

HyperCoil-Recon: A Hypernetwork-based Adaptive Coil Configuration Task Switching Network for MRI Reconstruction

Sriprabha Ramanarayanan^{1,2}, Mohammad Al Fahim¹, Rahul G. S.^{1,2}, Amrit Kumar Jethi¹, Keerthi Ram², and Mohanasankar Sivaprakasam^{1,2}

¹Indian Institute of Technology Madras (IITM), India

²Healthcare Technology Innovation Centre, IITM, India

Abstract

Parallel imaging, a fast MRI technique, involves dynamic adjustments based on the configuration i.e. number, positioning, and sensitivity of the coils with respect to the anatomy under study. Conventional deep learning-based image reconstruction models have to be trained or fine-tuned for each configuration, posing a barrier to clinical translation, given the lack of computational resources and machine learning expertise for clinicians to train models at deployment. Joint training on diverse datasets learns a single weight set that might underfit to deviated configurations. We propose, HyperCoil-Recon, a hypernetwork-based coil configuration task-switching network for multi-coil MRI reconstruction that encodes varying configurations of the numbers of coils in a multi-tasking perspective, posing each configuration as a task. The hypernetworks infer and embed task-specific weights into the reconstruction network, 1) effectively utilizing the contextual knowledge of common and varying image features among the various fields-of-view of the coils, and 2) enabling generality to unseen configurations at test time. Experiments reveal that our approach 1) adapts on the fly to various unseen configurations up to 32 coils when trained on lower numbers (i.e. 7 to 11) of randomly varying coils, and to 120 deviated unseen configurations when trained on 18 configurations in a single model, 2) matches the performance of coil configuration-specific models, and 3) outperforms configuration-invariant models with improvement margins of ~ 1 dB/0.03 and 0.3 dB/0.02 in PSNR / SSIM for knee and brain data. Our code is available at <https://github.com/sriprabhar/HyperCoil-Recon>

1. Introduction

Parallel Imaging (PI) is a widely used technique in reducing the acquisition time of Magnetic Resonance Imaging (MRI) [1]. Modern MRI scanners adopt PI as the default option to image the patient’s anatomy by acquiring frequency-domain (or k-space) measurements using an in-

strument called a receiver coil. They employ multiple receiver coils that simultaneously obtain under-sampled k-space data of different views of the anatomy being imaged to speed up the acquisition [2]. The main characteristics of PI are: 1) The captured views change for each scan and are dependent on the configuration (number and positioning) of the coils with respect to the anatomy being imaged [36]. 2) Different coils are typically sensitive to different but overlapping regions based on their interaction with the anatomy [36]. These factors indicate that PI is not only diversified by the multimodal nature of the MRI data [44, 32], (e.g. various contrasts) but also characterized by dynamic contextual adjustments in the coil configurations at scan time. Recently, deep learning methods have shown promising results over conventional methods like SENSE [28] and GRAPPA [14] to reconstruct images from under-sampled multi-coil k-space owing to their capabilities to learn complex representations from the data [25]. However, despite their success, standard deep neural networks (DNNs) remain restrictive under two training conditions.

1. Context-invariant or joint training: Conventional DNN models, when presented with images from diverse acquisition contexts while training, employ a naive data-loading logic that combines and shuffles the images from the contexts. This joint training process learns a fixed set of weights that contain features common across all the contexts considered during training [21, 23] (Figure 1). However, due to the diversity and complementarity of the multi-coil MRI, different coil configurations representing different numbers of coils, exhibit heterogeneous visual characteristics. As a result, the shared features learned using a single set of weights could underfit contexts that are deviated from the training data during inference [33]. For instance, the MIDL Calgary Campinas Multi-coil MRI reconstruction challenge [3] takes a clinical scenario of adding more coils in a scanner and raises concern about the generalizability of the DNN reconstruction models to the 32-coil dataset, when trained on the 12-coil dataset for a given anatomy. It is noted that 28% of the images with unseen coil configura-

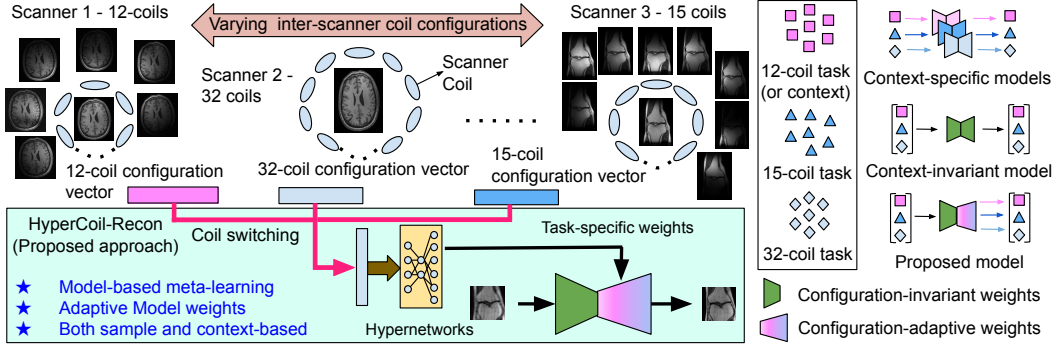


Figure 1: The hypernetwork-based coil-configuration task switching model for adaptive MRI reconstruction. Task-specific models need training for every coil configuration, while joint training has single shared weights, HyperCoil-Recon infers task-adaptive weights for the reconstruction network, enabling generalization to several unseen contexts without retraining.

tions at test time, assessed by expert observers were deemed to have poorer image quality when compared with the performance of the trained reference model [3].

2. **Context-specific training:** The architectures used to solve the MRI reconstruction tasks are often similar. Yet, the models have to be trained separately for a dataset with different acquisition settings (different coil configuration, anatomies, and contrasts) [8, 31] (Figure 1). Methods like test-time training adapt the model to the target distribution by back-propagation using an unsupervised loss on complete test data distribution [39]. These approaches might be infeasible for clinical translation, given the lack of large computational resources or machine-learning expertise for clinicians to train models at deployment time [8]. As PI involves overlapping fields-of-view of different coils for a given scan [9], different coil-switching configurations obtained from the presence or absence of various coils carry both common and varying image domain features among them. Utilizing this knowledge motivates the need for a meta-model that dynamically changes across varying coil configurations, enabling generality to new configurations in a zero-shot setting, without model updates at deployment.

