aliasing effect [9,11,12,43]. The network is trained by means of a loss function

$$\mathcal{L}_{\mathrm{DNN}}(\boldsymbol{x}, \hat{\boldsymbol{x}}) = \mathbf{E}_{\boldsymbol{x}} \left[\psi \mathbf{S}(\boldsymbol{x}, \hat{\boldsymbol{x}}) + (1-\psi)\mathbf{R}(\hat{\boldsymbol{x}})\right] \tag{2}$$

where $\mathbf{E}_{\boldsymbol{x}}$ is the expectation over the training inputs, $\mathbf{S}$ is a similarity function measuring the resemblance between $\boldsymbol{x}$ and $\hat{\boldsymbol{x}}$, $\mathbf{R}$ is a generic regularization term, and $\psi \in [0,1]$ is a scalar that weights the two terms of the loss. The first term ensures similarity between true ad reconstructed images, while the second takes various forms, depending on prior knowledge, e.g., [24] explores what regularization term better imposes the measurement constraint at training time.

Data-driven approaches have also been proposed to adapt the encoder to the class of input images so that a learned mask $\boldsymbol{M}$ captures the most salient k-space samples [44]. Furthermore, it has been demonstrated that the joint optimization of $\mathrm{Enc}(\cdot)$ and $\mathrm{Dec}(\cdot)$ benefits the overall model [22,23,27–29,48].

2.1. Reference architecture

One of the latest advances among the data-driven approaches that jointly optimizes the encoder and decoder has been proposed in [24] as an extension of Learning-based Optimization of the Under-sampling Pattern (LOUPE) [22,23]. This model jointly trains the two blocks where the encoder and the decoder respectively depend on trainable parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\theta}$, i.e., $\hat{\boldsymbol{x}} = \mathrm{Dec}_{\boldsymbol{\theta}}(\mathrm{Enc}_{\boldsymbol{\gamma}}(\boldsymbol{x}))$, where $\cdot_{\boldsymbol{\gamma}}$ and $\cdot_{\boldsymbol{\theta}}$ indicate the dependency on $\boldsymbol{\gamma}$ and $\boldsymbol{\theta}$.

According to this framework, the model follows the *autoencoder* paradigm: $\mathcal{A}_{\boldsymbol{\gamma},\boldsymbol{\theta}}(\cdot) = \mathrm{Dec}_{\boldsymbol{\theta}}(\mathrm{Enc}_{\boldsymbol{\gamma}}(\cdot))$ having as peculiarity an encoder that performs undersampling, while the objective of the decoder is to restore the original image, as depicted in Fig. 2.

At training time, instead of a binary mask, $\mathrm{Enc}_{\boldsymbol{\gamma}}(\cdot)$ uses $\boldsymbol{M}_{\boldsymbol{\gamma}} \in [0,1]^{d\times d}$ to promote gradient back-propagation. $\boldsymbol{M}_{\boldsymbol{\gamma}}$ is computed from the output of an auxiliary layer that sets the probability of each element of $\boldsymbol{y}$ to be included in $\boldsymbol{z}$. During inference, the probability mask controlled by $\boldsymbol{\gamma}$ is thresholded to obtain the final binary mask. The ad-hoc layer controlling $\boldsymbol{M}_{\boldsymbol{\gamma}}$ sets a desired value of s [23,24].

As depicted in Fig. 2, the decoder $\mathrm{Dec}_{\boldsymbol{\theta}}(\cdot)$ is a cascade of blocks producing two intermediate results and the final reconstruction. The first intermediate image is computed by means of the residual UNET [9], that returns $\overline{\boldsymbol{x}} \in \mathbb{C}^{d\times d}$

$$\overline{\boldsymbol{x}} = \mathcal{F}^{-1}[\boldsymbol{z}] + \mathrm{UNET}_{\boldsymbol{\theta}}(\mathcal{F}^{-1}[\boldsymbol{z}]) \tag{3}$$

As z is an undersampling of the k-space, $\mathcal{F}^{-1}(z)$ corresponds to the original image x corrupted by aliasing. The UNET aims to correct the

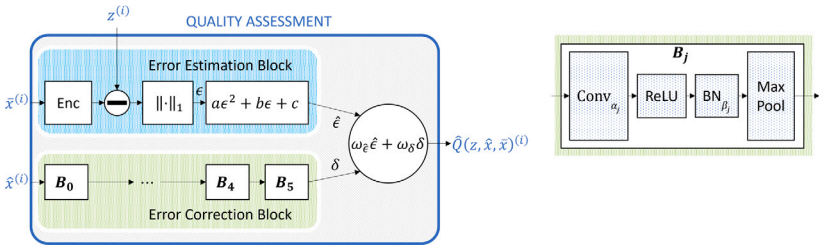

Fig. 3. Scheme of the quality assessment architecture, where two main blocks are implemented. The error estimation block computes a first quality guess by computing the measurement error on $\bar{x}$ (an intermediate output of the decoder architecture we adopted in this work) and fitting it to quadratic polynomial. The Error Correction Block adjusts the first guess estimation, with a cascade of six blocks $\mathbf{B}_0, \mathbf{B}_1, \ldots, \mathbf{B}_5$ comprising Convolutional layers that extract features from the reconstructed image $\hat{x}$. The two contributions are weighted together by means of a Dense layer. All the network parameters are jointly optimized at training time.

aliasing artifacts allowing $\bar{x}$ to be a better reconstruction compared to $\mathcal{F}^{-}1(z)$.

However, $\bar{x}$ may not be consistent with the measurements z as demonstrated in [24]. For this reason, an ad-hoc block forces the match between z and $\text{Enc}_{\gamma}(\bar{x})$ by projecting $\bar{x}$ into the space of the image consistent with the measurement (referred as "block C" in Fig. 2), thus producing a second intermediate reconstruction $\overset{\star}{x} \in \mathbb{R}^{d \times d}$. In details:

$$\overset{\star}{x} = \mathcal{F}^{-1}\left(\mathcal{F}(\bar{x}) \circ \left(\mathbf{1} - \boldsymbol{M}_{\gamma}\right) + z\right) \tag{4}$$

Finally, $\hat{x}$ is obtained by applying Dykstra's projection algorithm [24] (referred as "Projection" in Fig. 2), that iteratively projects the image into two convex sets: one representing the space of images consistent with the measured k-space (the measurement constraint $z = \text{Enc}_{\gamma}(\hat{x})$), and the other representing the real-valued image space (the hypercube $[0, 1]^{d \times d}$). Applying the Dykstra's projection algorithm ensures that the final image satisfies both constraints.

3. Incremental MRI acquisition

In this section, we evolve the classic acquisition procedure by presenting the novel *Incremental MRI Acquisition*.

For $s > 1$, $\mathcal{A}_{\gamma,\theta}$ inevitably introduces errors, causing quality degradation according to a given quality function $Q(\boldsymbol{x}, \hat{\boldsymbol{x}})$. Given $\boldsymbol{x}$ it is reasonable to say that there is a monotonic decreasing mapping between s and $Q(\boldsymbol{x}, \hat{\boldsymbol{x}})$.

One option is to find a maximum $\tilde{s}$ such that $Q(\boldsymbol{x}, \hat{\boldsymbol{x}}) \geq q_{\min}$ for all the $\boldsymbol{x} \in \mathcal{X}$, where $\mathcal{X}$ represents the set of possible input instances. Nevertheless, it could be convenient not to adopt a fixed $\tilde{s}$ but to identify a procedure that for any image $\boldsymbol{x}$ returns the highest s matching the minimum quality of service. To make this possible, we consider different undersampling ratios $S = \{s^{(1)}, \ldots, s^{(N)}\}$ so that every $s^{(i)}$ corresponds a different autoencoder $\mathcal{A}^{(i)}_{\gamma,\theta}$ and undersampling mask $\boldsymbol{M}^{(i)}$. Then, we provide a procedure that, for each incoming image, selects the highest undersampling ratio matching the quality constraint.

