



# Multi-Coil MRI Reconstruction Challenge—Assessing Brain MRI Reconstruction Models and Their Generalizability to Varying Coil Configurations

## OPEN ACCESS

### Edited by:

Torbjørn Vefferstad Ness,  
Norwegian University of Life  
Sciences, Norway

### Reviewed by:

Sen Ma,  
Q Bio, Inc, United States  
Zhen Qiu,  
Michigan State University,  
United States

### \*Correspondence:

Roberto Souza  
roberto.medeirosdeso@ucalgary.ca

### Specialty section:

This article was submitted to  
Methods Article,  
a section of the journal  
Frontiers in Neuroscience

**Received:** 13 April 2022

**Accepted:** 01 June 2022

**Published:** 06 July 2022

### Citation:

Beauferris Y, Teuwen J,  
Karkalousos D, Moriakov N, Caan M,  
Yiasemis G, Rodrigues L, Lopes A,  
Pedrini H, Rittner L, Dannecker M,  
Studeniyak V, Gröger F, Vyas D,  
Faghih-Roohi S, Kumar Jethi A,  
Chandra Raju J, Sivaprakasam M,  
Lasby M, Nogovitsyn N, Loos W,  
Frayne R and Souza R (2022)  
Multi-Coil MRI Reconstruction  
Challenge—Assessing Brain MRI  
Reconstruction Models and Their  
Generalizability to Varying Coil  
Configurations.  
*Front. Neurosci.* 16:919186.  
doi: 10.3389/fnins.2022.919186

Youssef Beauferris<sup>1,2,3</sup>, Jonas Teuwen<sup>4,5,6</sup>, Dimitrios Karkalousos<sup>7</sup>, Nikita Moriakov<sup>4,5</sup>,  
Matthan Caan<sup>7</sup>, George Yiasemis<sup>5,6</sup>, Livia Rodrigues<sup>8</sup>, Alexandre Lopes<sup>9</sup>, Helio Pedrini<sup>9</sup>,  
Letícia Rittner<sup>8</sup>, Maik Dannecker<sup>10</sup>, Viktor Studeniyak<sup>10</sup>, Fabian Gröger<sup>10</sup>, Devendra Vyas<sup>10</sup>,  
Shahrooz Faghih-Roohi<sup>10</sup>, Amrit Kumar Jethi<sup>11</sup>, Jaya Chandra Raju<sup>11</sup>,  
Mohanasankar Sivaprakasam<sup>11,12</sup>, Mike Lasby<sup>1,3</sup>, Nikita Nogovitsyn<sup>13,14</sup>,  
Wallace Loos<sup>2,3,15,16</sup>, Richard Frayne<sup>2,3,15,16</sup> and Roberto Souza<sup>1,2,3\*</sup>

<sup>1</sup> (AI)<sup>2</sup> Lab, Electrical and Software Engineering, University of Calgary, Calgary, AB, Canada, <sup>2</sup> Biomedical Engineering Graduate Program, University of Calgary, Calgary, AB, Canada, <sup>3</sup> Hotchkiss Brain Institute, University of Calgary, Calgary, AB, Canada, <sup>4</sup> Department of Medical Imaging, Radboud University Medical Center, Nijmegen, Netherlands, <sup>5</sup> Department of Radiation Oncology, Netherlands Cancer Institute, Amsterdam, Netherlands, <sup>6</sup> Innovation Centre for Artificial Intelligence – Artificial Intelligence for Oncology, University of Amsterdam, Amsterdam, Netherlands, <sup>7</sup> Department of Biomedical Engineering and Physics, Amsterdam University Medical Centre, University of Amsterdam, Amsterdam, Netherlands, <sup>8</sup> Medical Image Computing Lab, School of Electrical and Computer Engineering, University of Campinas, Campinas, Brazil, <sup>9</sup> Institute of Computing, University of Campinas, Campinas, Brazil, <sup>10</sup> Computer Aided Medical Procedures, Technical University of Munich, Munich, Germany, <sup>11</sup> Department of Electrical Engineering, Indian Institute of Technology Madras, Chennai, India, <sup>12</sup> Healthcare Technology Innovation Centre, Indian Institute of Technology Madras, Chennai, India, <sup>13</sup> Centre for Depression and Suicide Studies, St. Michael's Hospital, Toronto, ON, Canada, <sup>14</sup> Mood Disorders Program, Department of Psychiatry and Behavioural Neurosciences, McMaster University, Hamilton, ON, Canada, <sup>15</sup> Radiology and Clinical Neurosciences, University of Calgary, Calgary, AB, Canada, <sup>16</sup> Seaman Family MR Research Centre, Foothills Medical Center, Calgary, AB, Canada

Deep-learning-based brain magnetic resonance imaging (MRI) reconstruction methods have the potential to accelerate the MRI acquisition process. Nevertheless, the scientific community lacks appropriate benchmarks to assess the MRI reconstruction quality of high-resolution brain images, and evaluate how these proposed algorithms will behave in the presence of small, but expected data distribution shifts. The multi-coil MRI (MC-MRI) reconstruction challenge provides a benchmark that aims at addressing these issues, using a large dataset of high-resolution, three-dimensional, T1-weighted MRI scans. The challenge has two primary goals: (1) to compare different MRI reconstruction models on this dataset and (2) to assess the generalizability of these models to data acquired with a different number of receiver coils. In this paper, we describe the challenge experimental design and summarize the results of a set of baseline and state-of-the-art brain MRI reconstruction models. We provide relevant comparative information on the current MRI reconstruction state-of-the-art and highlight the challenges of obtaining generalizable models that are required prior to broader clinical adoption. The MC-MRI benchmark data, evaluation code, and current challenge leaderboard are publicly available. They provide

an objective performance assessment for future developments in the field of brain MRI reconstruction.

**Keywords:** machine learning, magnetic resonance imaging (MRI), benchmark, image reconstruction, inverse problems, brain imaging

## 1. INTRODUCTION

Brain magnetic resonance imaging (MRI) is a commonly used diagnostic imaging modality. It is a non-invasive technique that provides images with excellent soft-tissue contrast. Brain MRI produces a wealth of information, which often leads to a definitive diagnosis of a number of neurological conditions, such as cancer and stroke. Furthermore, it is broadly adopted in neuroscience and other research domains. MRI data acquisition occurs in the Fourier or spatial-frequency domain, more commonly referred to as  $k$ -space. Image reconstruction consists of transforming the acquired  $k$ -space raw data into interpretable images. Traditionally, data is collected following the Nyquist sampling theorem (Lustig et al., 2008), and for a single-coil acquisition, a simple inverse Fourier Transform operation is often sufficient to reconstruct an image. However, the fundamental physics, practical engineering aspects, and biological tissue response factors underlying the MRI data acquisition process make fully sampled acquisitions inherently slow. These limitations represent a crucial drawback when MRI is compared to other medical imaging modalities, impact both patient tolerance of the procedure and throughput, and more broadly neuroimaging research.

Parallel imaging (PI) (Pruessmann et al., 1999; Griswold et al., 2002; Deshmane et al., 2012) and compressed sensing (CS) (Lustig et al., 2007; Liang et al., 2009) are two proven approaches that are able to reconstruct high-fidelity images from sub-Nyquist sampled acquisitions. PI techniques leverage the spatial information available across multiple, spatially distinct, receiver coils to allow the reconstruction of undersampled  $k$ -space data. Techniques, such as generalized autocalibrating partially parallel acquisition (GRAPPA) (Griswold et al., 2002), which operates in the  $k$ -space domain, and sensitivity encoding for fast MRI (SENSE) (Pruessmann et al., 1999), which works in the image domain, are currently used clinically. CS methods leverage image sparsity properties to improve reconstruction quality from undersampled  $k$ -space data. Some CS techniques, such as compressed SENSE (Liang et al., 2009), have also seen clinical adoption. Those PI and CS methods that have been approved for routine clinical use are generally restricted to relatively conservative acceleration factors (e.g.,  $R = 2\times$  to  $3\times$  acceleration). Currently employed comprehensive brain MRI scanning protocols, even those that use PI and CS, typically require between 30 and 45 min per patient procedure. Longer procedural times increase patient discomfort, thus lessening the likelihood of patient acceptance. It also increases susceptibility to both voluntary and involuntary motion artifacts.

In 2016, the first deep-learning-based MRI reconstruction models were presented (Sun et al., 2016; Wang et al., 2016). The excellent initial results obtained by these models caught

the attention of the MR imaging community, and subsequently, dozens of deep-learning-based MRI reconstruction models were proposed (cf., Sun et al., 2016; Wang et al., 2016; Kwon et al., 2017; Schlemper et al., 2017, 2018, 2019; Dedmari et al., 2018; Eo et al., 2018a,b; Gözcü et al., 2018; Hammernik et al., 2018; Quan et al., 2018; Seitzer et al., 2018; Yang et al., 2018; Zhang et al., 2018; Zhu et al., 2018; Akçakaya et al., 2019; Mardani et al., 2019; Pawar et al., 2019; Qin et al., 2019; Souza and Frayne, 2019; Zeng et al., 2019; Hosseini et al., 2020; Sriram et al., 2020b; Zhou and Zhou, 2020). Many of these studies demonstrated superior quantitative results from deep-learning-based methods compared to non-deep-learning-based MRI reconstruction algorithms (Schlemper et al., 2017; Hammernik et al., 2018; Knoll et al., 2020). These new methods are also capable of accelerating MRI examinations beyond traditional PI and CS methods. There is good evidence that deep-learning-based MRI reconstruction methods can accelerate MRI examinations by factors greater than 5 (Zbontar et al., 2018; Souza et al., 2020a).

A significant drawback, that hinders the progress of the brain MRI reconstruction field, is the lack of benchmark datasets. Importantly, the lack of benchmarks makes the comparison of different methods challenging. The fastMRI effort (Zbontar et al., 2018) is an important initiative that provides large volumes of raw MRI  $k$ -space data. The initial release of the fastMRI dataset provided two-dimensional (2D) MR acquisitions of the knee. A subsequent release added 2D brain MRI data with 5 mm slice thickness, which was used for the 2020 *fastMRI* challenge (Muckley et al., 2021). The *Calgary-Campinas* (Souza et al., 2018) initiative contains numerous sets of brain imaging data. For the purposes of this benchmark, we expanded the *Calgary-Campinas* initiative to include MRI raw data from three-dimensional (3D), high-resolution acquisitions. High-resolution images are crucial for many neuroimaging applications. Also importantly, 3D acquisitions allow for undersampling along two phase encoding dimensions, instead of one for 2D imaging. This potentially allows for further MRI acceleration. These  $k$ -space datasets correspond to either 12- or 32-channel data.

The goals of the multi-coil MRI (MC-MRI - <https://www.ccdataset.com/mr-reconstruction-challenge>) reconstruction challenge are to provide benchmarks that help improve the quality of brain MRI reconstruction, facilitate comparison of different reconstruction models, better understand the difficulties related to clinical adoption of these models, and investigate the upper limits of MR acceleration. The specific objectives of the challenge are as follows:

1. Compare the performance of different brain MRI reconstruction models on a large dataset, and
2. Assess the generalizability of these models to datasets acquired with different coils.