Inspired by hypernetworks [15, 40], a recently emerged deep learning technique in providing adaptability, information sharing, and data-efficiency in DNNs for multi-tasking [6], we propose a clinically motivated setting for multi-coil MRI reconstruction, called the coil-configuration task switching neural network. We pose each coil configuration as a task that encodes varying coil-switching combinations *i.e.* the presence and absence of one or more coils in the multi-coil MRI data. Our approach employs small task-conditioned hypernetworks as meta-learners [17] that infer task-specific weights and embed these latent representations of various coil configurations into an encoder-decoder-based network for MR image reconstruction (Figure 1). The encoder is task-invariant and learns shared features across coil configurations while the decoder is task-adaptive. The

hypernetworks 1) infuse the coil configuration-specific embeddings in all layers of the decoder to integrate both task-specific and task-invariant features, 2) offer interesting insight into the relationship between various coil-configuration tasks, and 3) enable the model to adapt, on the fly, to various unseen coil configurations across contrasts, anatomies, and datasets in a single forward pass during inference. Our contributions are,

1. We propose HyperCoil-Recon, an adaptive coil configuration task-switching network for test-time on-the-fly adaptation to varying coils in multi-coil MR image reconstruction. The proposed method adopts a multi-tasking approach by posing each coil configuration as a task and learning from both task embeddings and under-sampled images.

2. As the first endeavor to facilitate task generalization in multi-coil MRI reconstruction, we focus on analyzing the capabilities of such a unified model when combining multiple contrasts and anatomies and presenting insights into the relationship between various coil configurations.

3. Our method uses hypernetworks as meta-learners to enrich the training process and embeds coil-configuration-specific information into the reconstruction network.

4. Our experiments reveal that our approach (i) adapts on the fly to various unseen configurations up to 32 coils when trained on lower numbers (*i.e.* 7 to 11) of randomly varying coils, generalizes to 120 deviated unseen configurations when trained on 18 configurations, in a single unified model, (ii) matches the performance of coil configuration-specific models, and (iii) outperforms configuration-invariant models with improvement margins of ~ 1 dB / 0.03 and 0.3 dB / 0.02 in PSNR / SSIM for knee and brain anatomies, respectively.

2. Related Work

Hypernetworks: Hypernetworks are neural networks that generate weights for another neural network, known as the

target or primary task-oriented network. Hypernetworks have shown promising results in a variety of deep learning problems, including transfer learning [4], continual [40] and meta-learning [41], causal inference [5], and neural architecture search [4]. In image restoration tasks, hypernetworks are used in decoupled learning [12] to solve multi-parameterized imaging operators and controllable image restoration using a single network [7, 43]. The decoupled learning methods lack the inductive bias of convolution layers of the primary network to solve the imaging tasks, as the hypernetworks are the main source of weights for the primary network. In our approach, both the primary network and the hypernetworks are jointly learned, enabling both coil context-invariant image features and context-specific semantic features. The controllable image processing networks adopt hypernetworks as a tuning module to adjust the input parameters of the imaging operator (*e.g.* scale factor in image super-resolution). Different from this approach which uses a single parameter, in our work, the hypernetworks are conditioned with comprehensive task embeddings of the coil switching to infer task-specific weights. The hypernetworks have non-linear leaky ReLU layers, which prevent zero activations for negative layer weights, enabling expressivity to numerous contexts.

Adaptive MRI reconstruction networks: Several task-specific convolutional neural networks (CNNs) have been developed for MRI reconstruction, namely unrolled networks [38, 30, 35], attention mechanism [19], transformer-based networks [18], and variational networks [45]. Adaptive single-coil reconstruction networks include, (i) MAC-ReconNet [31], and Hyper-Recon [42], based on decouple learning, wherein the hypernetworks are driven based on scanner information and regularization hyper-parameters, respectively, and (ii) the universal under-sampled MRI reconstruction [23] and side-information guided networks [24], based on adaptive instance normalization (AdaIN) [37]. MAC-ReconNet uses layer-wise linear hypernetworks with limited-sized input layers for a deeper base network. In ours, we adopt wider non-linear hypernetworks to learn task relationships and they are not heavy due to the presence of a hidden bottleneck embedding layer. The number of hypernetworks depends only on the sub-sampling levels and not on the number of layers in the primary network. Unlike AdaIN, the hypernetworks in our method are more expressive [22] due to dense multiplicative interactions [20] via convolution of the dynamic weights with the CNN features.

3. Method

Problem Formulation: According to compressed sensing for MRI, the problem of recovering the desired complex-valued MR image, $x \in \mathbb{C}^N$, from the under-sampled (US) k-space measurements, $y \in \mathbb{C}^{M'}$ is ill-posed as $M' \ll N$ [27] and the optimization formulation is,

$$\min_x \sum_{j=1}^{n_c} \|MFS_jx - y_j\|_2^2 + \mathcal{R}(x) \quad (1)$$

Here, $\mathcal{R}(x)$ is the sparse regularization term, M , the 2-D under-sampling mask, F , the Fourier transform matrix, and S_j , $j = 1, 2, \dots, n_c$, the sensitivity maps of the n_c receiver coils. The proposed deep learning-based MRI reconstruction involves training a deep learning (DL) model on the average loss of all observed data samples from N_γ datasets. Each dataset, corresponding to the coil configuration task γ_i , consists of US or zero-filled (ZF) image input, and the fully-sampled (FS) image target pairs denoted as $D_i = (x_{US,i}, x_{FS,i})$, where $i = 1, 2, \dots, N_\gamma$ is the task index. This supervised task-aware joint training of the neural network, $f(x_{US,i}, \vec{\gamma}_i; \theta, \phi_H)$ consisting of the image reconstruction CNN, $f_{CNN}(x_{US}; \theta)$ with parameters θ and the hypernetwork, $f_H(\vec{\gamma}_i; \phi_H)$ with parameters ϕ , across multi-coil configuration tasks, γ_i is given by,

$$\theta^*, \phi_H^* = \underset{\theta, \phi_H}{\operatorname{argmin}} \mathbb{E}_{(x_{US,i}, x_{FS,i}) \in \bigcup_i \mathcal{D}_i} [\|x_{FS,i} - f(x_{US,i}, \vec{\gamma}_i; \theta, \phi_H)\|_2^2] \quad (2)$$