The problem of finding the highest $s \in S$ respecting the constraint $Q(\boldsymbol{x}, \hat{\boldsymbol{x}}) \geq q_{\min}$, for an image $\boldsymbol{x}$, is formulated as an optimization problem whose solution is referred to as *oracle-based*

$$\begin{aligned} s^* &= \underset{s \in S}{\text{argmax}}\, s \\ &\text{subject to} \quad Q(\boldsymbol{x}, \hat{\boldsymbol{x}}) \geq q_{\min} \end{aligned} \tag{5}$$

The above problem cannot be solved at inference time because $\boldsymbol{x}$ is unknown, hence we rewrite (5) as

$$\begin{aligned} s^* &= \underset{s \in S}{\text{argmax}}\, s \\ &\text{subject to} \quad \hat{Q}(\boldsymbol{z}, \hat{\boldsymbol{x}}, \ldots) \geq q_{\min} \end{aligned} \tag{6}$$

where $\hat{Q}(\boldsymbol{z}, \hat{\boldsymbol{x}}, \ldots)$, namely the *Quality Assessment*, is a proxy of $Q(\boldsymbol{x}, \hat{\boldsymbol{x}})$. $\hat{Q}(\boldsymbol{z}, \hat{\boldsymbol{x}}, \ldots)$ leverages available information, such as $\boldsymbol{z}$ and $\hat{\boldsymbol{x}}$, to produce an estimate of the image quality. It can also include other inputs, such as $\bar{\boldsymbol{x}}$ from (3).

Let us consider the elements of S in descending order so that $s^{(i)} > s^{(i+1)}$ and to build the sequence of masks $\{\boldsymbol{M}^{(1)}, \ldots, \boldsymbol{M}^{(N)}\}$ such that they respect the incremental constraint $\boldsymbol{M}^{(i)} - \boldsymbol{M}^{(i-1)} \in \{0, 1\}^{d \times d}$, i.e., any mask $\boldsymbol{M}^{(i)}$ selects all k-space samples selected by the previous mask $\boldsymbol{M}^{(i-1)}$ plus a batch of new points. Our proposal is to tackle (6) with an incremental approach where we use all $\mathcal{A}^{(i)}_{\gamma,\theta}(\boldsymbol{x})$ associated to $s^{(i)} \geq s^*$. The algorithm implementing this procedure is sketched in Fig. 2 and reported in Alg. 1 where $\mathbf{0}$ defines the null matrix, $\boldsymbol{M}^{(0)} = \mathbf{0}$, and get_sample$(\cdot)$ is the MRI sampling operation, modeled by (1) at design time, that takes the undersampling mask as an argument.

Algorithm 1 Incremental Adaptive Acquisition (slice)

Require: $q_{\min}$: a minimum quality level
Require: $\{\boldsymbol{M}^{(1)}, \ldots, \boldsymbol{M}^{(N)}\}$: a sequence of incremental masks with corresponding undersampling ratio $\{s^{(1)}, \ldots, s^{(N)}\}$, with $\boldsymbol{M}^{(0)} = \mathbf{0}$
  $i \leftarrow 1$
  $\hat{q} \leftarrow 0$
  $\boldsymbol{z} \leftarrow \mathbf{0}$
  while $\hat{q} < q_{\min}$ **do**
    $\Delta \boldsymbol{z}^{(i)} \leftarrow$ get_sample$(\boldsymbol{M}^{(i)} - \boldsymbol{M}^{(i-1)})$
    $\boldsymbol{z} \leftarrow \boldsymbol{z} + \Delta \boldsymbol{z}^{(i)}$
    $\hat{\boldsymbol{x}}, \bar{\boldsymbol{x}} \leftarrow \text{Dec}(\boldsymbol{z})$
    $\hat{q} \leftarrow \hat{Q}(\boldsymbol{z}, \hat{\boldsymbol{x}}, \bar{\boldsymbol{x}})$
    $i \leftarrow i + 1$
  end while
  $s^* = s^{(i)}$
  return $s^*, \hat{\boldsymbol{x}}$

With respect to the encoder in [24], we here need to force the incremental constraint on the sequence of binary masks characterizing the sequences of autoencoders. We build the set $\{\boldsymbol{M}^{(1)}, \ldots, \boldsymbol{M}^{(N)}\}$ of incremental masks by training the corresponding $\{\mathcal{A}^{(1)}_{\gamma,\theta}, \ldots, \mathcal{A}^{(N)}_{\gamma,\theta}\}$ in a decremental fashion,[1] i.e., starting from $i = N$ (corresponding to the lowest undersampling ratio s) till $i = 1$ (the highest s). Note that every $\boldsymbol{M}^{(i)}$ is obtained starting from $\boldsymbol{M}^{(i+1)}$ by means of a layer embedded in the encoder that forces the ones in $\boldsymbol{M}^{(i)}$ to be a subset of the ones in $\boldsymbol{M}^{(i+1)}$.

To make our approach effective, $\hat{Q}(\cdots)$ must be able to guess the reconstruction quality at each step of the incremental acquisition scheme. Inspired by the latest advancement in NR-IQA [49,50], we implement the $\hat{Q}(\cdots)$ function with a DNN. Nevertheless, we diverge from standard NR-IQA by leveraging other information than $\hat{\boldsymbol{x}}$. Our architecture develops in two blocks, namely the *Error Estimation* block and the *Error Correction* block, whose scalar outputs are combined to produce the

[1] Numerical evidence show that training in a decremental fashion leads to better performance than following the incremental ordering.

estimated image quality. The resulting block scheme of the incremental acquisition, including the neural quality predictor, is sketched in Fig. 2.

The first block takes the error $\epsilon = \left\|\mathrm{Enc}\left(\overline{\boldsymbol{x}}\right) - \boldsymbol{z}\right\|_1$ where $\|\cdot\|_1$ is the standard ℓ_1 norm, and feeds it to a quadratic function to produce $\hat{\epsilon} = a\epsilon^2 + b\epsilon + c$. As already demonstrated in [24,51], $\hat{\epsilon}$ is a first rough proxy of the reconstruction quality. The Error Correction Block extracts

$$\boldsymbol{\gamma} = \mathrm{gpool}\left(\mathbf{B}_5\left(\mathbf{B}_4\left(\mathbf{B}_3\left(\mathbf{B}_2\left(\mathbf{B}_1\left(\mathbf{B}_0\left(\hat{\boldsymbol{x}}\right)\right)\right)\right)\right)\right)\right) \tag{7}$$

with gpool implementing a global average pooling that returns a scalar feature $\delta \in \mathbb{R}$; and each $\mathbf{B}_j$ implementing:

$$\mathbf{B}_j(\cdot) = \mathrm{pool}\left(\mathrm{BN}_{\beta_j}\left(\mathrm{ReLu}(\mathrm{Conv}_{\alpha_j}(\cdot))\right)\right) \tag{8}$$

where j is the index of the block; pool is the average pooling operation with window of 2×2; BN_β is the batch normalization operation with parameter β; ReLu is the Rectified Linear unit operation; and Conv_α is the convolution operation with kernel 5×5, parameterized by α. The number of filters of each Conv_{α_j} is $[10, 8, 6, 4, 2, 1]$.

Finally, the two branches are concatenated and fed to a Dense layer, parameterized by $\omega_{\hat{\epsilon}}$ and ω_γ that outputs the estimated quality:

$$\hat{\boldsymbol{Q}}(\boldsymbol{z}, \hat{\boldsymbol{x}}, \overline{\boldsymbol{x}}) = \omega_{\hat{\epsilon}}\hat{\epsilon} + \omega_\gamma \boldsymbol{\gamma} \tag{9}$$

All the parameters $a, b, c, \alpha_i, \beta_i, \omega_{\hat{\epsilon}}, \omega_\gamma$ are trainable and tuned at training time by minimizing the loss:

$$L_Q(a, b, c, \alpha, \beta, \omega_{\hat{\epsilon}}, \omega_\gamma) = \mathbf{E}_{\boldsymbol{x}}\left[\left\|Q(\boldsymbol{x}, \hat{\boldsymbol{x}}) - \hat{Q}(\boldsymbol{z}, \hat{\boldsymbol{x}}, \overline{\boldsymbol{x}})\right\|_1\right] \tag{10}$$

Since $\boldsymbol{z}$ and thus $\hat{\boldsymbol{x}}$ and $\overline{\boldsymbol{x}}$ depend on the mask $\boldsymbol{M}$, a quality estimator $\hat{Q}$ must be trained for each $\boldsymbol{M}^{(i)}$ in our approach.