The results presented in this report correspond to benchmark submissions received up to 20 November, 2021. Four baseline solutions and three new benchmark solutions were presented and discussed during an online session at the Medical Imaging Deep Learning Conference held on 9 July, 2020.<sup>1</sup> Two additional benchmark solutions were submitted after the online session. Collectively, these results provide a relevant performance summary of some state of the art MRI reconstruction approaches, including different model architectures, processing strategies, and emerging metrics for training and assessing reconstruction models. The MC-MRI reconstruction challenge is ongoing and open to new benchmark submissions.<sup>2</sup> A public code repository with instructions on how to load the data, extract the benchmark metrics, and baseline reconstruction models are available at <https://github.com/rmsouza01/MC-MRI-Rec>.

## 2. MATERIALS AND METHODS

### 2.1. Calgary-Campinas Raw MRI Dataset

The data used in this challenge were acquired as part of the Calgary Normative Study (McCreary et al., 2020), which is a multi-year, longitudinal project that investigates normal human brain aging by acquiring quantitative MRI data using a protocol approved by our local research ethics board. Raw data from T1-weighted volumetric imaging was acquired, anonymized, and incorporated into the *Calgary-Campinas* (CC) dataset (Souza et al., 2018). The publicly accessible dataset currently provides  $k$ -space data from 167 3D, T1-weighted, gradient-recalled echo, 1 mm<sup>3</sup> isotropic sagittal acquisitions collected on a clinical 3-T MRI scanner (Discovery MR750; General Electric Healthcare, Waukesha, WI). The brain scans are from presumed healthy subjects (mean  $\pm$  standard-deviation age: 44.5  $\pm$  15.5 years; range: 20 years to 80 years; 71/167 (42.5%) male).

The datasets were acquired using either a 12-channel (117 scans, 70.0%) or 32-channel receiver coil (50 scans, 30.0%). Acquisition parameters were TR/TE/TI = 6.3 ms / 2.6 ms / 650 ms (93 scans, 55.7%) or TR/TE/TI = 7.4 ms / 3.1 ms / 400 ms (74 scans, 44.3%), with 170 to 180 contiguous 1.0 mm slices and a field of view of 256 mm  $\times$  218 mm. The acquisition matrix size [ $N_x, N_y, N_z$ ] for each channel was [256, 218, 170–180], where  $x$ ,  $y$ , and  $z$  denote readout, phase-encode, and slice-encode directions, respectively. In the slice-encode ( $k_z$ ) direction, only 85% of the  $k$ -space data were collected; the remainder (15% of 170–180) was zero-filled. This partial acquisition technique is common practice in MRI. The average scan duration is 341 s. Because  $k$ -space undersampling only occurs in the phase-encode and slice-encode directions, the 1D inverse Fourier Transform (iFT) along  $k_x$  was automatically performed by the scanner, and hybrid ( $x, k_y, k_z$ ) datasets were provided. This pre-processing effectively allows the MRI reconstruction problem to be treated as a 2D problem (in  $k_y$  and  $k_z$ ). The partial Fourier reference data was reconstructed by taking the 2D iFT along the  $k_y - k_z$  plane for each individual

channel and combining these using the conventional square-root sum-of-squares algorithm (Larsson et al., 2003).

### 2.2. MC-MRI Reconstruction Challenge Description

The MC-MRI Reconstruction Challenge was designed to be an ongoing investigation that will be disseminated through a combination of in-person sessions at meetings and virtual sessions, supplemented by periodic online submissions and updates. The benchmark is readily extensible and more data, metrics, and research questions are expected to be added in further updates. Individual research groups are permitted to make multiple submissions. The processing of submissions is semi-automated, and it takes on average 48 h to generate an update of the benchmark leaderboard.

Currently, the MC-MRI reconstruction challenge is split into two separate tracks. Teams can decide whether to submit a solution to just one track or to both tracks. Each track has a separate leaderboard. The tracks are:

- **Track 01:** Teams had access to 12-channel data to train and validate their models. Models submitted are evaluated by only using the 12-channel test data.
- **Track 02:** Teams had access to 12-channel data to train and validate their models. Models submitted are evaluated for both the 12-channel and 32-channel test data.

In both tracks, the goal is to assess the brain MR image reconstruction quality and in particular note any loss of high-frequency details, especially at the higher acceleration rates. By having two separate tracks, we hoped to determine whether a generic reconstruction model trained on data from one coil would have decreased performance when applied to data from another coil.

Two MRI acceleration factors were tested:  $R = 5$  and  $R = 10$ . These factors were chosen intentionally to exceed the acceleration factors typically used clinically with PI and CS methods. A Poisson disc distribution sampling scheme, where the center of  $k$ -space was fully sampled within a circle of radius of 16 pixels to preserve the low-frequency phase information, was used to achieve these acceleration factors. For brevity, we have only reported the results for  $R = 5$ , but the online challenge leaderboard contains the results for both acceleration factors.

The training, validation, and test split of the challenge data are summarized in **Table 1**. The initial 50 and last 50 slices in each participant's image volume were removed because they have little anatomy present. The fully sampled  $k$ -space data of the training and validation sets were made public for teams to develop their models. Pre-undersampled  $k$ -space data corresponding to the test sets were provided for the teams for accelerations of  $R = 5$  and  $R = 10$ .

### 2.3. Quantitative Metrics

In order to measure the quality of the image reconstructions, three commonly used, quantitative performance metrics were selected: peak signal-to-noise ratio (pSNR), structural similarity (SSIM) index (Zhou et al., 2004), and visual information fidelity (VIF) (Sheikh and Bovik, 2006). The choice of performance

<sup>1</sup>See video of session at <https://www.ccdataset.com/mr-reconstruction-challenge/mc-mrrec-2020-midl-recording>.

<sup>2</sup>See current leaders for the individual challenge tracks at <https://www.ccdataset.com/>.

**TABLE 1** | Summary of the raw MRI  $k$ -space datasets used in the first edition of the challenge.

Coil	Category	# of datasets	# of slices
12-channel	Train	47	12,032
	Validation	20	5,120
	Test	50	7,800
32-channel	Test	50	7,800

Reported are the number of slices in the test sets after removal of the initial 50 and last 50 slices (see text).

metrics is challenging, and it is recognized that objective measures, such as pSNR, SSIM, and VIF may not correlate well with subjective human image quality assessments. Nonetheless, these metrics provide a broad basis to assess model performance in this challenge.

The pSNR is a metric commonly used for MRI reconstruction assessment and consists of the log ratio between the maximum value of the reference reconstruction and the root mean squared error (RMSE):

$$\begin{aligned} \text{pSNR}(y, \hat{y}) &= 20 \log_{10} \left( \frac{\max(y)}{\text{RMSE}} \right) \\ &= 20 \log_{10} \left( \frac{\max(y)}{\sqrt{\frac{1}{M} \sum_{i=1}^M [y(i) - \hat{y}(i)]^2}} \right), \quad (1) \end{aligned}$$

where  $y$  is the reference image,  $\hat{y}$  is the reconstructed image, and  $M$  is the number of pixels in the image. Higher pSNR values represent higher-fidelity image reconstructions. However, pSNR does not take into consideration the factors involved in human vision. For this reason, increased pSNR can suggest that reconstructions are of higher quality, when in fact they may not be as well-perceived by the human visual system.

Unlike pSNR, SSIM and VIF are metrics that attempt to model aspects of the human visual system. SSIM considers biological factors, such as luminance, contrast, and structural information. SSIM is computed using:

$$\text{SSIM}(x, \hat{x}) = \frac{(2\mu_x \mu_{\hat{x}} + c_1)(2\sigma_{x\hat{x}} + c_2)}{(\mu_x^2 + \mu_{\hat{x}}^2 + c_1)(\sigma_x^2 + \sigma_{\hat{x}}^2 + c_2)} \quad (2)$$

where  $x$  and  $\hat{x}$  represent corresponding image windows from the reference image and the reconstructed image, respectively;  $\mu_x$  and  $\sigma_x$  represent the mean and standard-deviation inside the image window,  $x$ ; and  $\mu_{\hat{x}}$  and  $\sigma_{\hat{x}}$  represent the mean and standard-deviation inside the reconstructed image window,  $\hat{x}$ . The constants  $c_1$  and  $c_2$  are used to avoid numerical instability. SSIM values for non-negative images are within  $[0, 1]$ , where 1 represents two identical images.

The visual information fidelity metric is based on natural scene statistics (Simoncelli and Olshausen, 2001; Geisler, 2008). VIF models the natural scene statistics based on a Gaussian scale mixture model in the wavelet domain, and additive white Gaussian noise is used to model the human visual system. The natural scene of the reference image is modeled into wavelet components (C) and the human visual system is modeled by

adding zero-mean white Gaussian noise in the wavelet domain (N), which results in the perceived reference image ( $E = C + N$ ). In the same way, the reconstructed image, which is called the distorted image, is also modeled by a natural scene model (D) and the human visual system model (N'), leading to the perceived distorted image ( $F = D + N'$ ). The VIF is given by the ratio of the mutual information of  $I(C, F)$  and  $I(C, E)$ :

$$\text{VIF} = \frac{I(C, F)}{I(C, E)}, \quad (3)$$

where  $I$  represents the mutual information.

Mason et al. (2019) investigated the VIF metric for assessing MRI reconstruction quality. Their results indicated that it has a stronger correlation with subjective radiologist opinion about MRI quality than other metrics such as pSNR and SSIM. The VIF Gaussian noise variance was set to 0.4 as recommended in Mason et al. (2019). All metrics were computed slice-by-slice in the test set. The reference and reconstructed images were normalized by dividing them by the maximum value of the reference image.

## 2.4. Visual Assessment

An expert observer (NN) with over 5 years of experience analyzing brain MR images and manually segmenting complex structures, such as the hippocampus and hypothalamus, visually inspected 25 randomly selected volumes for the 12-channel test set and other 25 volumes for the 32-channel test set for the best two submissions as determined from the quantitative metrics. The best two submissions were obtained by sorting the weighted average ranking. The weighted average ranking was generated by applying pre-determined weights to the ranking of the three individual quantitative metrics (0.4 for VIF, 0.4 for SSIM, and 0.2 for pSNR). We chose to give higher weights to VIF and SSIM because they have a better correlation with the human perception of image quality.

The visual assessment of the images was done by comparing the machine-learning-based reconstructions to the fully sampled reference images. This allowed the observer to distinguish between data acquisition related quality issues (e.g., motion) and problems associated with image reconstruction. The image quality assessment focused mostly on overall image quality and how well-defined was the contrast between white-matter, gray-matter, and other relevant brain structures. The goal of the visual assessment was to compare the quality of the reconstructed MR images against the fully sampled reference images and not to compare the quality of the different submissions, because the benchmark is ongoing and we wanted to account for potential observer memory bias effects (Kalm and Norris, 2018) in the qualitative metrics due to the difference between submission dates of the different solutions to the benchmark (i.e., future submissions will be visually assessed at different dates compared to current submissions).

## 2.5. Models

Track 01 of the challenge included four baseline models, selected from the literature. These models are the zero-filled reconstruction, the U-Net model (Jin et al., 2017), the WW-net model (Souza et al., 2020b), and the hybrid-cascade

model (Souza et al., 2019). To date, Track 01 has received six independent submissions from ResoNNance (Yiasemis et al., 2022a) (two different models), The Enchanted (two different models), TUMRI, and M-L UNICAMP teams.