Here, θ^* and ϕ_H^* are the optimized weights of the networks. The configuration vector input $\vec{\gamma}_i$ to the hypernetwork is a binary vector that enumerates the presence (binary 1) and absence (binary 0) of a coil (*i.e.* the corresponding sensitivity map and the k-space). In order to optimize Eq. 1 efficiently, the variable splitting method in VS-Net [10] introduces auxiliary splitting variables $u \in \mathbb{C}^N$ and $\{x_j \in \mathbb{C}^N\}_{j=1}^{n_c}$ and derives the final solution as follows:

$$\begin{aligned} u^{k+1} &= \operatorname{denoiser}(m^k) \\ x_j^{k+1} &= F^{-1}((\lambda M^T M + \alpha I)^{-1}(\alpha F S_j m^k + \lambda M^T y_j)) \\ m^{k+1} &= (\beta I + \alpha \sum_{j=1}^{n_c} S_j^H S_j)^{-1}(\beta u^{k+1} + \alpha \sum_{j=1}^{n_c} S_j^H x_j^{k+1}) \end{aligned} \quad (3)$$

The top equation converts the original problem (Eq. 1) to a denoising problem as in Eq. 2. The middle equation provides the data consistency to k-space for each coil. The bottom equation computes a weighted average of the results obtained from the first two equations. The proposed model (Fig. 2) for multi-coil MRI reconstruction follows the iterative setup formulated in Eq. 3 with N_b cascades of the three blocks: HyperCoil-Recon as denoiser block, data consistency block (DCB) and weighted average block (WAB). The DL model takes in the sensitivity-weighted US image ($m^0 = \sum_{j=1}^{n_c} S_j^H F^{-1} M^T y_j$) and its corresponding binary configuration vector as inputs. DCB uses the binary sampling mask (M), the under-sampled k-space data ($\{M^T y_j\}_{j=1}^{n_c}$) with the corresponding sensitivity maps for

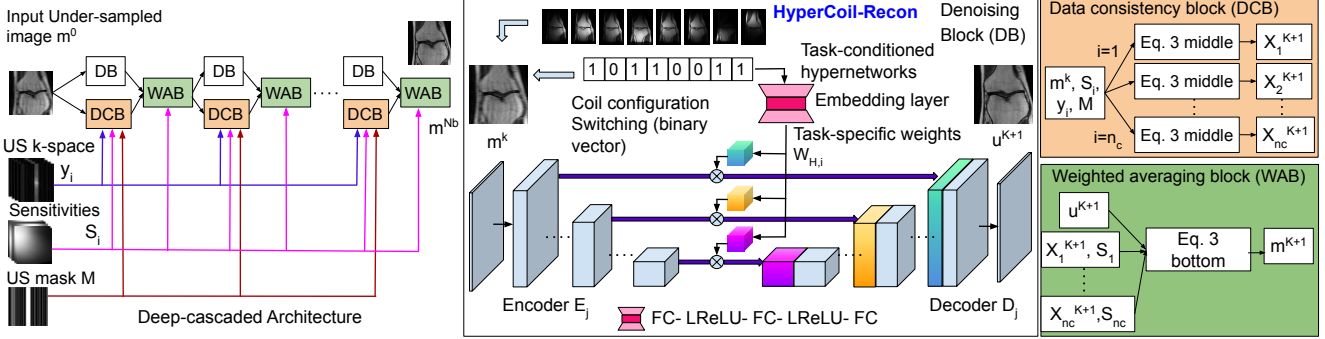


Figure 2: Deep Cascaded HyperCoil-Recon Architecture with the denoising (the hypernetworks and the reconstruction network), data consistency, and weighted-average blocks. FC - Fully connected layer, LReLU - Leaky ReLU non-linearity.

which the coil configuration switching vector has binary 1 value to provide k-space data consistency. WAB uses coil sensitivity maps to perform weighted averaging of the DCB outputs and combine it with the DL model output.

Architecture Details: The primary reconstruction CNN is the U-Net [34]. There are N_{sub} hypernetworks that take the coil configuration task vector as input and infer task-specific weights $W_{H,i}$ corresponding to the N_{sub} sub-sampling levels of the decoder in the primary CNN. The encoder features $E_1, E_2, \dots, E_{N_{sub}}$ of the CNN are convolved (shown as \otimes) with the N_{sub} task-specific weight vectors to obtain task-specific features, that are concatenated with the decoder features $D_1, D_2, \dots, D_{N_{sub}}$ at each level and the reconstructed image $x_{CNN,i}$ for the task, γ_i , is obtained (Fig. 2).

$$\begin{aligned}
 W_{H,i} &= (W_{H,1,i}, \dots, W_{H,N_{sub},i}) = f_H(\vec{\gamma}_i; \phi_H) \\
 F_{Dec,j,i} &= (E_j \otimes W_{H,j,i}) \parallel D_j, \forall j = 1, \dots, N_{sub} \\
 x_{CNN,i} &= Dec(F_{Dec,1,i}, \dots, F_{Dec,N_{sub},i})
 \end{aligned} \quad (4)$$

4. Experiments

4.1. Datasets and Implementation Details

Datasets: 1) The **Knee multi-coil dataset** [16] consists of coronal proton-density without (PD) and with fat suppression (PDFS) images, acquired from 20 patients. Each patient data has 20 slices of size 640×368 with 15 channels (coils) and sensitivity maps. We split the data into 10 slices with 200 slices each, for training and validation with Cartesian under-sampling. 2) **Calgary Campinas Challenge multi-coil dataset** [3]: This large-scale publicly accessible dataset provides k-space data from 167 3D, T1-weighted, gradient-recalled echo, 1 mm^3 isotropic sagittal brain scans collected on a clinical 3-T MRI scanner (Discovery MR750; GE Healthcare). There are two datasets, one with 12-channel (117 scans) and the other with 32-channel (50 scans) receiver coils, each with 170 to 180 contiguous slices of size 256×218 . The dataset consists of Poisson under-sampling masks for 5x and 10x acceleration.