A detailed visual representation of the quality assessment block is given in Fig. 3.

4. Numerical evidences

4.1. Dataset

We evaluate our method by adopting two publicly available datasets: NYU fastMRI dataset [52] and IXI dataset [53]. The first contains fully-sampled knee scans, the second contains fully-sampled brain scans. According to the block scheme in Fig. 2, we consider directly the reconstructed images provided by these datasets, rather than raw k-space. Here, we compute k-spaces by applying the Fourier transform to each image in every volume for both use cases considered.

In this work, we represent k-space as a 2D Cartesian grid. We evaluated our method using two sampling strategies: (*i*) Line-Constrained, which follow straight-line sampling patterns. This align with real 2D imaging, e.g., Spin Echo sequence with frequency encoding dimension on the x-axis and with the phase encoding direction on the y-axis of our k-space representation; (*ii*) Unconstrained, which offers more flexibility and allows for pure exploration of the potentialities of the model. While theoretically feasible by leveraging one frequency encoding dimension and two phase encoding dimensions, the deployability of our incremental acquisition for 3D imaging requires more investigations.

By combining datasets and sampling strategies we create four different use cases, referenced as (*i*) `FastMRI-U`; (*ii*) `FastMRI-C`; (*iii*) `IXI-U`; (*iv*) `IXI-C`, where `U` and `C` refer to Unconstrained and Line-Constrained, respectively. We work with the 320×320 images in FastMRI emulated single-coil proton density-weighted data. Similarly, we extract the 256×256 single-coil T1 data from the IXI dataset.[2]

Both datasets are naturally grouped in *volumes*, each volume being composed of grayscale slices. To address variations in intensity across independently acquired volumes, we normalize each volume by dividing pixel values by the maximum magnitude within that volume. This approach ensures consistent contrast levels between different patients without introducing distortions. FastMRI images are split into a training, validation, and test set comprising 50, 5, and 5 volumes, respectively. Each volume contains a number of slices in the range $[34, 45]$. IXI images are similarly split into sets comprising 166, 11, and 8 volumes. Each volume contains 150 slices.

`IXI-U` (`IXI-C`) images are similarly split into sets comprising 166 (166), 11 (12), and 8 (7) volumes. Each volume contains 150 slices. The additional validation volume of `IXI-C` was employed to estimate and compensate an unexpected bias error in the quality assessment.

Within the same volume, initial and final slices are characterized by a scarcity of anatomical features and the information content tend to accumulate in the center. Fig. 4 depicts slices $5, 15, 25, 35$ $(30, 60, 90, 130)$ of a volume belonging to the FastMRI (IXI) dataset, confirming such trend.

4.2. Training parameters

We use the Adam optimizer with an initial learning rate of 0.01, with an Early Stop and Reduce Learning Rate on Plateau to secure convergence. Since training produces probabilities with which a pixel or a line of the mask is non-zero, testing is performed by first drawing a binary mask according to such probabilities and keeping it fixed throughout the assessment.

The autoencoder network has approximately 31.5M parameters. The N Self-Assessment models, one for each $s \in \mathcal{S}$, have approximately 5.5 K total parameters each. We train every $\mathcal{A}^{(i)}_{\gamma,\theta}$ for l different weight initialization, then select the one having the highest average reconstruction quality. We set $l = 5$ for fastMRI and $l = 2$ for IXI.

For the two use cases `FastMRI-U` and `FastMRI-C` we adopt,[3]

$$\mathcal{S} = \{64, 32, 24, 20, 16, 12, 10, 8, 6, 5, 4, 3\}.$$

For `IXI-U` we consider

$$\mathcal{S} = \{64, 32, 24, 20, 16, 12, 10, 9, 8, 7.5, 7, 6.75, 6.5, 6.25, 6, 5.75, 5.5, 5.25, 5, 4.75, 4.5, 4.25, 4, 3\}.$$

while `IXI-C` refers to

$$\mathcal{S} = \{64, 32, 24, 20, 16, 12, 10, 9, 8, 7, 6.5, 6, 5.75, 5.5, 5.25, 5, 4.75, 4.5, 4.25, 4, 3.75, 3.5, 3.25, 3\}.$$

In a real use case, the choice of $q_{\min}$ is left to domain experts. For the sole purpose of testing, we tuned parameters $\mathcal{S}$ so that our incremental acquisition is competitive in the range $q_{\min} \in [38, 44]$. To modify this range of efficacy, it is sufficient to change $\mathcal{S}$.

4.3. Evaluation metrics

To assess the quality of each reconstruction we adopt three different metrics:

- Peak Signal-to-Noise Ratio:
$$\mathrm{PSNR}(\boldsymbol{x}, \hat{\boldsymbol{x}}) = 10\log_{10}\left(\frac{\max(\boldsymbol{x})^2}{\mathrm{MSE}}(\boldsymbol{x}, \hat{\boldsymbol{x}})\right) \tag{11}$$
where $\mathrm{MSE}(\boldsymbol{x}, \hat{\boldsymbol{x}}) = 1/n \sum_{j=0}^{n-1}([\boldsymbol{x}]_j - [\hat{\boldsymbol{x}}]_j)^2$ is the Mean Square Error between $\boldsymbol{x}$ and the decoder output and $[\boldsymbol{x}]_j$ refers to the jth pixel of the flattened matrix. Notice that, $\max(\boldsymbol{x}) = 1$.
- Inverse dB Mean Absolute Error:
$$\mathrm{MAE}^{-1}_{\mathrm{dB}}(\boldsymbol{x}, \hat{\boldsymbol{x}}) = 20\log_{10}\left(\frac{\max(\boldsymbol{x})}{\mathrm{MAE}(\boldsymbol{x}, \hat{\boldsymbol{x}})}\right) \tag{12}$$
where $\mathrm{MAE}(\boldsymbol{x}, \hat{\boldsymbol{x}}) = 1/n \sum_{j=0}^{n-1} |[\boldsymbol{x}]_j - [\hat{\boldsymbol{x}}]_j|$ is the Mean Absolute Error between $\boldsymbol{x}$ and the decoder output.

[2] Note that the proposed acquisition strategy can be extended to the case of multi-coil datasets and we plan to address this in future communications.

[3] For `FastMRI-C`, the set used in the final test is reduced to $\mathcal{S} = \{16, 12, 10, 8, 6, 5, 4, 3\}$, as the quality estimator failed to converge for the other 4 undersampling ratio.

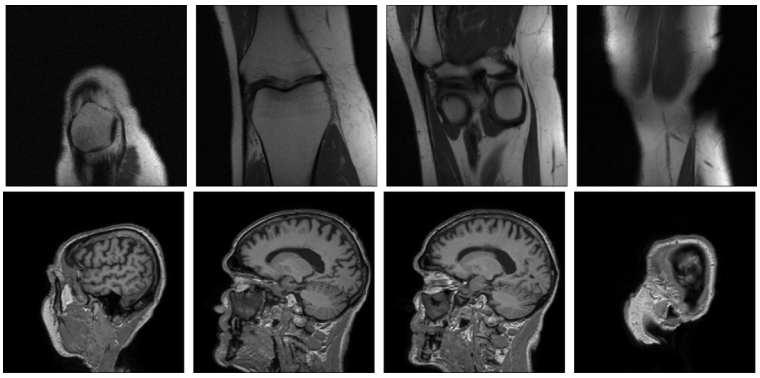

Fig. 4. Some slices from the first volume of both FastMRI and IXI datasets. Slice index for FastMRI are, in order, [5, 15, 25, 35], while for IXI slice indexes are [30, 60, 90, 130]. These sequences show how a volume is made by images of different information content. For improved visualization, all images have been linearly rescaled to enhance visual contrast.