The ResoNNance 1.0 model submission was a recurrent inference machine (Lønning et al., 2019), ResoNNance 2.0 was a recurrent variational network (Yiasemis et al., 2022b). The Enchanted 1.0 model was inspired by Lee et al. (2018), where they used magnitude and phase networks, followed by a VS-net architecture (Duan et al., 2019). The Enchanted 2.0 (Jethi et al., 2022) used an end-to-end variational network (Sriram et al., 2020a), and it was the only submission that used self-supervised learning (Chen et al., 2019) to initialize their model. The pretext task to initialize their models was the prediction of image rotations (Gidaris et al., 2018). TUMRI used a similar model to the WW-net, but they implemented complex-valued operations (Trabelsi et al., 2018). They used a linear combination of VIF and MS-SSIM (Wang et al., 2003) as their loss function. M-L UNICAMP used a hybrid model with parallel network branches operating in  $k$ -space and image domains. Links to the source code for the different models are available in the benchmark repository. Some of the Track 01 models were designed to work with a specific number of coil channels, thus they were not submitted to Track 02 of the challenge.

Track 02 of the challenge included two baseline models (zero-filled reconstruction and the U-Net model). ResoNNance and The Enchanted teams submitted two models each to Track 02. The models submitted by ResoNNance and The Enchanted teams were the same models that were used for Track 01 of the challenge. **Table 2** summarizes the processing domains (image,  $k$ -space, or dual/hybrid), the presence of elements, such as coil sensitivity estimation, data consistency, and the loss function used during training of the models. For more details about the models, we refer the reader to the source publications or the code repositories for the unpublished work.

## 3. RESULTS

### 3.1. Track 01

The quantitative results for Track 01 are summarized in **Table 3**. There were in total 10 models (four baseline and six submitted) in Track 01. The zero-filled and U-Net reconstructions had the worst results. The M-L UNICAMP, Hybrid Cascade, WW-net, and TUMRI models were next with similar results in terms of SSIM and pSNR. Notably, the TUMRI submission achieved the second-highest VIF metric. ResoNNance and The Enchanted teams' submissions achieved the highest overall scores on the quantitative metrics. The ResoNNance 2.0 submission had the best SSIM and pSNR metrics and the fourth-best VIF metric. The Enchanted 1.0 submission obtained the best VIF metric. The Enchanted 2.0 submission achieved the second-best SSIM metric, and the third-best VIF and pSNR metrics. Representative reconstructions resulting from the different models for  $R = 5$  are shown in **Figure 1**.

Twenty five images in the test set were visually assessed by our expert observer for the two best submission (ResoNNance

**TABLE 2 |** Summary of the submissions including processing domain, presence of coil sensitivity estimation (SE), presence of data consistency (DC), and basis of the training loss functions.

Model	Domain	SE	DC	Loss function
ResoNNance 2.0	Hybrid	Yes	Yes	MAE and SSIM
The Enchanted 2.0	Hybrid	Yes	Yes	Cross entropy (pretext) and SSIM (main task)
ResoNNance 1.0	Image	Yes	Yes	MAE and SSIM
The-Enchanted 1.0	Image	Yes	Yes	MSE (first step) and SSIM (second step)
TUMRI	Hybrid	No	Yes	MS-SSIM and VIF
WW-Net*	Hybrid	No	Yes	MSE
Hybrid-cascade*	Hybrid	No	Yes	MSE
M-L UNICAMP	Hybrid	No	Yes	MSE
U-Net*	Image	No	No	MSE
Zero-filled*	N/A	No	N/A	N/A

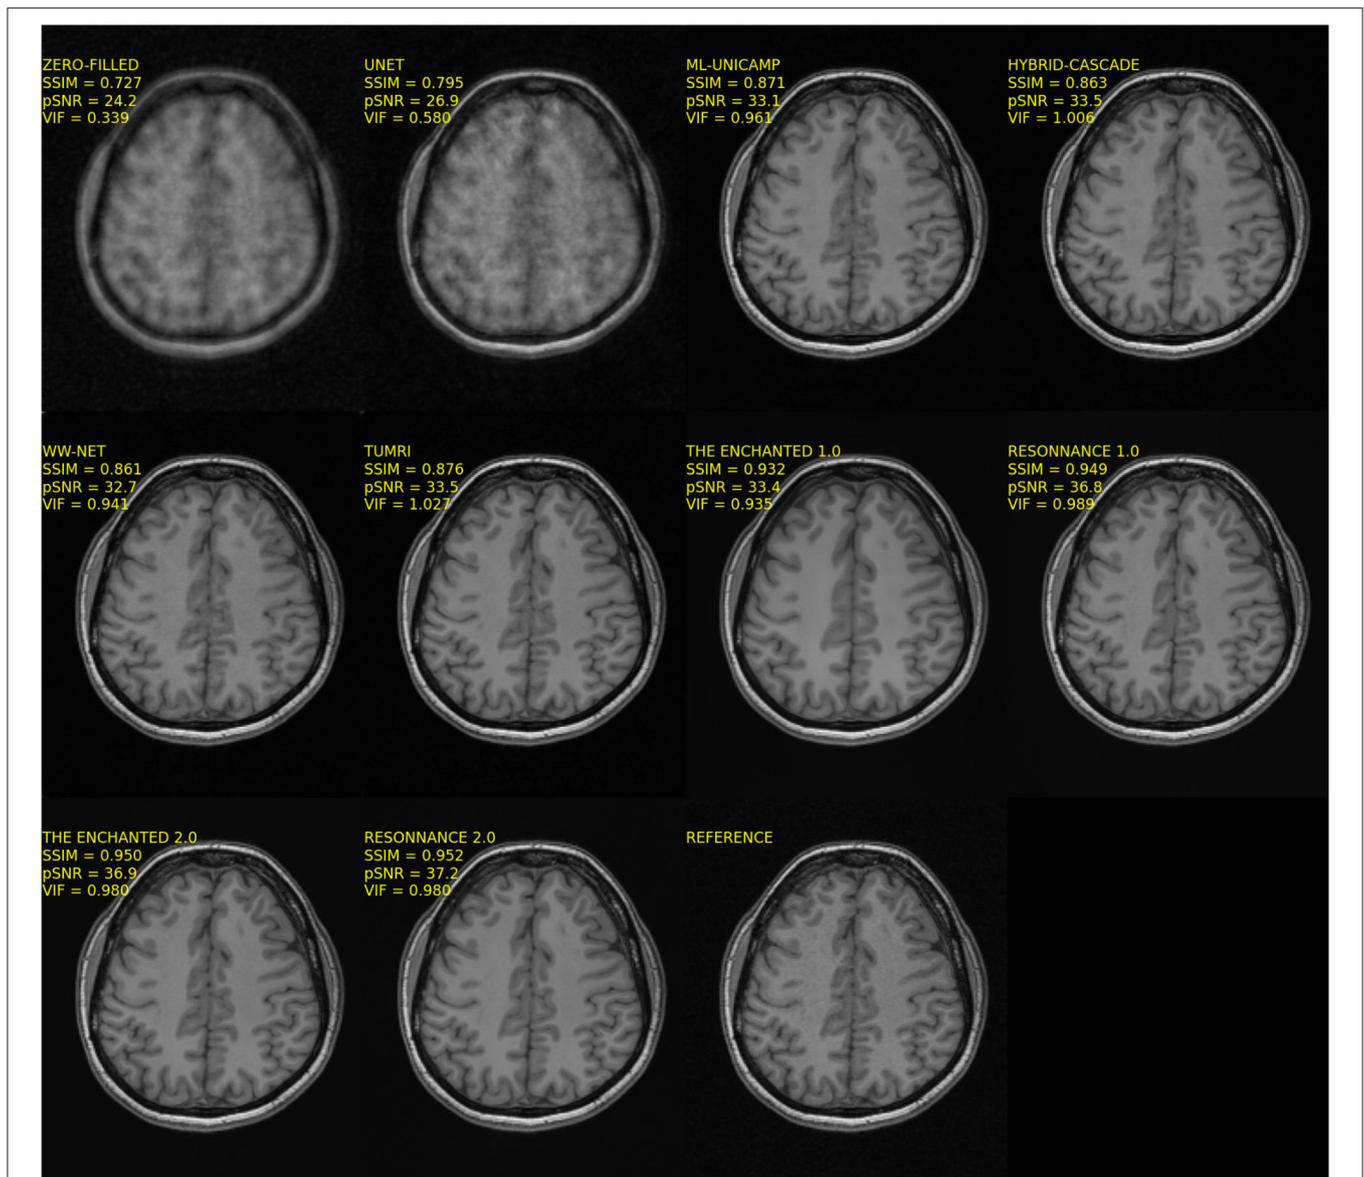
\* indicates a baseline model. Loss functions: Mean Absolute Error (MAE), Structural Similarity (SSIM), Mean Squared Error (MSE), Multi-Scale SSIM (MS-SSIM), and Visual Information Fidelity (VIF).

2.0 and The Enchanted 2.0). Out of the 50 images assessed by the expert observer, only two (4.0%) were deemed to have minor deviations, such as shape, intensity, and contrast between the reconstructed images and the reference (cf., **Figure 2A**). Twenty seven images (54.0%) were deemed to have a similar quality to the fully sampled reference, and 21 (42.0%) were rated as having similar quality when compared to the reference, but exhibited filtering of the noise in the image (cf., **Figure 2B**).

### 3.2. Track 02

Two teams, ResoNNance and The Enchanted, submitted a total of four models to Track 02 of the benchmark. Their results were compared to two baseline techniques. The models submitted to Track 02, except for the U-Net baseline, which has a higher input dimension (i.e., the input dimensions depends on the number of receiver coils), was the same as the models submitted for Track 01, so for the 12-channel test dataset, the results are the same as in Track 01 (see **Table 3**).

The results for Track 02 using the 32-channel test set are summarized in **Table 4**. For the 32-channel test dataset, The Enchanted 2.0 submission obtained the best VIF and pSNR metrics, and the second-best SSIM score. The ResoNNance 2.0 submission obtained the best SSIM metric, second-best pSNR, and third-best VIF metrics. The ResoNNance 1.0 submission obtained the third-best SSIM and pSNR metrics, and second-best VIF. The Enchanted 1.0 submission obtained the fourth-best



**FIGURE 1** | Representative reconstructions of the different models submitted to Track 01 (i.e., 12-channel) of the challenge for  $R = 5$ . Note that the reconstructions from the top four methods, ResoNNance 1.0 and 2.0, and The Enchanted 1.0 and 2.0, try to match the noise pattern seen in the background of the reference image, while ML-UNICAMP, Hybrid-cascade, WW-net, and TUMRI seem to have partially filtered this background noise.