Implementation Details: We represent each coil configuration task vector as a binary vector that encodes the coil switching. For example, in a 15-coil dataset, a 9-coil task switching vector, 111001010101110 denotes that the sensitivities and k-space data of coils 4, 5, 7, 9, 11, and 15 are absent while the rest are present. For a given number of coils (task), 7, 9, 11, and 12, we augment the task with several randomly varying combinations (sub-tasks) of binary-valued vectors, embed them in a 32-bit vector (maximum embedding vector size) initialized with 1's and feed to the hypernetworks to support multiple configurations in one training. We have implemented the models using Pytorch v1.12, trained for 100 epochs with 5 cascades on a 24 GB Nvidia RTX-3090 and L1 as the loss function between the predicted and the ground truth (GT) fully sampled images. Our evaluation metrics are PSNR and SSIM measures.

4.2. Results and discussion

Our experiments include, 1) Generalization to unseen coil configurations when trained on few configurations, 2) Task relationship, 3) Performance Comparison with other multi-coil MRI reconstruction architectures on large-scale clinical datasets, 4) Comparison with other adaptive MRI reconstruction methods for multi-modal acquisition contexts, and 5) An ablative study.

4.2.1 Generalization to Unseen Coil Configurations

We train the models on a few (18) tasks and assess the generality of the models to several (120) deviated unseen tasks that share the same label space. From the 15-coil knee dataset, we create the 18 training tasks or configurations using combinations of 7, 10, and 12 coils, PD and PDFS contrasts, and 4x, 5x, and 8x acceleration factors. We compare our model with the jointly trained model using only the images (coil-configuration task-invariant model or CCTIM) and models trained for a specific number of coils (coil-configuration task-specific model or CCTSM).

Table 1: Quantitative comparison of the generality of the proposed model for unseen configurations with CCTIM and CCTSM when combining 7, 10, & 12 coils, 4, 5, & 8x acceleration, and contrasts - PD & PDFS. ZF - Zero-filled (*i.e.* US) image. For *e.g.*, the task PD7 - proton density MRI with 7x acceleration. Green and blue - the first and second-best metrics, respectively.

| Coils | Scanner context | ZF | CCTIM | HyperCoil-Recon | CCTSM | Coils | Scanner context | ZF | CCTIM | HyperCoil-Recon | CCTSM |
|-------|-----------------|------------|------------|-----------------|------------|-------|-----------------|------------|-------------|-----------------|------------|
| | | PSNR/SSIM | PSNR/SSIM | PSNR/SSIM | PSNR/SSIM | | | PSNR/SSIM | PSNR/SSIM | PSNR/SSIM | PSNR/SSIM |
| 7 | PD7 | 18.47/.717 | 28.64/.780 | 29.38/.821 | 27.93/.768 | 13 | PD7 | 25.44/.822 | 30.80/.799 | 31.53/.862 | 30.88/.845 |
| | PD9 | 18.46/.714 | 28.28/.770 | 29.03/.814 | 27.58/.756 | | PD9 | 25.32/.816 | 30.35/.791 | 31.00/.853 | 30.44/.837 |
| | PDFS7 | 21.91/.602 | 29.85/.733 | 30.30/.746 | 29.76/.732 | | PDFS7 | 28.30/.760 | 31.90/.766 | 32.39/.789 | 32.07/.782 |
| | PDFS9 | 21.89/.598 | 29.56/.726 | 30.00/.740 | 29.49/.725 | | PDFS9 | 28.20/.753 | 31.52/.757 | 32.04/.781 | 31.73/.774 |
| 9 | PD7 | 20.17/.761 | 29.61/.798 | 30.20/.835 | 30.19/.828 | 14 | PD7 | 27.01/.823 | 30.87/.798 | 31.52/.861 | 31.79/.847 |
| | PD9 | 20.14/.757 | 29.21/.789 | 29.78/.828 | 29.79/.818 | | PD9 | 26.82/.823 | 30.41/.789 | 30.99/.852 | 31.25/.838 |
| | PDFS7 | 23.57/.665 | 30.90/.752 | 31.23/.766 | 31.27/.768 | | PDFS7 | 29.75/.775 | 31.93/.766 | 32.46/.790 | 32.58/.784 |
| | PDFS9 | 23.54/.660 | 30.59/.746 | 30.86/.759 | 31.00/.760 | | PDFS9 | 29.60/.768 | 31.55/.757 | 32.01/.784 | 32.17/.776 |
| 11 | PD7 | 22.44/.797 | 30.41/.800 | 30.83/.843 | 30.37/.818 | 15 | PD7 | 28.88/.836 | 30.87/.797 | 31.35/.861 | 32.29/.885 |
| | PD9 | 22.38/.792 | 29.98/.792 | 30.32/.833 | 29.95/.810 | | PD9 | 28.61/.829 | 30.42/.788 | 30.81/.852 | 31.68/.876 |
| | PDFS7 | 25.75/.719 | 31.59/.762 | 31.90/.778 | 31.56/.761 | | PDFS7 | 31.20/.786 | 31.90/0.764 | 32.43/.793 | 32.91/.803 |
| | PDFS9 | 25.69/.713 | 31.25/.755 | 31.50/.771 | 31.25/.753 | | PDFS9 | 31.00/.779 | 31.50/0.756 | 32.04/.785 | 32.46/.794 |




Figure 3: Left & Middle: Plots showing more unseen coil configurations when combining 7, 10, and 12 as the number of coils, acceleration factors 4x, 5x, and 8x, and contrasts - PD and PDFS. Right: Matrix plot showing the inter-task relationship. Tasks with neighboring coil configurations exhibit more similarity, while far-apart configurations exhibit lesser similarity.