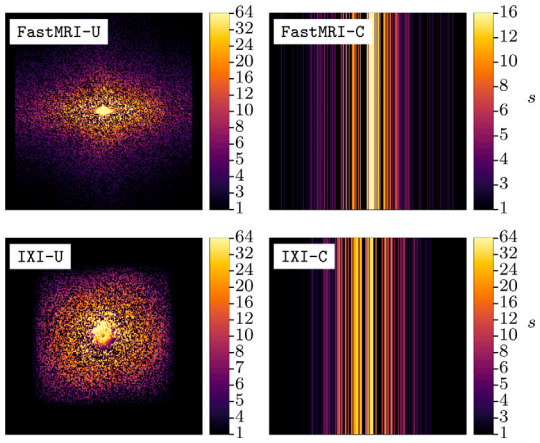

Fig. 5. Sampling pattern for the four use-cases. Each pixel of the image is a sample of the k-space, that is painted proportionally with the undersampling ratio s at which it is acquired. The higher the associated undersampling ratio (the sooner it is acquired), the brighter the pixel.

- Structural Similarity Index:

$$\mathrm{SSIM}(\boldsymbol{x}, \hat{\boldsymbol{x}}) = \frac{(2\mu_{\boldsymbol{x}}\mu_{\hat{\boldsymbol{x}}} + c_1)(2\sigma_{\boldsymbol{x}\hat{\boldsymbol{x}}} + c_2)}{(\mu_{\boldsymbol{x}}^2 + \mu_{\hat{\boldsymbol{x}}}^2 + c_1)(\sigma_{\boldsymbol{x}}^2 + \sigma_{\hat{\boldsymbol{x}}}^2 + c_2)} \tag{13}$$

where $\mu_{\boldsymbol{x}}$ and $\mu_{\hat{\boldsymbol{x}}}$ are defined as *local* averages for the original and reconstructed images, $\sigma_{\boldsymbol{x}}$ and $\sigma_{\hat{\boldsymbol{x}}}$ are the *local* variances, while $\sigma_{\boldsymbol{x}\hat{\boldsymbol{x}}}$ is the *local* covariance between the decoder output and the ground-truth. *local* means computed on sliding windows with size 11×11 where each pixel is weighted with a Gaussian weight with standard deviation 1.5 pixels. In the above formula, $c_1 = (k_1 D)^2$ and $c_2 = (k_2 D)^2$, where D is the dynamic range of the pixel values, $k_1 = 0.01$ and $k_2 = 0.03$ as suggested in [54].

4.4. Introduction to experiments

To demonstrate the effectiveness of our proposed method, we conducted a series of experiments designed to evaluate different aspects of the acquisition and reconstruction process. The experiments revolve around three main axis:

- Plausibility: we display all the sampling patterns (acquisition masks) in a single figure, with patterns corresponding to higher undersampling ratios (s). We also display some real reconstructed images.
- Quantify the Advantages: we show the improvements in terms of acceleration achieved by our method for different reconstruction quality targets. We explore such advantages both in terms of reconstruction quality and undersampling ratios for a fixed acquisition target.
- Quality Estimation Analysis: we assess the performances of the quality estimator and the validity of the chosen metric for quality estimation.

In this section, we evaluate three distinct approaches to MRI acquisition:

- Oracle-based approach: This approach assumes perfect knowledge of the target quality for the reconstructed image and selects the optimal undersampling ratio s to meet the specified quality. The values of s are derived according to the solution of Eq. (5). This serves as a theoretical upper bound for comparison.
- Incremental approach: In this method, the undersampling ratio is progressively increased based on real-time quality estimates, as outlined in Algorithm 1. This approach reflects a more practical and adaptive acquisition process that adjusts based on the slice being acquired.
- Classic acquisition approach: This strategy uses a single, fixed undersampling mask with a predetermined value $\tilde{s}$. In this method, the undersampling ratio does not change during the acquisition, providing a baseline for comparison.

For consistent comparison across all methods, we define the minimum quality threshold, $q_{\min}$, as the lowest acceptable reconstruction accuracy observed over the test set. The oracle-based approach sets the benchmark, while the incremental approach represents a realistic, adaptive solution that dynamically adjusts the acquisition process.

4.5. Results

In this section, we present the experimental results comparing the performance of the proposed methods across different acquisition strategies. The results are discussed in terms of undersampling ratios, reconstruction quality, and quality estimation accuracy. Below, we detail key observations from our experiments.

Fig. 5 associates to every frequency in $\boldsymbol{y}$ its corresponding s, for all the four use cases. Frequencies associated with higher s are acquired before frequencies associated with lower s. We notice that central frequencies are always acquired first. Also, drastic structure changes between use cases become evident, highlighting the adaptability of the model to different classes of signal.

Fig. 5 illustrates the relationship between the frequency in y and its corresponding undersampling ratio s for all four use cases. Frequencies associated with higher s values are acquired first, with central frequencies always being prioritized. This reflects the adaptability of the model to different signal classes, as evidenced by the structural changes between use cases.

Results in terms of s, as a function of $q_{\min}$ $\left[\mathrm{MAE}_{dB}^{-1}\right]$, are shown in Fig. 6. For all approaches, $q_{\min}$ is monotonic with respect to s. The

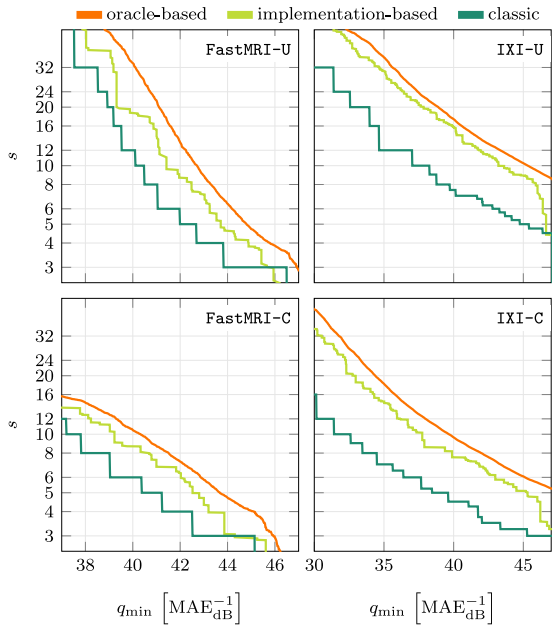

Fig. 6. $\bar{s}$ (average s over the test set) for incremental acquisition and $\tilde{s}$ for the classic acquisition, are plotted for many different $q_{\min}$. $\bar{s}$ is displayed for both the oracle-based and incremental acquisition; and for the four selected use cases.

Table 1
$\bar{s}$ for incremental and classic acquisitions, for different $q_{\min}$, for all the use cases.

$q_{\min}$	$\bar{s}$: incremental/classic			
$[\mathrm{MAE}_{\mathrm{dB}}^{-1}]$	FastMRI-U	FastMRI-C	IXI-U	IXI-C
38	49.4 / 32.0	12.6 / 8.0	20.1 / 10.00	9.3 / 5.25
39	39.0 / 20.0	11.2 / 8.0	17.4 / 8.00	8.6 / 5.00
40	19.6 / 12.0	8.7 / 6.0	15.8 / 7.50	8.2 / 4.50
41	15.4 / 8.0	7.4 / 5.0	13.0 / 7.00	7.4 / 4.50
42	9.1 / 5.0	6.3 / 4.0	12.0 / 6.75	6.8 / 3.75
43	7.1 / 4.0	4.7 / 3.0	10.9 / 6.00	6.2 / 3.50
44	4.7 / 3.0	3.9 / 3.0	9.7 / 5.50	5.6 / 3.25

reported undersampling ratio values are computed as the inverse of the average of the acquisition rates r over the entire test set:

$$\bar{s} = \left(\frac{1}{L} \sum_{j=0}^{L-1} r_j^* \right)^{-1} \tag{14}$$

with L representing the magnitude of the test set and $r^* = 1/s^*$.

Reported results evidence that, fixed $q_{\min}$, an undisputed superiority of the incremental acquisition in terms of undersampling ratio can be seen for all four use cases. The distance between the incremental and the oracle-based curves is a first index of the quality assessment block performances. Because $\hat{Q}(\boldsymbol{z}, \hat{\boldsymbol{x}}, \overline{\boldsymbol{x}})$ inevitably introduces errors, the incremental curve tends to be lower than the oracle-based curve.

Table 1 compares, for all the use cases, the incremental approach $\bar{s}$ with the classic approach $\tilde{s}$, for some fixed reference $q_{\min}$. These numerical results give a more precise idea of the magnitude of the improvement that the incremental acquisition brings. As expected, the table also highlights how the improvements for the Line-Constrained cases are smaller than the improvements for the Unconstrained cases. It is reasonable to link this performance drop with the fewer degrees of freedom that the Line-Constrained trajectories leave to the encoder.