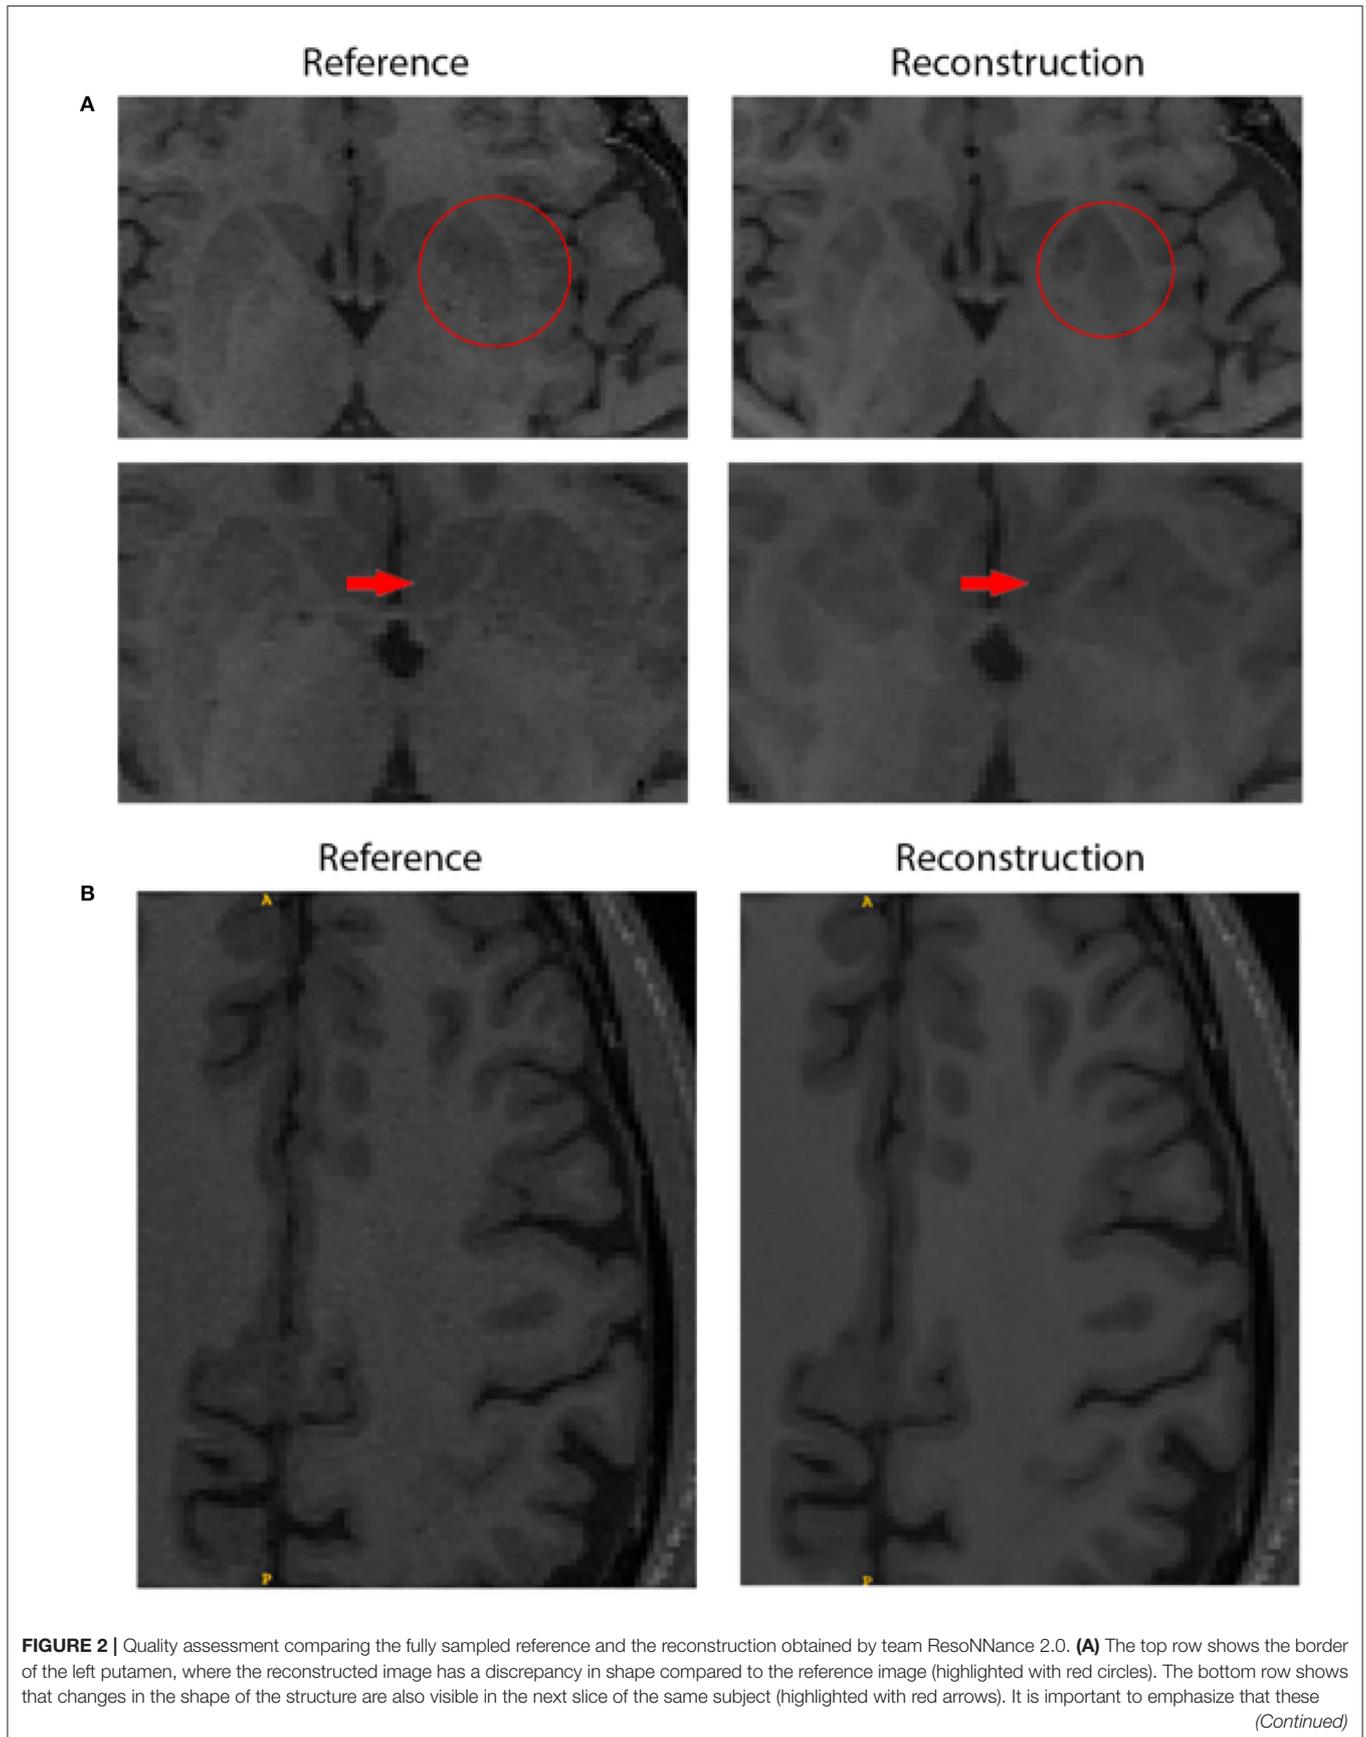
SSIM and VIF, and fifth-best pSNR. The zero-filled and U-Net reconstructions obtained the worse results. Representative reconstructions resulting from the different models are depicted in **Figure 3**.

Twenty five images in the test set were visually assessed by our expert observer for the two best submissions (ResoNNance 2.0 and The Enchanted 2.0). Out of the 50 images assessed by the expert observer, 14 (28.0%) were deemed to have deviations from common anatomical borders. A total of 34 images (68.0%) were deemed to have a similar quality to the fully sampled reference, and only two images (4.0%) were rated as having similar quality

when compared to the reference, but exhibited filtering of the noise in the image.

#### 4. DISCUSSION

The first track of the challenge compared ten different reconstruction models (**Table 3**). As expected, the zero-filled reconstruction, which does not involve any training from the data, universally had the poorest results. The second worst technique was the U-Net model, which used as input the channel-wise zero-filled reconstruction and tried to recover the



**FIGURE 2** | discrepancies are not restricted to the putamen, but a systematic evaluation of where these changes occur is out of scope for this work. **(B)** Illustration of a case where the expert observed rated that the deep-learning-based reconstruction improved image quality. In this figure, we can see smoothing of cortical white matter without loss of information as no changes appeared in the pattern of gyrification within cortical gray matter.

**TABLE 3** | Summary of the Track 01 results for  $R = 5$ .

Model	SSIM	pSNR (dB)	VIF
ResoNNance 2.0	<b>0.941 ± 0.029</b>	<b>35.7 ± 1.8</b>	0.957 ± 0.034
The Enchanted 2.0	0.937 ± 0.033	34.9 ± 2.4	0.973 ± 0.036
ResoNNance 1.0	0.936 ± 0.031	35.3 ± 1.8	0.960 ± 0.035
The-Enchanted 1.0	0.912 ± 0.034	30.3 ± 2.8	<b>0.993 ± 0.176</b>
TUMRI	0.868 ± 0.044	32.5 ± 1.7	0.989 ± 0.045
WW-net*	0.870 ± 0.043	32.5 ± 1.7	0.929 ± 0.049
Hybrid-cascade*	0.860 ± 0.044	32.7 ± 1.6	0.954 ± 0.042
M-L UNICAMP	0.868 ± 0.044	32.4 ± 1.7	0.918 ± 0.053
U-Net*	0.779 ± 0.039	26.8 ± 1.7	0.642 ± 0.068
Zero-filled*	0.726 ± 0.045	25.2 ± 1.5	0.518 ± 0.066

The best value for each metric and acceleration is emboldened. Mean ± standard deviation are reported. \* indicates a baseline model.

**TABLE 4** | Summary of the Track 02 results for  $R = 5$  using the 32-channel test set.

Model	SSIM	pSNR (dB)	VIF
ResoNNance 2.0	<b>0.961 ± 0.027</b>	38.3 ± 2.2	0.955 ± 0.036
The Enchanted 2.0	0.960 ± 0.037	<b>38.34 ± 3.2</b>	<b>1.024 ± 0.034</b>
ResoNNance 1.0	0.947 ± 0.033	37.7 ± 2.9	0.992 ± 0.030
The Enchanted 1.0	0.907 ± 0.046	30.1 ± 2.7	0.834 ± 0.236
U-Net*	0.832 ± 0.058	31.5 ± 2.6	0.804 ± 0.045
Zero-filled*	0.780 ± 0.041	26.4 ± 1.5	0.472 ± 0.064

The best value for each metric and acceleration is emboldened. Mean ± standard deviation are reported. \* indicates a baseline model.

high-fidelity image. The employed U-Net (Jin et al., 2017) model did not include any data consistency steps. The remaining eight models all include a data consistency step, which seems to be an essential step for high-fidelity image reconstruction, as has been previously highlighted in Schlemper et al. (2017) and Eo et al. (2018a).

The M-L UNICAMP model explored parallel architectures that operated both in the  $k$ -space and image domains. M-L UNICAMP had the eighth-lowest pSNR and VIF metrics, and the seventh-lowest SSIM score. In contrast, the top ranked methods were either cascaded networks (Hybrid-cascade, WW-net, TUMRI, The Enchanted 1.0 and 2.0) or recurrent methods (ResoNNance 1.0 and 2.0).