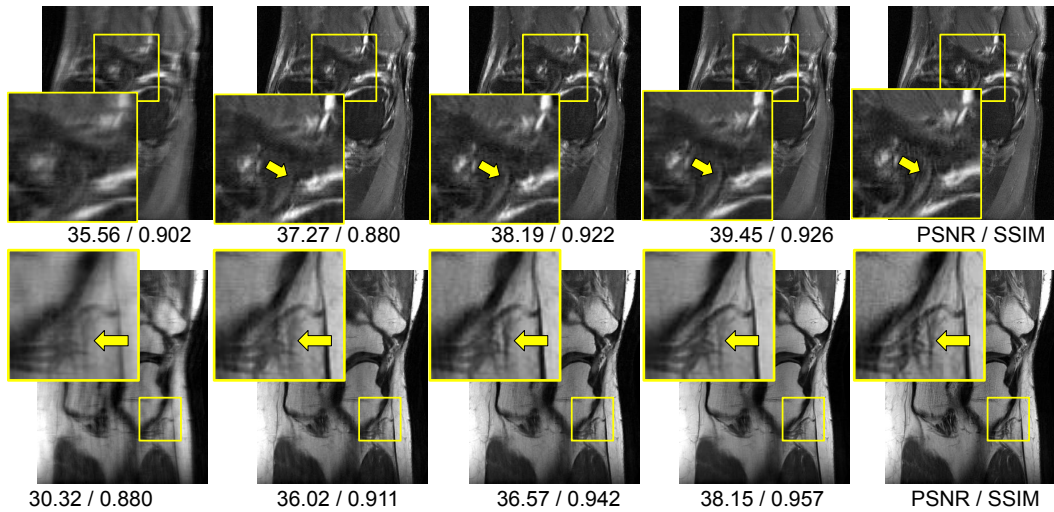


Figure 4: From the left: Qualitative comparison of the images of the ZF, CCTIM, HyperCoil-Recon, and CCTSM for an unseen case with 15 coils and 4x acceleration when trained on 7, 10, and 12 coils, 4x, 5x, and 8x acceleration and contrasts - PD (bottom) and PDFS (top). The quality of the proposed model is better than CCTIM and matches that of CCTSM.

Table 1 shows the performance of the models for 24 unseen coil configurations with unseen acceleration factors (7x and 9x). From the table, our observations are as follows. (i)

HyperCoil-Recon consistently performs better than the CCTIM for all the coil configurations. (ii) HyperCoil-Recon outperforms CCTSM for almost all the configurations in

Table 2: Quantitative comparison of HyperCoil-Recon with other multi-coil MRI reconstruction methods large-scale clinical data. The column pairs are the evaluation results of the 7-9-11 model on the 12-coil unseen task (same dataset), the 7-9-11 model on the 32-coil unseen task (unseen dataset), and the 15-28-31 model on the 32-coil task (reference for column 2).

| Method | 12-coil testing on 7-9-11 model (seen dataset & unseen coils) | | 32-coil testing on 7-9-11 model (unseen dataset and unseen coils) | | 32-coil on 15-28-31 model (seen dataset & unseen coils) | |
|----------------------|--|---------------|--|---------------|--|---------------|
| | 5x | 10x | 5x | 10x | 5x | 10x |
| | PSNR/ SSIM | PSNR/ SSIM | PSNR/ SSIM | PSNR/ SSIM | PSNR/ SSIM | PSNR/ SSIM |
| ZF | 25.20 / 0.726 | 23.70 / 0.632 | 26.40 / 0.780 | 25.50 / 0.745 | 26.40 / 0.780 | 25.50 / 0.745 |
| VS-Net [10] | 27.61 / 0.862 | 26.08 / 0.802 | 19.03 / 0.707 | 18.61 / 0.671 | 31.04 / 0.921 | 29.39 / 0.869 |
| VarNet [16] | 26.06 / 0.839 | 24.57 / 0.789 | 19.19 / 0.673 | 17.85 / 0.615 | 28.99 / 0.885 | 27.49 / 0.828 |
| RecurrentVarNet [45] | 29.17 / 0.908 | 27.46 / 0.855 | 18.81 / 0.707 | 18.50 / 0.675 | 31.83 / 0.925 | 30.82 / 0.893 |
| KIKI-Net [11] | 31.13 / 0.911 | 29.24 / 0.865 | 23.25 / 0.853 | 22.50 / 0.815 | 32.44 / 0.943 | 30.53 / 0.896 |
| DC-Unet [38] | 31.42 / 0.914 | 29.14 / 0.865 | 23.13 / 0.859 | 22.90 / 0.824 | 33.92 / 0.948 | 31.72 / 0.908 |
| SWin [18] | 29.26 / 0.895 | 27.33 / 0.834 | 19.68 / 0.756 | 19.43 / 0.721 | 32.48 / 0.939 | 30.39 / 0.889 |
| EnchantedNet [3] | 31.29 / 0.912 | 29.00 / 0.861 | 26.63 / 0.887 | 26.47 / 0.846 | 33.29 / 0.937 | 30.75 / 0.889 |
| HyperCoil-Recon | 33.54 / 0.928 | 30.39 / 0.878 | 31.15 / 0.924 | 29.55 / 0.882 | 36.13 / 0.955 | 32.52 / 0.910 |

SSIM and performs competitively in PSNR. (iii) In some cases, lower (*e.g.* 13-coil) coil tasks of HyperCoil-Recon are better than higher (*e.g.* 14-coil) coil tasks of CCTSM. Joint training in CCTIM lacks context awareness as the model learns a single set of shared weights from the image samples only. Although CCTSMs operate as expert models for higher coil configurations, the knowledge gained using only the primary network is still inadequate at lower coil configurations wherein the receiving coils might not capture the anatomy of interest adequately. Furthermore, at higher acceleration factors, the CCTIM becomes comparable to ZF performance. HyperCoil-Recon benefits from 1) the low-level image features of the primary CNN, and 2) additional knowledge of context embeddings via the task-switching hypernetworks, which help to build the relationship between various contexts and enables the model to interpolate and extrapolate to unseen contexts between known contexts. Figure 3 shows more unseen coil configurations wherein our model generalizes better than CCTIM and CCTSM. The visual results for HyperCoil-Recon (Figure 4) show better recovery of details over CCTIM and CCTSM.

4.2.2 Understanding Task Relationships

We interpret the role of the hypernetworks when jointly training with the primary reconstruction network with a multi-tasking objective to study if the embeddings learned by the hypernetworks carry information useful for task relationships. To analyze this perspective, we compute the similarity measure, SIM between two task embedding layer output vectors τ_i and τ_j (second embedding hidden layer after LReLU in the hypernetwork shown in Fig. 2) based on cosine similarity as $SIM(\tau_i, \tau_j) = 1 - \frac{\tau_i^T \tau_j}{\|\tau_i\| \|\tau_j\|}$.

Figure 3 (right) shows the matrix plot of the similarity between tasks. The darker the blue color region, the better the similarity between the tasks (*i.e.* SIM less than 0.4) corresponding to each row-column pair. Each task denotes the number of coils used for reconstruction. We make two interesting observations. 1. The tasks that are closer in

the number of coils exhibit higher similarity (*e.g.*, Task 12 (seen) and 15 (unseen)) and those farther exhibit lesser similarity (Task 7 with Tasks 11 to 15). 2. Task similarity is denser in the region with a higher number of coils. We believe that when coils share common features, the tasks have common patterns in the conditioning, and the task embeddings might be localized closer to each other in the latent embedding space. This indicates that the hypernetworks share the knowledge across the tasks, facilitating adaptive coil configuration in a single unified model.