Fig. 7 displays the reconstruction quality Q for the acquisition of 5 volumes of the FastMRI and IXI test set with fixed $q_{\min} = 40$ dB. On the x-axis are the indexes of every slice of the 5 volumes while on the y-axis is the reconstruction quality in terms of $\mathrm{MAE}_{\mathrm{dB}}^{-1}$. The periodicity reflects the anatomical volume structure, as observable in Fig. 4. As the slices move through the volume, they transition from less detailed regions that are simpler to reconstruct (i.e., higher PSNR) to more complex areas (i.e, lower PSNR), and back to simpler regions. The concatenation of 5 volumes results in pseudo-periodic variations in PSNR. This effect is more pronounced in the classic acquisition method, which applies a fixed undersampling ratio across all slices. Conversely, the incremental acquisition curve adapts s to the slice complexity and tends to have a PSNR closer to $q_{\min}$.

The classic and the incremental acquisition curves distance themselves more in the volume tails than in the volume centers. This demonstrates the ability of the incremental acquisition to adapt to the information content of slices, in fact, higher s are adopted when the information content is poor (at the extremities of the volumes). Due to the error in the assessment of the achieved quality of reconstruction, the incremental tends to over-sample the central slices, hence diverging from the optimal oracle-based trend, which is well spread just above the target quality line.

To further stress this aspect, Fig. 8 visually compares the resulting undersampling ratios s of the three approaches for $q_{\min} = 40$ dB. Here, the curve associated with the classic acquisition becomes a horizontal line ($\tilde{s}$), while the curves associated with the adopted approaches display an accentuated step-like trend. The incremental is always below the oracle-based, and only sporadically below the classic acquisition horizontal line. Once again, this highlights the adaptability and superiority of the incremental acquisition. The figure also reports the average undersampling ratios over the whole test set for both adapted solutions with two dashed lines.

We report results in terms of $\mathrm{MAE}_{\mathrm{dB}}^{-1}$ because MAE characterizes the loss function used to train $\mathcal{A}_{\gamma,\theta}$ (as in [24]). Further to that, we discuss the mapping between $\mathrm{MAE}_{\mathrm{dB}}^{-1}$ and the other two metrics presented in Section 4.3. The scatter plot in Fig. 9 associates, for every slice in the test set, the reconstruction quality in terms of $\mathrm{MAE}_{\mathrm{dB}}^{-1}$ with the reconstruction quality in terms of both PSNR (first row) and SSIM (second row). Such comparison is done for images acquired with three different target qualities (represented with a different color).

Interestingly enough, these results confirm that the incremental acquisition, by forcing a target $q_{\min}$ in terms of $\mathrm{MAE}_{\mathrm{dB}}^{-1}$, influences the reconstructed slices in terms of both PSNR and SSIM. We also observe that, in case of FastMRI, there is a strong connection between the $\mathrm{MAE}_{\mathrm{dB}}^{-1}$ and the other two metrics.

Finally, we investigate the performances of the quality assessment $\hat{Q}(\boldsymbol{z}, \hat{\boldsymbol{x}}, \overline{\boldsymbol{x}})$ when estimating $\mathrm{MAE}_{\mathrm{dB}}^{-1}$. A figure of merit is obtained by computing the absolute error, i.e., we compute $\zeta = \frac{1}{L} \sum_{j=0}^{L-1} |Q(\boldsymbol{x}_j, \hat{\boldsymbol{x}}_j) - \hat{Q}(\boldsymbol{z}_j, \hat{\boldsymbol{x}}_j, \overline{\boldsymbol{x}}_j)|$. The results for all the quality assessment models are displayed in Fig. 10. The plots show, in general, a growing difficulty in the quality estimation with the growth of s.

The encoder–decoder network processes images in 4.3 ms (FastMRI) and 2.7 ms (IXI) on average using an NVIDIA A100 GPU. The quality estimator adds 0.78 ms (FastMRI) and 0.36 ms (IXI), resulting in total processing times of 5.1 ms and 3.1 ms, respectively. In terms of Multiply and Accumulate (MAC) operations, it requires 85 M MAC and 54 M MAC to estimate the quality of a Fast MRI and IXI image respectively.

Before conclusion, in Fig. 11 we visually compare some slices acquired with classic and incremental acquisitions, where $q_{min} = 43$ dB. Every image also reports (in red) the undersampling ratio s and the real reconstruction quality. As expected, classic-acquired slices share the same s. Conversely, the acceleration s in incrementally-acquired slices varies depending on the information content. We observe that when a sufficiently high $q_{\min}$ is chosen, incremental acquisition enables faster acquisitions at the cost of a negligible quality degradation. In fact, as observed in the figure, no degradation is perceivable even when the qualities differ more than 3 dB and similar behavior is observed in the rest of the dataset. Fig. 12 further proves the analysis. In this figure, we visualize the absolute difference between the reconstructed image and the ground truth, where the color scale (from black to pink to white) highlights the areas of minimal to maximal differences, respectively. In fact, incremental acquisition introduces no significant perceivable degradation, with differences being minimal even when reconstruction quality varies by more than 3 dB.

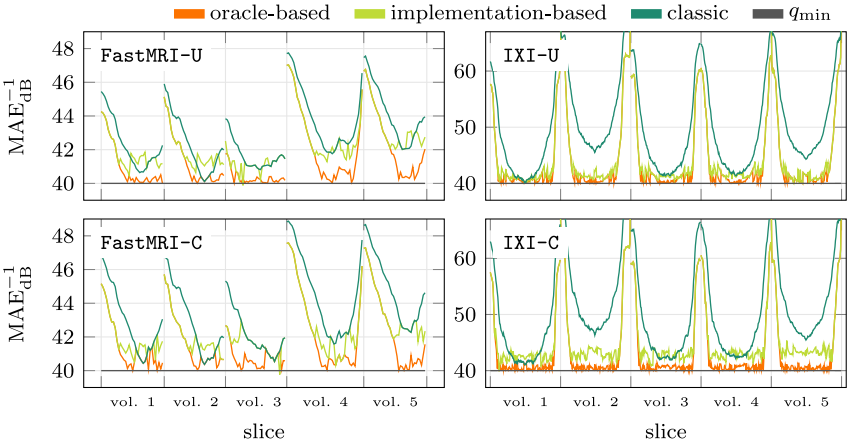

Fig. 7. Quality of the slices with classic, incremental, and oracle-based incremental acquisitions, given a target quality $q_{\min} = 40\mathrm{MAE}_{\mathrm{dB}}^{-1}$ (solid horizontal line) for all use cases FastMRI-U, FastMRI-C, IXI-U, IXI-C. FastMRI displays all the volumes of the test set, while IXI displays a subset of five volumes.

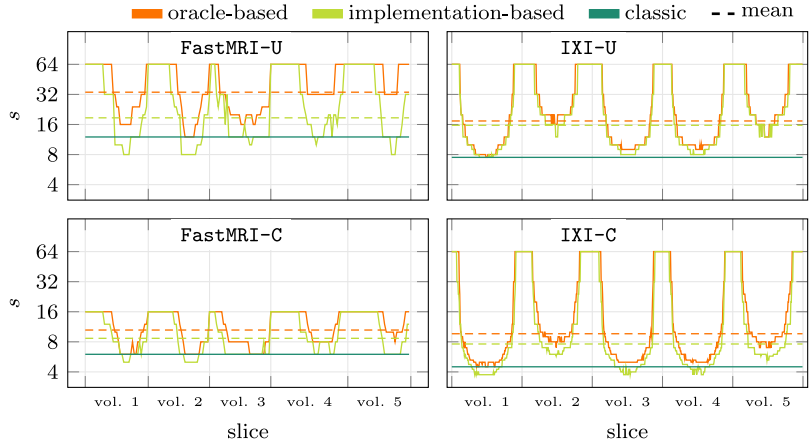

Fig. 8. Undersampling ratio s of the slices with classic, incremental and oracle-based incremental acquisitions, given a target quality $q_{\min} = 40\mathrm{MAE}_{\mathrm{dB}}^{-1}$ for all use cases FastMRI-U, FastMRI-C, IXI-U, IXI-C. The average s over the whole test set is reported with dashed lines. FastMRI displays all the volumes of the test set, while IXI displays a subset of five volumes.