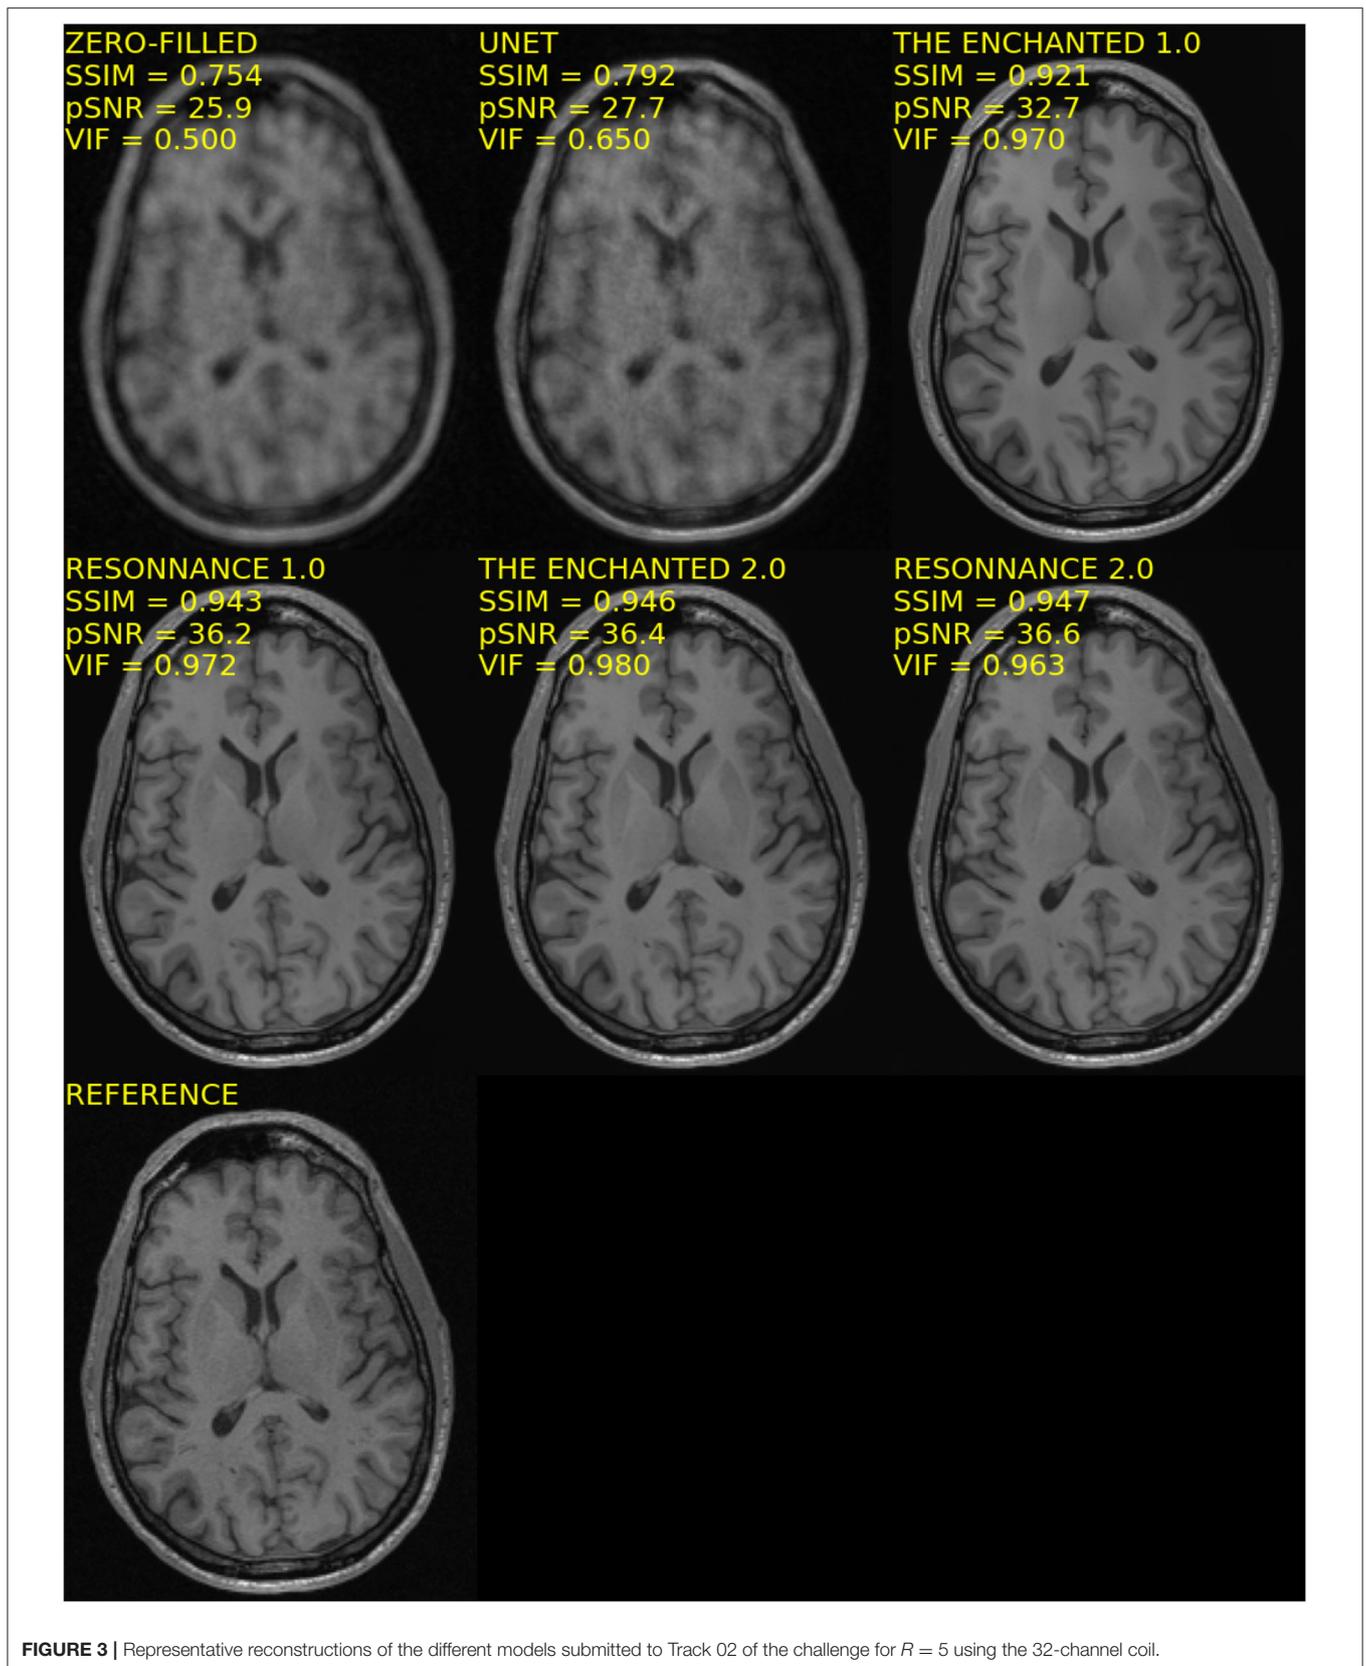
The top four models in the benchmark were the ResoNNance 1.0 and 2.0 and The Enchanted 1.0 and 2.0 submissions. These four models estimated coil sensitivities and combined the coil channels, which made these models flexible and capable of working with datasets acquired with an arbitrary number of receiver coils. The top two models ResoNNance 2.0 and Enchanted 2.0 are hybrid models. They are followed in rank

by ResoNNance 1.0 and Enchanted 1.0, which are image-domain methods. The other better performing models (M-L UNICAMP, Hybrid Cascade, WW-net, and TUMRI) used an approach that receives all coil channels as input, making these models tailored to a specific coil configuration (i.e., number of channels). Though the methods that combined the channels before reconstruction using coil sensitivity estimations similarly to Sriram et al. (2020a), such as from ResoNNance and The Enchanted teams, demonstrated the best results so far, it is still unclear if this approach is superior to models that do not combine the channels before reconstruction. A recent work (Sriram et al., 2020b) indicated that the separate channel approach may be advantageous compared to models that combine the  $k$ -space channels before reconstruction.

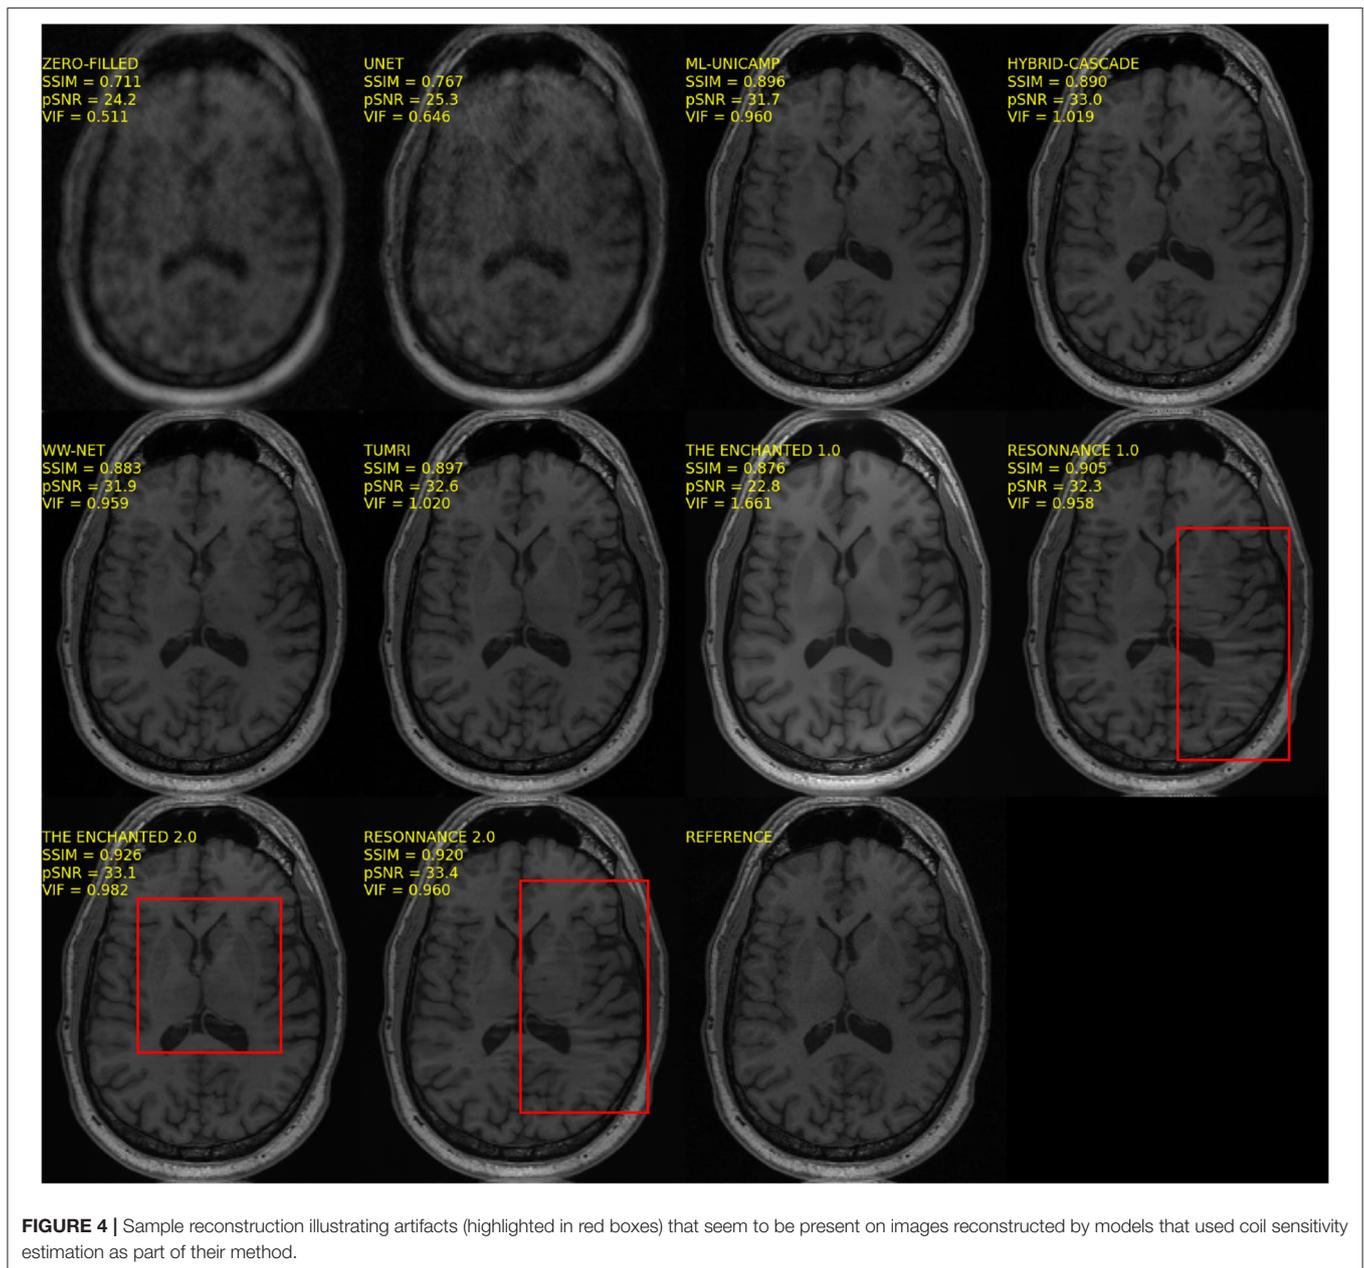
All of the models submitted to the MC-MRI Reconstruction Challenge had a relatively narrow input convolutional layer (e.g., 64 filters), which may have resulted in the loss of relevant information. In Sriram et al. (2020b), they used 15-channel data and the first layer had 384 filters. Another advantage of models that receive all channels as input is that they seem more robust to artifacts that can occur in the reconstructed images due to problems in coil sensitivity estimation. This finding was observed in our visual assessment mostly in methods that involved coil sensitivity estimation (ResoNNance and The Enchanted—Figure 4). Similar artifacts were not observed in images produced on models that do not require coil sensitivity estimation.