4.2.3 Feature Reusability

We evaluate the models on the Calgary-Campinas large-scale clinical datasets of the brain anatomy. The goal of this study is to provide an objective benchmark of various multi-coil MRI reconstruction methods with respect to the generalizability of the models across the 32-coil and the 12-coil datasets. For each reconstruction method shown in Table 2, we train two models - 1) Training with combinations of 7, 9, and 11 coils using the 12-coil Calgary dataset. 2) Training with combinations of 15, 28, and 31 coils using the 32-coil Calgary dataset. We refer to these models as the 7-9-11 and 15-28-31 models, respectively. Table 2 shows three evaluation cases, 1) the 7-9-11 model on 12-coil unseen configuration within the same dataset, 2) the 7-9-11 model on the 32-coil unseen task on the 32-coil unseen dataset, and 3) the 15-28-31 model on 32-coil configuration within the same dataset. We take the 15-28-31 model of each method as the reference model to assess the performance of the corresponding 7-9-11 model in Case 2. From the table, our observations are as follows. 1) The proposed approach performs better than other reconstruction methods for all three cases. 2) The HyperCoil-Recon 7-9-11 model is comparable with the corresponding 15-28-31 model on the 32-coil task, while other models perform poorly. These observations reveal that the task-conditioned hypernetworks in the proposed model exhibit meta-learning [13] capabilities wherein the hypernetworks contain reusable features




Figure 5: Qualitative comparison of the HyperCoil-Recon with other multi-coil MRI reconstruction architectures for the unseen 12-coil task using the 7-9-11 model of the same dataset. Top: ZF image, VS-Net [10], VarNet [16], Recurrent VarNet [45], KIKI-Net [11]. Bottom: SWin [18], DC-UNet [38], EnchantedNet [3], HyperCoil-Recon, and GT image.




Figure 6: Qualitative comparison of the HyperCoil-Recon with other methods for the unseen 32-coil Calgary dataset with unseen 32-coil configuration on the 7-9-11 model. Top: ZF image, VS-Net, VarNet, Recurrent VarNet, KIKI-Net. Bottom: SWin, DC-UNet, EnchantedNet, HyperCoil-Recon, HyperCoil-Recon (15-28-31 reference model), and GT image.

[29, 32] for new tasks deviated from the training tasks. Figure 5 shows that the reconstructed images of our model exhibit better visual quality compared to other methods for a case with the highlighted region pointing to faint structures containing subtle image gradients. Figure 6 also shows that our method gives better quality for 32 coils over other methods and is on par with the 32-coil reference model.

4.2.4 Expressivity in Multi-modal Scenario

We combine multimodal contexts in a single model with multiple anatomies, PD knee and T1 brain, and multiple coil configurations (7, 10, and 12 coils), with 5x undersampling. We compare our model with other adaptive MRI

reconstruction [35] architectures - MAC-ReconNet [31] and AdaIN [23]. Table 3 and Figure 7 show the results for the brain and knee for multiple coil configurations. We note that HyperCoil-Recon performs better than MAC-ReconNet and AdaIN for almost all the coil configurations both in PSNR and SSIM. MAC-ReconNet lacks task-invariant low-level image features as the dynamic weight prediction hypernetworks infer all the weights of the primary reconstruction network. On the other hand, AdaIN normalizes the features to compensate for the distribution shift using scale and shift operations. The hypernetworks in our model exhibit dense multiplicative interactions [20] via convolution operations between predicted task-specific weights and the CNN features. These interactions provide comprehensive

Table 3: Quantitative comparison of the HyperCoil-Recon with other adaptive MRI reconstruction methods - MAC-ReconNet [31] and AdaIN [23] under multi-modal scenario when combining (during training) different anatomies (12-coil T1 brain and 15-coil PD knee), with different contrast and different coil configurations (7, 10 and 12 coils respectively).

| Coils | Brain 12 coils | | | Coils | Knee 15 coils | | |
|-------|-------------------|---------------|-----------------|-------|-------------------|---------------|-----------------|
| | MAC-ReconNet [31] | AdaIN [23] | HyperCoil-Recon | | MAC-ReconNet [31] | AdaIN [23] | HyperCoil-Recon |
| | PSNR / SSIM | PSNR / SSIM | PSNR / SSIM | | PSNR / SSIM | PSNR / SSIM | PSNR / SSIM |
| 7 | 24.95 / 0.815 | 23.49 / 0.793 | 26.67 / 0.844 | 10 | 30.53 / 0.857 | 29.44 / 0.835 | 31.98 / 0.859 |
| 8 | 25.28 / 0.826 | 23.88 / 0.805 | 27.04 / 0.853 | 11 | 31.17 / 0.860 | 30.19 / 0.845 | 32.36 / 0.864 |
| 9 | 25.70 / 0.838 | 24.23 / 0.817 | 27.45 / 0.862 | 12 | 31.58 / 0.862 | 30.92 / 0.853 | 32.68 / 0.868 |
| 10 | 26.37 / 0.852 | 24.64 / 0.827 | 27.80 / 0.869 | 13 | 31.89 / 0.860 | 31.39 / 0.855 | 32.87 / 0.868 |
| 11 | 27.08 / 0.863 | 25.35 / 0.842 | 28.38 / 0.877 | 14 | 32.18 / 0.857 | 31.70 / 0.853 | 33.08 / 0.869 |
| 12 | 28.01 / 0.875 | 26.65 / 0.861 | 29.07 / 0.885 | 15 | 32.25 / 0.851 | 31.75 / 0.845 | 33.00 / 0.869 |




Figure 7: Qualitative comparison of the HyperCoil-Recon with other adaptive MRI reconstruction architectures - MAC-ReconNet and AdaIN for a multimodal scenario that combines 7, 10, and 12 coils of T1 brain and PD knee datasets. From left: AdaIN, MAC-ReconNet, HyperCoil-Recon, and the GT images for the brain and knee anatomies.