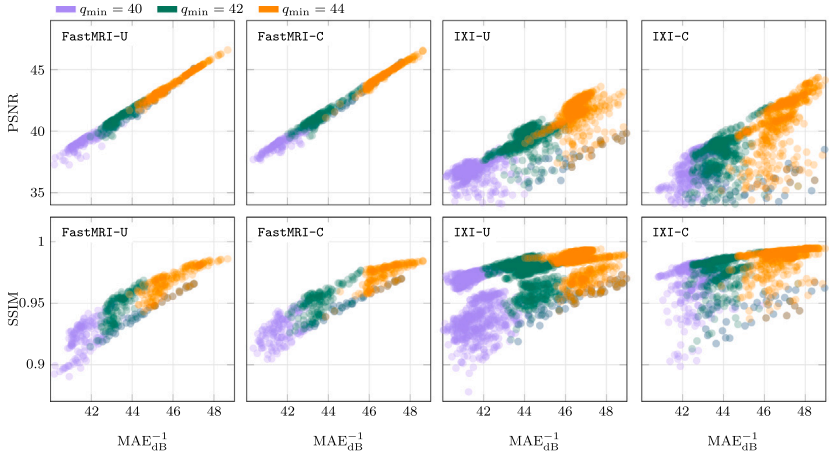

Fig. 9. PSNR and SSIM as a function of $\mathrm{MAE}_{\mathrm{dB}}^{-1}$, for every sample of the test set, for incremental acquisition, for three different target qualities.

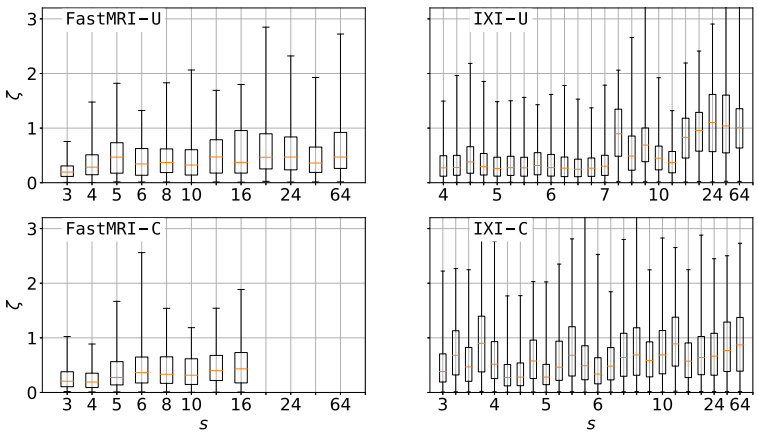

Fig. 10. Estimation error in terms of ζ or all the implemented quality estimators. The central line within each box represents the median, while the edges of the box correspond to the 1st and 3rd quartiles. Whiskers extend to the 1st and 99th percentiles.

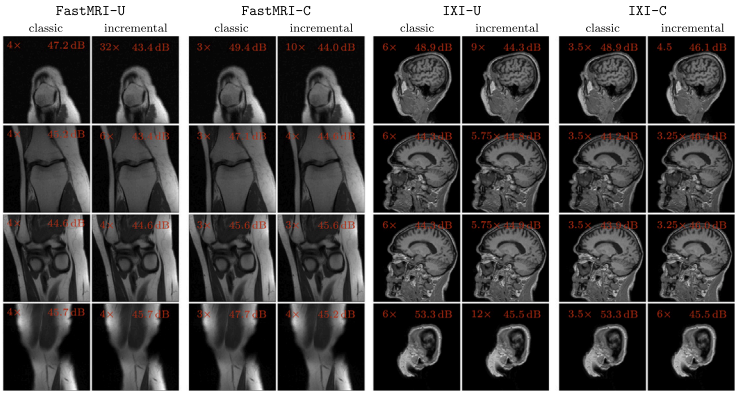

Fig. 11. Classically and incrementally acquired test slices with indexes $[5, 15, 25, 35]$ for the FastMRI dataset, and $[30, 60, 90, 130]$ for the IXI dataset. In red, the acquisition undersampling ratio s and the real reconstruction quality $Q(\boldsymbol{x}, \hat{\boldsymbol{x}})$[dB]. For improved visualization, all images have been linearly rescaled to enhance visual contrast.

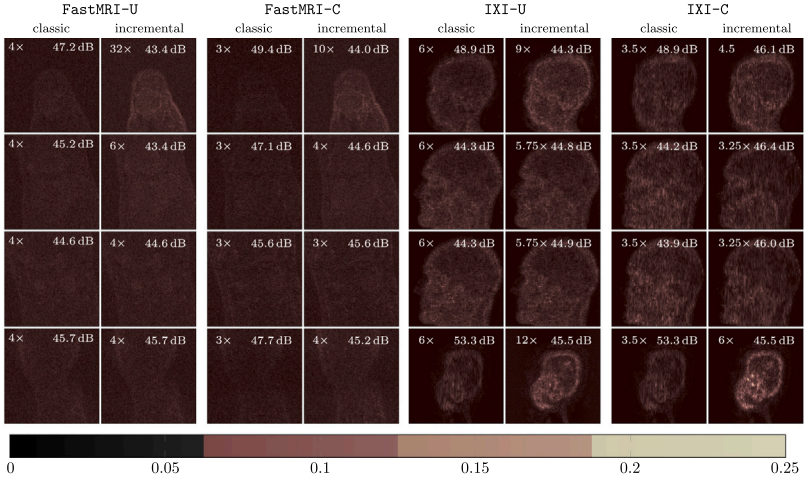

Fig. 12. Absolute difference between classically and incrementally acquired slices w.r.t. the ground truth slices. We use slices from the test set which indexes are $[5, 15, 25, 35]$ for the FastMRI dataset, and $[30, 60, 90, 130]$ for the IXI dataset. In red, the acquisition undersampling ratio s and the real reconstruction quality $Q(\boldsymbol{x}, \hat{\boldsymbol{x}})$[dB]. For improved visualization, all images have been linearly rescaled to enhance visual contrast.

5. Conclusion

We present a novel acquisition method able to adapt the sampling rate of the inference acquisition. We divide the generic slice acquisition into N temporally ordered batches, where each batch acquires a portion of the k-space. At the end of every batch, a new reconstructed slice is computed and its reconstruction quality is estimated via a neural predictor. Because every batch brings more information, the quality monotonically improves batch after batch. The acquisition terminates after a certain target quality is reached with a number of iterations that is, in general, lower than N.

We test the method on a deep neural state-of-the-art encoder–decoder [24] demonstrating how our acquisition method is able to reach the target quality with a sensibly higher acceleration than the

standard acquisition. In particular, tests have been carried out on the FastMRI dataset and on the IXI dataset for both Unconstrained and Line-Constrained sampling strategies.

Because our method relies on the use of a quality estimator, we give a detailed description of the adopted architecture. We also prove that, in the considered cases, the adoption of the proposed neural quality estimator returns an overall performance close to the one observed with an oracle-based solution.

Our current approach focuses on single-coil reconstructions for simplicity, but integrating multi-channel coils could enhance performance by utilizing the spatial sensitivity profiles of each coil. Additionally, while our method outputs real-valued images, this may lead to the loss of phase information present in complex-valued images. A potential enhancement would be to produce complex-valued outputs. These are promising directions for future work.

The total processing time achieved by our method is 5.1 ms for the knee slices of 320 × 320 samples and 3.1 ms for brain slices of 256 × 256 samples. These runtimes, which include both reconstruction and quality assessment, are shorter than the typical Echo Times (TE) reported for conventional MRI sequences, which range from 22 to 65 ms for conventional Spin Echo sequence knee acquisition protocols [30]. This suggests that our method could be potentially integrated into real-time MRI acquisition pipelines, enabling dynamic adjustments to the undersampling ratio.