In our study, we also noted variability in the ranking across metrics (Table 3). For example, The Enchanted 1.0 submission had the best VIF score, but only the fourth-best SSIM and seventh-highest pSNR metrics. This variability reinforces the importance of including many benchmarks that can summarize the result of multiple submissions by using a consistent set of multiple metrics. Studies that use a single image quality metric, for example, are potentially problematic if the chosen measure masks specific classes of performance issues. While imperfect, the use of a composite score based on metric rankings attempts to reduce this inherent variability by examining multiple performance measures.

Visual inspection of the reconstructed MR images (cf., Figures 1, 3) indicates that with some models and for some samples in the test set, the reconstructed background noise is different from the background noise in the reference images. This observation, particularly with the ResoNNance and The Enchanted teams' submissions, leads to questions on whether the evaluated quantitative metrics are best suited to determine the reconstruction quality. Given a noisy reference image, a noise-free reconstruction will potentially achieve lower pSNR, SSIM, and VIF than the same reconstruction with added noise. This finding is contrary to human visual perception, where noise impacts the image quality negatively and is, in



**FIGURE 3** | Representative reconstructions of the different models submitted to Track 02 of the challenge for  $R = 5$  using the 32-channel coil.

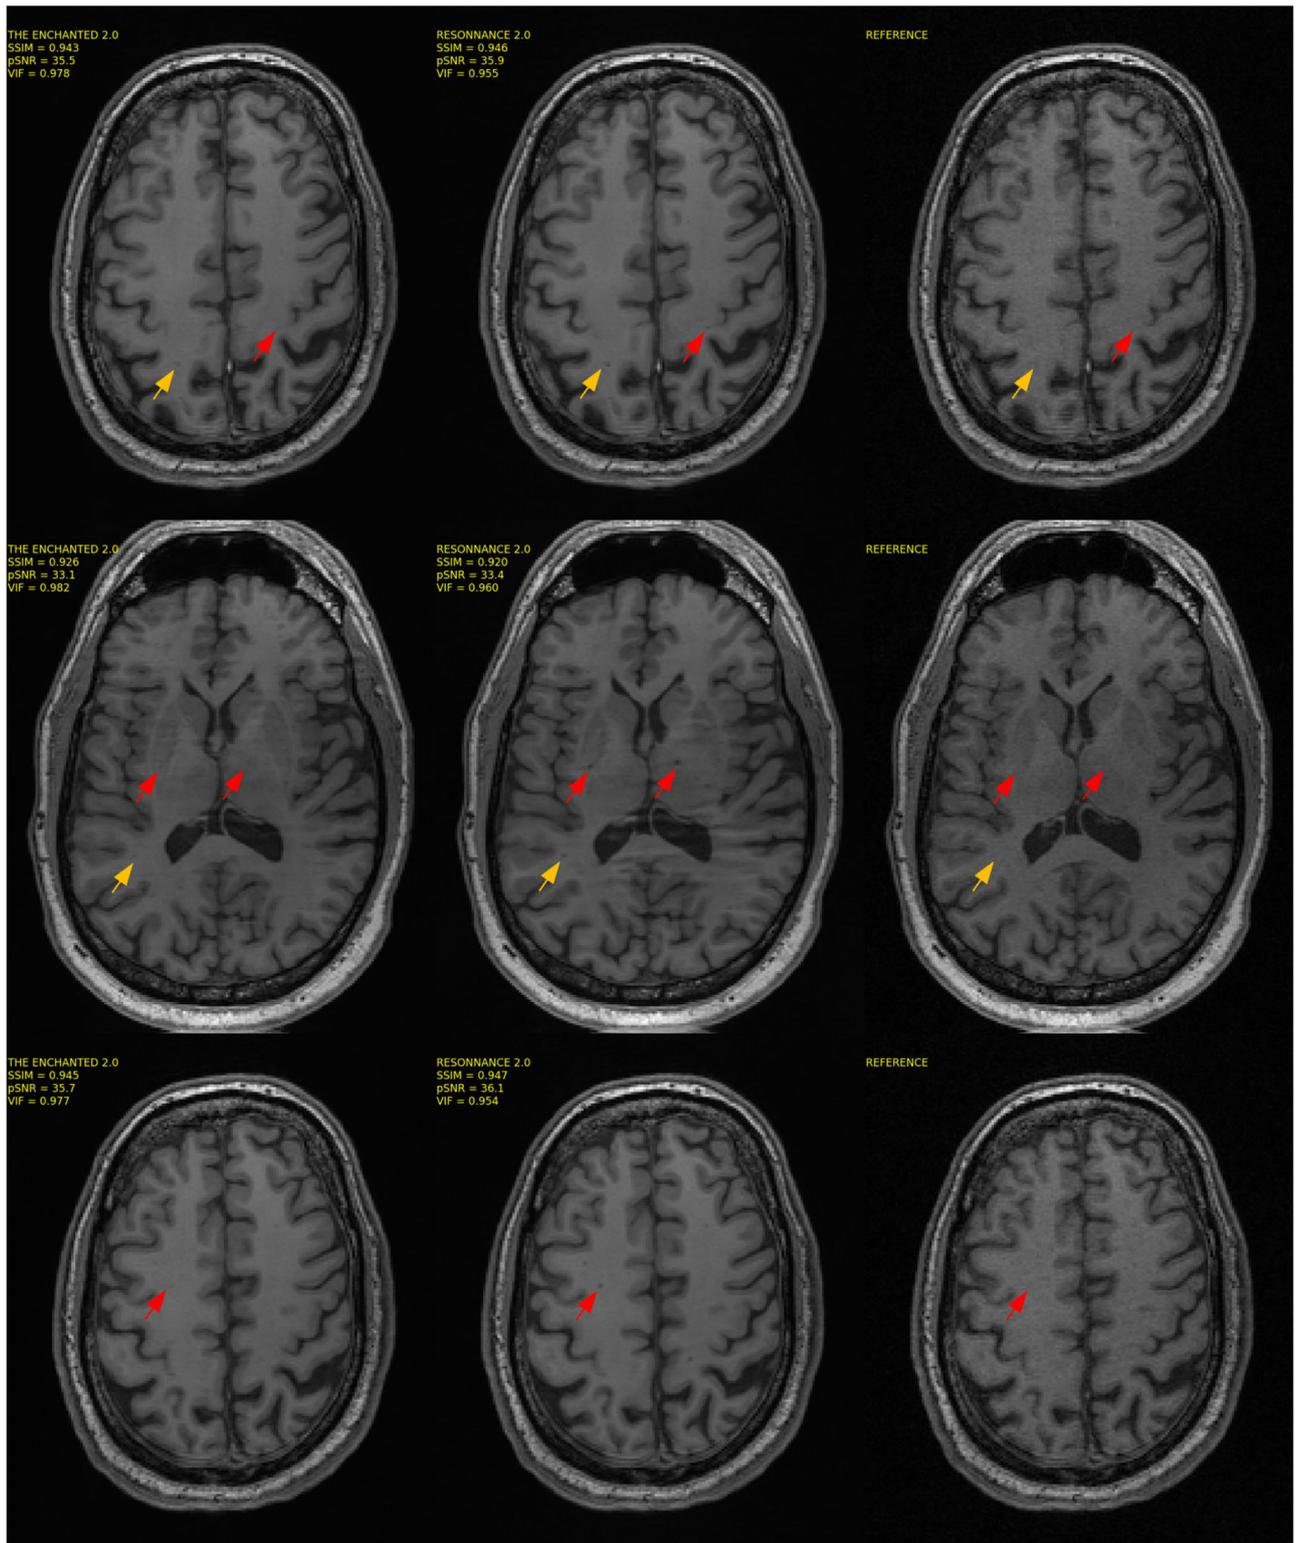


general, undesired. During the expert visual assessment, 23 of 50 (46.0%) reconstructions were rated higher than the fully sampled reference due to the fact that the brain anatomical borders in these images were preserved, but the image noise was filtered out.

All trainable baseline models and the model submitted by M-L UNICAMP used mean squared error (MSE) as their cost function. The model submitted by TUMRI was trained using a combination of multi-scale SSIM (MS-SSIM) (Wang et al., 2003) and VIF as their cost function. The model The Enchanted 1.0 has two components in their cost function: (1) their model was trained using MSE as the cost function with the target being the coil-combined complex-valued fully sampled reference and then (2) their Down-Up network (Yu et al., 2019) received as

input the absolute value of the reconstruction obtained in the previous stage, and the reference was the square-root sum-of-squares fully sampled reconstruction. The Down-Up network was trained using SSIM as the loss function. The model The Enchanted 2.0 is the only model that was pre-trained using a self-supervised learning pretext task of predicting rotations. The pretext task was trained using cross-entropy as the loss function. The main task (i.e., reconstruction task) was trained using SSIM as the loss function.

The ResoNNance 1.0 and 2.0 models used a combination of SSIM and mean absolute error (MAE) as the training loss function, which is a combination that has been shown to be effective for image restoration (Zhao et al., 2016). Because



**FIGURE 5 |** Three sample reconstructions, one per row, for the top two models. The Enchanted 2.0 and ResoNNance 2.0 and the reference are illustrated. The arrows in the figure indicate regions of interest that indicate deviations between the deep-learning-based reconstructions and the fully sampled reference.

the background in the images is quite substantial and SSIM is a bounded metric that is computed across image patches, this observation causes models trained using SSIM as part of their loss function to try to match the background noise in their reconstructions. This observation may offer a potential explanation for why the models submitted by The Enchanted and ResoNNance teams were able to preserve the noise pattern in their reconstructions. Metrics that are based on visual perception are important and evaluating the possibility of using these types of metrics as part of the loss functions is an interesting research avenue for the field of MRI reconstruction.