CRediT authorship contribution statement

Filippo Martinini: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Mauro Mangia:** Writing – original draft, Visualization, Validation, Supervision, Methodology, Formal analysis, Data curation, Conceptualization. **Alex Marchioni:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Methodology, Formal analysis, Data curation, Conceptualization. **Gianluca Setti:** Supervision, Project administration, Conceptualization. **Riccardo Rovatti:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Research partly funded by PNRR - M4C2 - Investimento 1.3, Partenariato Esteso PE00000013 - "FAIR - Future Artificial Intelligence Research" - Spoke 8 "Pervasive AI", funded by the European Commission under the NextGeneration EU programme.

The authors would like to thank Santissima Margherita Mottola da MRI for her valuable assistance.

Data availability

Data will be made available on request.

References

[1] T. Soldatos, J.P. Shah, A. Chhabra, 3-dimensional (3D) isotropic MRI of the shoulder - advantages over 2D MRI, Semin. Roentgenol. (2024) http://dx.doi.org/10.1053/j.ro.2024.06.004, URL: https://www.sciencedirect.com/science/article/pii/S0037198X24000567.

[2] D.L. Donoho, Compressed sensing, IEEE Trans. Inform. Theory 52 (4) (2006) 1289–1306, http://dx.doi.org/10.1109/TIT.2006.871582, URL: http://ieeexplore.ieee.org/document/1614066/.

[3] E. Candes, T. Tao, Decoding by linear programming, IEEE Trans. Inform. Theory 51 (12) (2005) 4203–4215, http://dx.doi.org/10.1109/TIT.2005.858979, URL: http://ieeexplore.ieee.org/document/1542412/.

[4] M. Lustig, D.L. Donoho, J. Santos, J. Pauly, Compressed sensing MRI, IEEE Signal Process. Mag. 25 (2) (2008) 72–82, http://dx.doi.org/10.1109/MSP.2007.914728, URL: http://ieeexplore.ieee.org/document/4472246/.

[5] P. Li, H. Ge, P. Geng, Signal and image reconstruction with tight frames via unconstrained l1 l2 analysis minimizations, Signal Process. 203 (2023) 108755, http://dx.doi.org/10.1016/j.sigpro.2022.108755, URL: https://www.sciencedirect.com/science/article/pii/S0165168422002948.

[6] J. Liu, Y. Pan, M. Li, Z. Chen, L. Tang, C. Lu, J. Wang, Applications of deep learning to MRI images: A survey, Big Data Min. Anal. 1 (1) (2018) 1–18, http://dx.doi.org/10.26599/BDMA.2018.9020001.

[7] S. Ma, W. Yin, Y. Zhang, 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, 2008, pp. 1–8, http://dx.doi.org/10.1109/CVPR.2008.4587391, URL: http://ieeexplore.ieee.org/document/4587391/.

[8] Z. He, K. Hong, J. Zhou, D. Liang, Y. Wang, Q. Liu, Deep frequency-recurrent priors for inverse imaging reconstruction, Signal Process. 190 (2022) 108320, http://dx.doi.org/10.1016/j.sigpro.2021.108320, URL: https://www.sciencedirect.com/science/article/pii/S0165168421003571.

[9] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 9351, 2015, pp. 234–241, http://dx.doi.org/10.1007/978-3-319-24574-4_28.

[10] K.H. Jin, M.T. McCann, E. Froustey, M. Unser, Deep convolutional neural network for inverse problems in imaging, IEEE Trans. Image Process. 26 (9) (2017) 4509–4522, http://dx.doi.org/10.1109/TIP.2017.2713099.

[11] L. Sun, Z. Fan, X. Fu, Y. Huang, X. Ding, J. Paisley, A deep information sharing network for multi-contrast compressed sensing MRI reconstruction, IEEE Trans. Image Process. 28 (12) (2019) 6141–6153, http://dx.doi.org/10.1109/TIP.2019.2925288.

[12] E. Cha, G. Oh, J.C. Ye, Geometric approaches to increase the expressivity of deep neural networks for MR reconstruction, IEEE J. Sel. Top. Sign. Proces. 14 (6) (2020) 1292–1305, http://dx.doi.org/10.1109/JSTSP.2020.2982777.

[13] S.U. Dar, M. Yurt, M. Shahdloo, M.E. Ildız, B. Tınaz, T. Çukur, Prior-guided image reconstruction for accelerated multi-contrast MRI via generative adversarial networks, IEEE J. Sel. Top. Sign. Proces. 14 (6) (2020) 1072–1087, http://dx.doi.org/10.1109/JSTSP.2020.3001737.

[14] G. Oh, B. Sim, H. Chung, L. Sunwoo, J.C. Ye, Unpaired deep learning for accelerated MRI using optimal transport driven CycleGAN, IEEE Trans. Comput. Imaging 6 (2020) 1285–1296, http://dx.doi.org/10.1109/TCI.2020.3018562.

[15] Z. Chen, Y. Fu, Y. Xiang, Y. Zhu, A novel MR image denoising via LRMA and NLSS, Signal Process. 185 (2021) 108109, http://dx.doi.org/10.1016/j.sigpro.2021.108109, URL: https://www.sciencedirect.com/science/article/pii/S016516842100147X.

[16] L. Qiusheng, F. Xiaoyu, S. Baoshun, Z. Xiaohua, Compressed sensing MRI based on the hybrid regularization by denoising and the epigraph projection, Signal Process. 170 (2020) 107444, http://dx.doi.org/10.1016/j.sigpro.2019.107444, URL: https://www.sciencedirect.com/science/article/pii/S0165168419304955.

[17] A. Majumdar, Solving inverse problems with autoencoders on learnt graphs, Signal Process. 190 (2022) 108300, http://dx.doi.org/10.1016/j.sigpro.2021.108300, URL: https://www.sciencedirect.com/science/article/pii/S0165168421003376.

[18] L. Manovi, L. Capelli, A. Marchioni, F. Martinini, G. Setti, M. Mangia, R. Rovatti, SVD-based peephole and clustering to enhance trustworthiness in DNN classifiers, in: 2024 IEEE 6th International Conference on AI Circuits and Systems, AICAS, 2024, pp. 129–133, http://dx.doi.org/10.1109/AICAS59952.2024.10595919.

[19] U. Gamper, P. Boesiger, S. Kozerke, Compressed sensing in dynamic MRI, Magn. Reson. Med. 59 (2) (2008) 365–373, http://dx.doi.org/10.1002/mrm.21477, URL: http://doi.wiley.com/10.1002/mrm.21477.

[20] Z. Wang, G. Arce, Variable density compressed image sampling, IEEE Trans. Image Process. 19 (1) (2010) 264–270, http://dx.doi.org/10.1109/TIP.2009.2032889, URL: http://ieeexplore.ieee.org/document/5256257/.

[21] J.P. Haldar, D. Hernando, Zhi-Pei Liang, Compressed-sensing MRI with random encoding, IEEE Trans. Med. Imaging 30 (4) (2011) 893–903, http://dx.doi.org/10.1109/TMI.2010.2085084, URL: http://ieeexplore.ieee.org/document/5599301/.

[22] C.D. Bahadir, A.V. Dalca, M.R. Sabuncu, Learning-based optimization of the under-sampling pattern in MRI, in: Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 2019, pp. 780–792, http://dx.doi.org/10.1007/978-3-030-20351-1_61.

[23] C.D. Bahadir, A.Q. Wang, A.V. Dalca, M.R. Sabuncu, Deep-learning-based optimization of the under-sampling pattern in MRI, IEEE Trans. Comput. Imaging 6 (2020) 1139–1152, http://dx.doi.org/10.1109/TCI.2020.3006727.

[24] F. Martinini, M. Mangia, A. Marchioni, R. Rovatti, G. Setti, A deep learning method for optimal undersampling patterns and image recovery for MRI exploiting losses and projections, IEEE J. Sel. Top. Sign. Proces. 16 (4) (2022) 713–724, http://dx.doi.org/10.1109/JSTSP.2022.3171082.

[25] J. Xie, J. Zhang, Y. Zhang, X. Ji, PUERT: Probabilistic under-sampling and explicable reconstruction network for CS-MRI, IEEE J. Sel. Top. Sign. Proces. 16 (4) (2022) 737–749, http://dx.doi.org/10.1109/JSTSP.2022.3170654.