For  $R = 5$ , the top three models: ResoNNance 2.0, The Enchanted 2.0, and ResoNNance 1.0 produced the most visually pleasing reconstructions and also had the top performing metrics. It is important to emphasize that  $R = 5$  in the challenge is relative to the 85% of  $k$ -space that was sampled in the slice-encode ( $k_z$ ) direction. If we consider the equivalent full  $k$ -space, the acceleration factor would be  $R = 5.9$ . Based on the Track 01 results, we would say that an acceleration between 5 and 6 might be feasible to be incorporated into a clinical setting for a single-sequence MR image reconstruction model. Further analysis of the image reconstructions by a panel of radiologists is needed to better assess clinical value before achieving a definite conclusion.

The second track of the challenge compared six different reconstruction models (Tables 3, 4). The models, The Enchanted 2.0 and ResoNNance 2.0, achieved the best overall results. For the 12-channel test set (Figure 1), the results were the same as the results they obtained in Track 01 of the challenge since the models were the same. More interesting are the results for the 32-channel test set. Though the metrics for the 32-channel test set are higher than the 12-channel test set, by visually inspecting the quality of the reconstructed images, it is clear that 32-channel image reconstructions are of poorer quality compared to 12-channel reconstructions (Figure 3). In total, 28% of the 32-channel images assessed by the expert observer were deemed to have poorer quality when compared the reference against 4% of the 12-channel images rated. This fact raises concerns about the generalizability of the reconstruction models across different coils. Potential approaches to mitigate this issue is to include representative data collected with different coils in the training and validation sets or employ domain adaptation techniques (Kouw and Loog, 2019), such as data augmentation strategies, that simulate data acquired under different coil configurations, to make the models more generalizable.

Though the generalization of learned MR image reconstruction models and their potential for transfer learning has been previously assessed (Knoll et al., 2019), the results from Track 02 of our challenge indicate that there is still room for improvements. Interestingly, the model The Enchanted 2.0 is the only model that employed self-supervised learning, which seems to have had a positive impact on the model generalizability for the 32-channel test data.

One important finding that we noticed during the visual assessment of the images is that some of the reconstructed images enhanced hypointensity regions within the brain white matter, while in others images, these hypointensities were blurred out of the image (cf., Figure 5). In many cases, it was unclear

from the fully sampled reference whether this hypointensity region corresponded to noise in the image or if it indicated the presence of relevant structures, such as lesions that appear as dark spots in T1-weighted images. This finding is critical especially when targeting diseases that often present small lesions. Further investigation is necessary to determine its potential impact before the clinical adoption of these reconstruction models.

## 5. SUMMARY

The MC-MRI reconstruction challenge provided an objective benchmark for assessing brain MRI reconstruction and the generalizability of models across datasets collected with different coils using a high-resolution, 3D dataset of T1-weighted MR images. Track 01 compared ten reconstruction models and Track 02 compared six reconstruction models. The results indicated that although the quantitative metrics are higher for the test data not seen during training (i.e., 32-channel data), visual inspection indicated that these reconstructed images had poorer quality. This conclusion that current models do not generalize well across datasets collected using different coils indicates a promising research field in the coming years that is very relevant for the potential clinical adoption of deep-learning-based MR image reconstruction models. The results also indicated the difficulty of reconstructing finer details in the images, such as lacunes. The MC-MRI reconstruction challenge continues and the organizers of the benchmark will periodically incorporate more data, which will potentially allow to train deeper models. As a long-term benefit of this challenge, we expect that the adoption of these deep-learning-based MRI reconstruction models in the clinical and research environments will be streamlined.

## DATA AVAILABILITY STATEMENT

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found below: [www.ccdataset.com](http://www.ccdataset.com).

## ETHICS STATEMENT

The studies involving human participants were reviewed and approved by Conjoint Health Research Ethics Board (CHREB), which reviews applications from researchers affiliated with the Faculties of Kinesiology, Medicine and Nursing at the University of Calgary. Approval number is REB 15-1285. The patients/participants provided their written informed consent to participate in this study.

## AUTHOR CONTRIBUTIONS

YB, WL, RF, and RS were responsible for data preparation, design of the challenge, implementation of the baseline models, and writing the bulk of the manuscript. JT, DK, NM, MC, and GY were responsible for submitting the ResoNNance reconstruction models. LRo, AL, HP, and LRi were responsible for the M-L UNICAMP submission. MD, VS, FG, DV, and

SF-R submitted the TUMRI solution and helped with the VIF analysis. AK, JC, and MS submitted The Enchanted models. ML was responsible for creating the data repository and automating the extraction of metrics and challenge leader board creation. NN was the medical expert responsible for the visual assessment of the images. All authors reviewed the manuscript, provided relevant feedback across multiple rounds of reviews, and contributed to the article and approved the submitted version.

## FUNDING

RF thanks the Canadian Institutes for Health Research (CIHR, FDN-143298) for supporting the Calgary Normative Study and acquiring the raw datasets. RF and RS thank the Natural Sciences and Engineering Research Council (NSERC - RGPIN/02858-2021 and RGPIN-2021-02867) for providing ongoing operating support for this project. We also acknowledge the infrastructure funding provided by the Canada Foundation of Innovation (CFI). The organizers of

the challenge also acknowledge Nvidia for providing a Titan V Graphics Processing Unit and Amazon Web Services for providing computational infrastructure that was used by some of the teams to develop their models. DK and MC were supported by the STAIRS project under the Top Consortium for Knowledge and Innovation-Public, Private Partnership (TKI-PPP) program, co-funded by the PPP Allowance made available by Health Holland, Top Sector Life Sciences & Health. HP thanks the National Council for Scientific and Technological Development (CNPq #309330/2018-1) for the research support grant. LRI also thank the National Council for Scientific and Technological Development (CNPq #313598/2020-7) and São Paulo Research Foundation (FAPESP #2019/21964-4) for the support.

## ACKNOWLEDGMENTS

The authors would like to acknowledge that a preprint of this manuscript is available on *arXiv* (<https://arxiv.org/abs/2011.07952>).

## REFERENCES

- Akçakaya, M., Moeller, S., Weingärtner, S., and Uğurbil, K. (2019). Scan-specific robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction: database-free deep learning for fast imaging. *Magnet. Reson. Med.* 81, 439–453. doi: 10.1002/mrm.27420
- Chen, L., Bentley, P., Mori, K., Misawa, K., Fujiwara, M., and Rueckert, D. (2019). Self-supervised learning for medical image analysis using image context restoration. *Med. Image Anal.* 58, 101539. doi: 10.1016/j.media.2019.101539
- Dedmari, M. A., Conjeti, S., Estrada, S., Ehses, P., Stöcker, T., and Reuter, M. (2018). “Complex fully convolutional neural networks for MR image reconstruction,” in *International Workshop on Machine Learning for Medical Image Reconstruction* (Granada), 30–38. doi: 10.1007/978-3-030-00129-2\_4
- Deshmane, A., Gulani, V., Griswold, M. A., and Seiberlich, N. (2012). Parallel MR imaging. *J. Magnet. Reson. Imaging* 36, 55–72. doi: 10.1002/jmri.23639
- Duan, J., Schlemper, J., Qin, C., Ouyang, C., Bai, W., Biffi, C., et al. (2019). “VSNet: variable splitting network for accelerated parallel MRI reconstruction,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Shenzhen), 713–722. doi: 10.1007/978-3-030-32251-9\_78
- Eo, T., Jun, Y., Kim, T., Jang, J., Lee, H.-J., and Hwang, D. (2018a). KIKI-Net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images. *Magnet. Reson. Med.* 80, 2188–2201. doi: 10.1002/mrm.27201
- Eo, T., Shin, H., Kim, T., Jun, Y., and Hwang, D. (2018b). “Translation of 1D inverse Fourier Transform of k-space to an image based on deep learning for accelerating magnetic resonance imaging,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Granada: Springer), 241–249. doi: 10.1007/978-3-030-00928-1\_28
- Geisler, W. S. (2008). Visual perception and the statistical properties of natural scenes. *Annu. Rev. Psychol.* 59, 167–192. doi: 10.1146/annurev.psych.58.110405.085632
- Gidaris, S., Singh, P., and Komodakis, N. (2018). “Unsupervised representation learning by predicting image rotations,” in *International Conference on Learning Representations* (Vancouver, BC).
- Gözcü, B., Mahabadi, R. K., Li, Y.-H., Ilcak, E., Cukur, T., Scarlett, J., et al. (2018). Learning-based compressive MRI. *IEEE Trans. Med. Imaging* 37, 1394–1406. doi: 10.1109/TMI.2018.2832540
- Griswold, M., Jakob, P., Heidemann, R., Nittka, M., Jellus, V., Wang, J., et al. (2002). Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnet. Reson. Med.* 47, 1202–1210. doi: 10.1002/mrm.10171
- Hammernik, K., Klatzer, T., Kobler, E., Recht, M. P., Sodickson, D. K., Pock, T., et al. (2018). Learning a variational network for reconstruction of accelerated MRI data. *Magnet. Reson. Med.* 79(6):3055–3071. doi: 10.1002/mrm.26977
- Hosseini, S. A. H., Yaman, B., Moeller, S., Hong, M., and Akçakaya, M. (2020). Dense recurrent neural networks for accelerated mri: History-cognizant unrolling of optimization algorithms. *IEEE J. Select. Top. Signal Process.* 14, 1280–1291. doi: 10.1109/JSTSP.2020.3003170
- Jethi, A., Souza, R., Ram, K., and Sivaprakasam, M. (2022). “Improving fast MRI reconstructions with pretext learning in low-data regime,” in *2022 44rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)* (Glasgow: IEEE), 1–4.
- Jin, K., McCann, M., Froustey, E., and Unser, M. (2017). Deep convolutional neural network for inverse problems in imaging. *IEEE Trans. Image Process.* 26, 4509–4522. doi: 10.1109/TIP.2017.2713099
- Kalm, K., and Norris, D. (2018). Visual recency bias is explained by a mixture model of internal representations. *J. Vis.* 18, 1–1. doi: 10.1167/18.7.1
- Knoll, F., Hammernik, K., Kobler, E., Pock, T., Recht, M. P., and Sodickson, D. K. (2019). Assessment of the generalization of learned image reconstruction and the potential for transfer learning. *Magnet. Reson. Med.* 81, 116–128. doi: 10.1002/mrm.27355
- Knoll, F., Murrell, T., Sriram, A., Yakubova, N., Zbontar, J., Rabbat, M., et al. (2020). Advancing machine learning for mr image reconstruction with an open competition: overview of the 2019 fastMRI challenge. *Magnet. Reson. Med.* 84, 3054–3070. doi: 10.1002/mrm.28338
- Kouw, W. M., and Loog, M. (2019). A review of domain adaptation without target labels. *IEEE Trans. Pattern Anal. Mach. Intell.* 43, 766–785. doi: 10.1109/TPAMI.2019.2945942
- Kwon, K., Kim, D., and Park, H. (2017). A parallel MR imaging method using multilayer perceptron. *Med. Phys.* 44, 6209–6224. doi: 10.1002/mp.12600
- Larsson, E. G., Erdogmus, D., Yan, R., Principe, J. C., and Fitzsimmons, J. R. (2003). SNR-optimality of sum-of-squares reconstruction for phased-array magnetic resonance imaging. *J. Magnet. Reson.* 163, 121–123. doi: 10.1016/S1090-7807(03)00132-0
- Lee, D., Yoo, J., Tak, S., and Ye, J. C. (2018). Deep residual learning for accelerated MRI using magnitude and phase networks. *IEEE Trans. Biomed. Eng.* 65, 1985–1995. doi: 10.1109/TBME.2018.2821699
- Liang, D., Liu, B., Wang, J., and Ying, L. (2009). Accelerating SENSE using compressed sensing. *Magnet. Reson. Med.* 62, 1574–1584. doi: 10.1002/mrm.22161

- Lønning, K., Putzky, P., Sonke, J.-J., Reneman, L., Caan, M. W., and Welling, M. (2019). Recurrent inference machines for reconstructing heterogeneous MRI data. *Med. Image Anal.* 53, 64–78. doi: 10.1016/j.media.2019.01.005
- Lustig, M., Donoho, D., and Pauly, J. M. (2007). Sparse MRI: the application of compressed sensing for rapid MR imaging. *Magnet. Reson. Med.* 58, 1182–1195. doi: 10.1002/mrm.21391
- Lustig, M., Donoho, D. L., Santos, J. M., and Pauly, J. M. (2008). Compressed sensing MRI. *IEEE Signal Process. Mag.* 25, 72–82. doi: 10.1109/MSP.2007.914728
- Mardani, M., Gong, E., Cheng, J. Y., Vasanawala, S. S., Zaharchuk, G., Xing, L., et al. (2019). Deep generative adversarial neural networks for compressive sensing MRI. *IEEE Trans. Med. Imaging* 38, 167–179. doi: 10.1109/TMI.2018.2858752
- Mason, A., Rioux, J., Clarke, S. E., Costa, A., Schmidt, M., Keough, V., et al. (2019). Comparison of objective image quality metrics to expert radiologists? scoring of diagnostic quality of MR images. *IEEE Trans. Med. Imaging* 39, 1064–1072. doi: 10.1109/TMI.2019.2930338
- McCreary, C. R., Salluzzi, M., Andersen, L. B., Gobbi, D., Lauzon, L., Saad, F., et al. (2020). Calgary Normative Study: design of a prospective longitudinal study to characterise potential quantitative MR biomarkers of neurodegeneration over the adult lifespan. *BMJ Open* 10, e038120. doi: 10.1136/bmjopen-2020-038120
- Muckley, M. J., Riemenschneider, B., Radmanesh, A., Kim, S., Jeong, G., Ko, J., et al. (2021). Results of the 2020 fastMRI challenge for machine learning MR image reconstruction. *IEEE Trans. Med. Imaging* 40, 2306–2317. doi: 10.1109/TMI.2021.3075856
- Pawar, K., Chen, Z., Shah, N. J., and Egan, G. F. (2019). A deep learning framework for transforming image reconstruction into pixel classification. *IEEE Access* 7, 177690–177702. doi: 10.1109/ACCESS.2019.2959037
- Pruessmann, K., Weiger, M., Scheidegger, M., and Boesiger, P. (1999). SENSE: sensitivity encoding for fast MRI. *Magnet. Reson. Med.* 42, 952–962. doi: 10.1002/(SICI)1522-2594(199911)42:5<952::AID-MRM16>3.0.CO;2-S
- Qin, C., Schlemper, J., Caballero, J., Price, A. N., Hajnal, J. V., and Rueckert, D. (2019). Convolutional recurrent neural networks for dynamic MR image reconstruction. *IEEE Trans. Med. Imaging* 38, 280–290. doi: 10.1109/TMI.2018.2863670
- Quan, T. M., Nguyen-Duc, T., and Jeong, W.-K. (2018). Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic loss. *IEEE Trans. Med. Imaging* 37, 1488–1497. doi: 10.1109/TMI.2018.2820120
- Schlemper, J., Caballero, J., Hajnal, J. V., Price, A. N., and Rueckert, D. (2017). A deep cascade of convolutional neural networks for dynamic MR image reconstruction. *IEEE Trans. Med. Imaging* 37, 491–503. doi: 10.1109/TMI.2017.2760978
- Schlemper, J., Oksuz, I., Clough, J., Duan, J., King, A., Schanbel, J., et al. (2019). “dAUTOMAP: decomposing AUTOMAP to achieve scalability and enhance performance,” in *International Society for Magnetic Resonance in Medicine (Paris)*.
- Schlemper, J., Yang, G., Ferreira, P., Scott, A., McGill, L.-A., Khalique, Z., et al. (2018). “Stochastic deep compressive sensing for the reconstruction of diffusion tensor cardiac MRI,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 295–303. doi: 10.1007/978-3-030-00928-1\_34
- Seitzer, M., Yang, G., Schlemper, J., Oktay, O., Würfl, T., Christlein, V., et al. (2018). “Adversarial and perceptual refinement for compressed sensing MRI reconstruction,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention (Granada)*, 232–240. doi: 10.1007/978-3-030-00928-1\_27
- Sheikh, H. R., and Bovik, A. C. (2006). Image information and visual quality. *IEEE Trans. Image Process.* 15, 430–444. doi: 10.1109/TIP.2005.859378
- Simoncelli, E. P., and Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annu. Rev. Neurosci.* 24, 1193–1216. doi: 10.1146/annurev.neuro.24.1.1193
- Souza, R., Beauferris, Y., Loos, W., Lebel, R. M., and Frayne, R. (2020a). Enhanced deep-learning-based magnetic resonance image reconstruction by leveraging prior subject-specific brain imaging: Proof-of-concept using a cohort of presumed normal subjects. *IEEE J. Select. Top. Signal Process.* 14, 1126–1136. doi: 10.1109/JSTSP.2020.3001525
- Souza, R., Bento, M., Nogovitsyn, N., Chung, K. J., Loos, W., Lebel, R. M., et al. (2020b). Dual-domain cascade of U-Nets for multi-channel magnetic resonance image reconstruction. *Magnet. Reson. Imaging* 71, 140–153. doi: 10.1016/j.mri.2020.06.002
- Souza, R., and Frayne, R. (2019). “A hybrid frequency-domain/image-domain deep network for magnetic resonance image reconstruction,” in *IEEE Conference on Graphics, Patterns and Images*, 257–264. doi: 10.1109/SIBGRAPI.2019.00042
- Souza, R., Lebel, R. M., and Frayne, R. (2019). “A hybrid, dual domain, cascade of convolutional neural networks for magnetic resonance image reconstruction,” in *International Conference on Medical Imaging with Deep Learning*, 437–446.
- Souza, R., Lucena, O., Garrafa, J., Gobbi, D., Saluzzi, M., Appenzeller, S., et al. (2018). An open, multi-vendor, multi-field-strength brain MR dataset and analysis of publicly available skull stripping methods agreement. *Neuroimage* 170, 482–494. doi: 10.1016/j.neuroimage.2017.08.021
- Sriram, A., Zbontar, J., Murrell, T., Defazio, A., Zitnick, C. L., Yakubova, N., et al. (2020a). “End-to-end variational networks for accelerated MRI reconstruction,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention (Lima)*, 64–73. doi: 10.1007/978-3-030-59713-9\_7
- Sriram, A., Zbontar, J., Murrell, T., Zitnick, C. L., Defazio, A., and Sodickson, D. K. (2020b). “GrappaNet: combining parallel imaging with deep learning for multi-coil MRI reconstruction,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (IEEE)*, 14315–14322. doi: 10.1109/CVPR42600.2020.01432
- Trabelsi, C., Bilaniuk, O., Zhang, Y., Serdyuk, D., Subramanian, S., Santos, J. F., et al. (2018). “Deep complex networks,” in *International Conference on Learning Representations (Vancouver, BC)*.
- Wang, S., Su, Z., Ying, L., Peng, X., Zhu, S., Liang, F., et al. (2016). “Accelerating magnetic resonance imaging via deep learning,” in *IEEE International Symposium on Biomedical Imaging (Melbourne, VIC: IEEE)*, 514–517. doi: 10.1109/ISBI.2016.7493320
- Wang, Z., Simoncelli, E. P., and Bovik, A. C. (2003). “Multiscale structural similarity for image quality assessment,” in *The Thirty-Seventh Asilomar Conference on Signals, Systems Computers (Pacific Grove, CA)*, 1398–1402. doi: 10.1109/ACSSC.2003.1292216
- Yang, G., Yu, S., Dong, H., Slabaugh, G., Dragotti, P. L., Ye, X., et al. (2018). DAGAN: deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction. *IEEE Trans. Med. Imaging* 37, 1310–1321. doi: 10.1109/TMI.2017.2785879
- Yang, Y., Sun, J., Li, H., and Xu, Z. (2016). “Deep ADMM-Net for compressive sensing MRI,” in *Proceedings of the 30th International Conference on Neural Information Processing Systems (Red Hook, NY: Curran Associates)*, 10–18.
- Yiasemis, G., Moriakov, N., Karkaloulos, D., Caan, M., and Teuwen, J. (2022a). Direct: Deep image reconstruction toolkit. *J. Open Sour. Softw.* 7, 4278. doi: 10.21105/joss.04278537
- Yiasemis, G., Sánchez, C. I., Sonke, J.-J., and Teuwen, J. (2022b). “Recurrent variational network: a deep learning inverse problem Solver applied to the task of accelerated MRI reconstruction,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Yu, S., Park, B., and Jeong, J. (2019). “Deep iterative down-up CNN for image denoising,” in *IEEE Conference on Computer Vision and Pattern Recognition Workshops (Long Beach, CA: IEEE)*. doi: 10.1109/CVPRW.2019.00262
- Zbontar, J., Knoll, F., Sriram, A., Muckley, M. J., Bruno, M., Defazio, A., et al. (2018). fastMRI: an open dataset and benchmarks for accelerated MRI. *arXiv preprint arXiv:1811.08839*.
- Zeng, K., Yang, Y., Xiao, G., and Chen, Z. (2019). A very deep densely connected network for compressed sensing MRI. *IEEE Access* 7, 85430–85439. doi: 10.1109/ACCESS.2019.2924604
- Zhang, P., Wang, F., Xu, W., and Li, Y. (2018). “Multi-channel generative adversarial network for parallel magnetic resonance image reconstruction in k-space,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention (Granada)*, 180–188. doi: 10.1007/978-3-030-00928-1\_21
- Zhao, H., Gallo, O., Frosio, I., and Kautz, J. (2016). Loss functions for image restoration with neural networks. *IEEE Trans. Comput. Imaging* 3, 47–57. doi: 10.1109/TCI.2016.2644865
- Zhou, B., and Zhou, S. K. (2020). “DuDoRNet: learning a dual-domain recurrent network for fast MRI reconstruction with deep T1 prior,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4273–4282. doi: 10.1109/CVPR42600.2020.00433

- Zhou, W., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* 13, 600–612. doi: 10.1109/TIP.2003.819861
- Zhu, B., Liu, J. Z., Cauley, S. F., Rosen, B. R., and Rosen, M. S. (2018). Image reconstruction by domain-transform manifold learning. *Nature* 555, 487–492. doi: 10.1038/nature25988

**Conflict of Interest:** MC is a shareholder of Nico.lab International Ltd.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

**Publisher's Note:** All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of

the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

*Copyright © 2022 Beauferris, Teuwen, Karkalousos, Moriakov, Caan, Yiasemis, Rodrigues, Lopes, Pedrini, Rittner, Dannecker, Studenyak, Gröger, Vyas, Faghieh-Roohi, Kumar Jethi, Chandra Raju, Sivaprakasam, Lasby, Nogovitsyn, Loos, Frayne and Souza. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.*