[26] Z. Wang, B. Li, W. Xia, C. Shen, M. Hou, H. Chen, Y. Liu, J. Zhou, Y. Zhang, Leaders: Learnable deep radial subsampling for mri reconstruction, in: 2022 IEEE 19th International Symposium on Biomedical Imaging, ISBI, 2022, pp. 1–5, http://dx.doi.org/10.1109/ISBI52829.2022.9761544.

[27] I.A. Huijben, B.S. Veeling, R.J. van Sloun, Learning sampling and model-based signal recovery for compressed sensing MRI, in: ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, 2020, pp. 8906–8910, http://dx.doi.org/10.1109/ICASSP40776.2020.9053331.

[28] H.K. Aggarwal, M. Jacob, J-MoDL: Joint model-based deep learning for optimized sampling and reconstruction, IEEE J. Sel. Top. Sign. Proces. 14 (6) (2020) 1151–1162, http://dx.doi.org/10.1109/JSTSP.2020.3004094.

[29] L. Pineda, S. Basu, A. Romero, R. Calandra, M. Drozdzal, Active MR k-space sampling with reinforcement learning, in: A.L. Martel, P. Abolmaesumi, D. Stoyanov, D. Mateus, M.A. Zuluaga, S.K. Zhou, D. Racoceanu, L. Joskowicz (Eds.), Medical Image Computing and Computer Assisted Intervention – MICCAI 2020, Springer International Publishing, Cham, 2020, pp. 23–33.

[30] P.M. Johnson, et al., Deep learning reconstruction enables prospectively accelerated clinical knee MRI, Radiology 220425, http://dx.doi.org/10.1148/radiol.220425, arXiv:https://doi.org/10.1148/radiol.220425 PMID: 36648347.

[31] A. Marchioni, F. Martinini, L. Manovi, S. Cortesi, R. Rovatti, G. Setti, M. Mangia, Adapted compressed sensing with incremental encoder and deep performance predictor for low-power sensor node design, in: 2023 IEEE International Instrumentation and Measurement Technology Conference, I2MTC, 2023, pp. 1–6, http://dx.doi.org/10.1109/I2MTC53148.2023.10175954.

[32] Z. Huang, S. Ravishankar, Single-pass object-adaptive data undersampling and reconstruction for MRI, IEEE Trans. Comput. Imaging 8 (2022) 333–345, http://dx.doi.org/10.1109/TCI.2022.3167454.

[33] T. Yin, Z. Wu, H. Sun, A.V. Dalca, Y. Yue, K.L. Bouman, End-to-end sequential sampling and reconstruction for MRI, in: Proceedings of Machine Learning for Health, in: Proceedings of Machine Learning Research, vol. 158, PMLR, 2021, pp. 261–281.

[34] K.H. Jin, M.A. Unser, K.M. Yi, Self-supervised deep active accelerated MRI, 2019, arXiv:1901.04547.

[35] V. Edupuganti, M. Mardani, S. Vasanawala, J. Pauly, Uncertainty quantification in deep MRI reconstruction, IEEE Trans. Med. Imaging 40 (1) (2021) 239–250, http://dx.doi.org/10.1109/TMI.2020.3025065.

[36] A. Mittal, A.K. Moorthy, A.C. Bovik, No-reference image quality assessment in the spatial domain, IEEE Trans. Image Process. 21 (12) (2012) 4695–4708, http://dx.doi.org/10.1109/TIP.2012.2214050.

[37] A. Mittal, R. Soundararajan, A.C. Bovik, Making a "Completely Blind" image quality analyzer, IEEE Signal Process. Lett. 20 (3) (2013) 209–212, http://dx.doi.org/10.1109/LSP.2012.2227726.

[38] S. Liu, K.-H. Thung, W. Lin, P.-T. Yap, D. Shen, Real-time quality assessment of pediatric MRI via semi-supervised deep nonlocal residual neural networks, IEEE Trans. Image Process. 29 (2020) 7697–7706, http://dx.doi.org/10.1109/TIP.2020.2992079.

[39] A. Rehman, Z. Wang, Reduced-reference image quality assessment by structural similarity estimation, IEEE Trans. Image Process. 21 (8) (2012) 3378–3389, http://dx.doi.org/10.1109/TIP.2012.2197011.

[40] T. Küstner, M. Jandt, A. Liebgott, L. Mauch, P. Martirosian, F. Bamberg, K. Nikolaou, S. Gatidis, F. Schick, B. Yang, Automatic motion artifact detection for whole-body magnetic resonance imaging, in: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, 2018, pp. 995–999, http://dx.doi.org/10.1109/ICASSP.2018.8462414.

[41] Q. Yan, D. Gong, Y. Zhang, Two-stream convolutional networks for blind image quality assessment, IEEE Trans. Image Process. 28 (5) (2019) 2200–2211, http://dx.doi.org/10.1109/TIP.2018.2883741.

[42] G.R. Chaithya, Z. Ramzi, P. Ciuciu, Learning the sampling density in 2D SPARKLING MRI acquisition for optimized image reconstruction, in: 2021 29th European Signal Processing Conference, EUSIPCO, 2021, pp. 960–964, http://dx.doi.org/10.23919/EUSIPCO54536.2021.9616336.

[43] C. Qin, J. Schlemper, J. Caballero, A.N. Price, J.V. Hajnal, D. Rueckert, Convolutional recurrent neural networks for dynamic MR image reconstruction, IEEE Trans. Med. Imaging 38 (1) (2019) 280–290, http://dx.doi.org/10.1109/TMI.2018.2863670.

[44] F. Sherry, M. Benning, J.C. De los Reyes, M.J. Graves, G. Maierhofer, G. Williams, C.-B. Schönlieb, M.J. Ehrhardt, Learning the sampling pattern for MRI, IEEE Trans. Med. Imaging 39 (12) (2020) 4310–4321, http://dx.doi.org/10.1109/TMI.2020.3017353.

[45] J.C. Ye, Compressed sensing MRI: a review from signal processing perspective, BMC Biomed. Eng. 30 (8) (2019) http://dx.doi.org/10.1186/s42490-019-0006-z.

[46] D.L. Donoho, Compressed sensing, IEEE Trans. Inform. Theory 52 (4) (2006) 1289–1306, http://dx.doi.org/10.1109/TIT.2006.871582.

[47] E. Candes, T. Tao, Decoding by linear programming, IEEE Trans. Inform. Theory 51 (12) (2005) 4203–4215, http://dx.doi.org/10.1109/TIT.2005.858979.

[48] Z. Wu, T. Yin, Y. Sun, R. Frost, A. van der Kouwe, A.V. Dalca, K.L. Bouman, Learning task-specific strategies for accelerated MRI, IEEE Trans. Comput. Imaging 10 (2024) 1040–1054, http://dx.doi.org/10.1109/TCI.2024.3410521.

[49] K. Ma, W. Liu, K. Zhang, Z. Duanmu, Z. Wang, W. Zuo, End-to-end blind image quality assessment using deep neural networks, IEEE Trans. Image Process. 27 (3) (2018) 1202–1213, http://dx.doi.org/10.1109/TIP.2017.2774045.

[50] S. Bosse, D. Maniry, K.-R. Müller, T. Wiegand, W. Samek, Deep neural networks for no-reference and full-reference image quality assessment, IEEE Trans. Image Process. 27 (1) (2018) 206–219, http://dx.doi.org/10.1109/TIP.2017.2760518.

[51] F. Martinini, M. Mangia, F. Pareschi, R. Rovatti, G. Setti, Compressed sensing inspired neural decoder for undersampled MRI with self-assessment, in: 2021 IEEE Biomedical Circuits and Systems Conference, BioCAS, 2021, pp. 01–06, http://dx.doi.org/10.1109/BioCAS49922.2021.9644958.

[52] Z. Jure, et al., fastMRI: An open dataset and benchmarks for accelerated MRI, 2018, arXiv:arXiv:1811.08839.

[53] IXI dataset. https://brain-development.org/ixi-dataset/ (Accessed 21 February 2023).

[54] Z. Wang, A. Bovik, H. Sheikh, E. Simoncelli, Image quality assessment: from error visibility to structural similarity, IEEE Trans. Image Process. 13 (4) (2004) 600–612, http://dx.doi.org/10.1109/TIP.2003.819861.