Table 4: Ablative study for an increasing number of DWP hypernetworks in the bottleneck and top layers of the decoder across acceleration factors for coronal PDFS knee

| Model | 5x | 8x |
|---------------------------------------|--------------|--------------|
| | PSNR/ SSIM | PSNR/ SSIM |
| No DWP | 32.39/ 0.823 | 30.56/ 0.798 |
| DWP (only in bottleneck layer) | 33.00/ 0.890 | 30.93/ 0.861 |
| DWP (both in top & bottleneck layers) | 33.65/ 0.893 | 31.34/ 0.862 |




Figure 8: An ablative study of cases with no DWP, DWP on the bottleneck layer and all the decoder layers. The residual images with respect to GT images are shown.

task-specific embeddings over the scale and bias operations of AdaIN and help to achieve mode-specific inductive bias [32, 26] when integrating diverse streams of contextual information in a single model.

4.2.5 Ablative Study

We perform an ablative study to understand the role of the hypernetworks on the top levels of the decoder and the bottleneck layer. We consider three cases across varying accelerations - 1) No dynamic weight prediction (DWP), 2) DWP only in the bottleneck layer, and 3) DWP in all layers *i.e.* bottleneck and top layers of the decoder. Table 4) shows the performance improvement with an increase in the number of hypernetworks as the contextual knowledge is improved. The qualitative results are shown in Figure 8.

5. Summary and Conclusion

We introduce a unified coil-configuration task-switching CNN in a multi-tasking perspective to infuse the knowledge

of dynamic coil configurations in multi-coil MRI reconstruction. Unlike conventional DL models, which scale with the number of coil configurations, the proposed network is simple and parameter-efficient by design. With each configuration posed as a task, our model uses hypernetworks to infer task-specific weights that are embedded in the primary reconstruction network, enabling on-the-fly adaptation to multiple coil configurations. Our experiments demonstrate the efficacy of our approach in terms of high expressivity by interpolating to several unseen configurations, better performance over other reconstruction methods, context-specific and context-invariant training methods, and insights into relationships between configurations. Our future direction is to extend the work for self-supervised learning for more PI settings in a single model.

References

- [1] H. M. Ahmed, R. E. Gabr, Y. M. Kadah, and A. M. Youssef. A new method for data acquisition and image reconstruction in parallel magnetic resonance imaging. In *2008 Cairo International Biomedical Engineering Conference*, pages 1–4, 2008.
- [2] Tim Bakker, Matthew Muckley, Adriana Romero-Soriano, Michal Drozdal, and Luis Pineda. On learning adaptive acquisition policies for undersampled multi-coil mri reconstruction. In Ender Konukoglu, Bjoern Menze, Archana Venkataraman, Christian Baumgartner, Qi Dou, and Shadi Albarqouni, editors, *Proceedings of The 5th International Conference on Medical Imaging with Deep Learning*, volume 172 of *Proceedings of Machine Learning Research*, pages 63–85. PMLR, 06–08 Jul 2022.
- [3] Youssef Beauferris, Jonas Teuwen, Dimitrios Karkalousos, Nikita Moriakov, Matthan Caan, George Yiasemis, Livia Rodrigues, Alexandre Lopes, Helio Pedrini, Letícia Rittner, Maik Dannecker, Viktor Studenyak, Fabian Gröger, Devendra Vyas, Shahrooz Faghih-Roohi, Amrit Kumar Jethi, Jaya Chandra Raju, Mohanasankar Sivaprakasam, Mike Lasby, Nikita Nogovitsyn, Wallace Loos, Richard Frayne, and Roberto Souza. Multi-coil mri reconstruction challenge—assessing brain mri reconstruction models and their generalizability to varying coil configurations. *Frontiers in Neuroscience*, 16, 2022.
- [4] Andrew Brock, Theodore Lim, James M. Ritchie, and Nick Weston. SMASH: one-shot model architecture search through hypernetworks. *CoRR*, abs/1708.05344, 2017.
- [5] Vinod Kumar Chauhan, Soheila Molaei, Marzia Hoque Tania, Anshul Thakur, Tingting Zhu, and David A. Clifton. Adversarial de-confounding in individualised treatment effects estimation, 2023.
- [6] Vinod Kumar Chauhan, Jiandong Zhou, Ping Lu, Soheila Molaei, and David A. Clifton. A brief review of hypernetworks in deep learning, arXiv preprint arXiv:2306.06955, 2023.
- [7] Dongdong Chen, Qingnan Fan, Jing Liao, Angelica Aviles-Rivero, Lu Yuan, Nenghai Yu, and Gang Hua. Controllable Image Processing via Adaptive FilterBank Pyramid. *IEEE Trans. Image Process.*, 29:8043–8054, 2020.
- [8] Steffen Czolbe and Adrian V. Dalca. Neuralizer: General neuroimage analysis without re-training, arXiv, 2305.02644, 2023.
- [9] Dean Darnell, Trong-Kha Truong, and Allen W. Song. Recent advances in radio-frequency coil technologies: Flexible, wireless, and integrated coil arrays. *Journal of Magnetic Resonance Imaging*, 55(4):1026–1042, 2022.
- [10] Jinming Duan, Jo Schlemper, Chen Qin, Cheng Ouyang, Wenjia Bai, Carlo Biffi, Ghalib Bello, Ben Statton, Declan P. O’Regan, and Daniel Rueckert. VS-Net: Variable Splitting Network for Accelerated Parallel MRI Reconstruction. In *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, pages 713–722, 2019.
- [11] Taejoon Eo, Yohan Jun, Taeseong Kim, Jinseong Jang, Ho-Joon Lee, and Dosik Hwang. Kiki-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images. *Magn. Reson. Med.*, 80(5):2188–2201, 2018.
- [12] Q. Fan, D. Chen, L. Yuan, G. Hua, N. Yu, and B. Chen. A General Decoupled Learning Framework for Parameterized Image Operators. *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 1–1, 2019.
- [13] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *Proc. ICML*, volume 70, pages 1126–1135, 06–11 Aug 2017.
- [14] Mark A. Griswold, Peter M. Jakob, Robin M. Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (grappa). *Magn. Reson. Med.*, 47(6):1202–1210, 2002.
- [15] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
- [16] Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P. Recht, Daniel K. Sodickson, Thomas Pock, and Florian Knoll. Learning a Variational Network for Reconstruction of Accelerated MRI Data. *Magn. Reson. Med.*, 79(6):3055–3071, 2018.
- [17] Rama Chellappa Hien Van Nguyen, Ronald Summers. *Metalearning with Medical Imaging and Health Informatics Applications*. The Elsevier and MICCAI Society book Series, 2022.
- [18] Jiahao Huang, Yingying Fang, Yinzhe Wu, Huanjun Wu, Zhifan Gao, Yang Li, Javier Del Ser, Jun Xia, and Guang Yang. Swin transformer for fast mri. *Neurocomputing*, 493:281–304, 2022.
- [19] Q. Huang, D. Yang, P. Wu, H. Qu, J. Yi, and D. Metaxas. MRI Reconstruction Via Cascaded Channel-Wise Attention Network. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 1622–1626, April 2019.
- [20] Siddhant M. Jayakumar, Jacob Menick, Wojciech M. Czarnecki, Jonathan Schwarz, Jack W. Rae, Simon Osindero, Yee Whye Teh, Tim Harley, and Razvan Pascanu. Multiplicative interactions and where to find them. In *International Conference on Learning Representations*, 2020.
- [21] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 105–114, 2017.
- [22] Etai Littwin, Tomer Galanti, Lior Wolf, and Greg Yang. On infinite-width hypernetworks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NeurIPS’20, Red Hook, NY, USA, 2020. Curran Associates Inc.
- [23] Xinwen Liu, Jing Wang, Feng Liu, and S. Kevin Zhou. Universal undersampled mri reconstruction. In Marleen de Bruijne, Philippe C. Cattin, Stéphane Cotin, Nicolas Padoy, Stefanie Speidel, Yefeng Zheng, and Caroline Essert, editors,

- Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, pages 211–221, Cham, 2021. Springer International Publishing.
- [24] Xinwen Liu, Jing Wang, Cheng Peng, Shekhar S. Chandra, Feng Liu, and S. Kevin Zhou. Undersampled mri reconstruction with side information-guided normalisation. In Linwei Wang, Qi Dou, P. Thomas Fletcher, Stefanie Speidel, and Shuo Li, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, pages 323–333, Cham, 2022. Springer Nature Switzerland.
- [25] Alexander Selvikvåg Lundervold and Arvid Lundervold. An overview of deep learning in medical imaging focusing on MRI. *Zeitschrift für Medizinische Physik*, 29(2):102 – 127, 2019.
- [26] Yao Ma, Shilin Zhao, Weixiao Wang, Yaoman Li, and Irwin King. Multimodality in meta-learning: A comprehensive survey. *Knowledge-Based Systems*, 250:108976, 2022.
- [27] Balamurali Murugesan, Sriprabha Ramanarayanan, Sricharan Vijayarangan, Keerthi Ram, Naranamangalam R Jagannathan, and Mohanasankar Sivaprakasam. A deep cascade of ensemble of dual domain networks with gradient-based t1 assistance and perceptual refinement for fast mri reconstruction. *Computerized Medical Imaging and Graphics*, 91:101942, 2021.
- [28] Klaas P. Pruessmann, Markus Weiger, Markus B. Scheidegger, and Peter Boesiger. Sense: Sensitivity encoding for fast mri. *Magn. Reson. Med.*, 42(5):952–962, 1999.
- [29] Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? towards understanding the effectiveness of MAML. *CoRR*, abs/1909.09157, 2019.
- [30] Sriprabha Ramanarayanan, Balamurali Murugesan, Keerthi Ram, and Mohanasankar Sivaprakasam. DC-WCNN: A Deep Cascade of Wavelet Based Convolutional Neural Networks for MR Image Reconstruction. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 1069–1073, 2020.
- [31] Sriprabha Ramanarayanan, Balamurali Murugesan, Keerthi Ram, and Mohanasankar Sivaprakasam. MAC-ReconNet: A Multiple Acquisition Context based Convolutional Neural Network for MR Image Reconstruction using Dynamic Weight Prediction. In *MIDL*, volume 121, pages 696–708, 06–08 Jul 2020.
- [32] Sriprabha Ramanarayanan, Arun Palla, Keerthi Ram, and Mohanasankar Sivaprakasam. Generalizing supervised deep learning mri reconstruction to multiple and unseen contrasts using meta-learning hypernetworks. *Applied Soft Computing*, page 110633, 2023.
- [33] James Requeima, Jonathan Gordon, John Bronskill, Sebastian Nowozin, and Richard E. Turner. *Fast and Flexible Multi-Task Classification Using Conditional Neural Adaptive Processes*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- [34] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, pages 234–241, 2015.
- [35] Jo Schlemper, Jose Caballero, Joseph V. Hajnal, Anthony Price, and Daniel Rueckert. A Deep Cascade of Convolutional Neural Networks for MR Image Reconstruction. In *Proc. Inf. Process. Med. Imaging*, pages 647–658, 2017.
- [36] Anuroop Sriram, Jure Zbontar, Tullie Murrell, C. Lawrence Zitnick, Aaron Defazio, and Daniel K. Sodickson. Grappanet: Combining parallel imaging with deep learning for multi-coil mri reconstruction. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14303–14310, 2020.
- [37] Guolei Sun, Thomas Probst, Danda Pani Paudel, Nikola Popovi, Menelaos Kanakis, Jagruti Patel, Dengxin Dai, and Luc Van Gool. Task switching network for multi-task learning. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8271–8280, 2021.
- [38] Liyan Sun, Zhiwen Fan, Xinghao Ding, Yue Huang, and John Paisley. Joint CS-MRI Reconstruction and Segmentation with a Unified Deep Network. In *Proc. Inf. Process. Med. Imaging*, pages 492–504, 2019.
- [39] Jeya Maria Jose Valanarasu, Pengfei Guo, Vibashan VS, and Vishal M. Patel. On-the-fly test-time adaptation for medical image segmentation, ARXIV.2203.05574, 2022.
- [40] Johannes von Oswald, Christian Henning, João Sacramento, and Benjamin F. Grewe. Continual learning with hypernetworks. *CoRR*, abs/1906.00695, 2019.
- [41] Risto Vuorio, Shao-Hua Sun, Hexiang Hu, and Joseph J. Lim. Multimodal model-agnostic meta-learning via task-aware modulation. *CoRR*, abs/1910.13616, 2019.
- [42] Alan Q. Wang, Adrian V. Dalca, and M. Sabuncu. Regularization-Agnostic Compressed Sensing MRI Reconstruction with Hypernetworks. *ArXiv*, abs/2101.02194, 2021.
- [43] Wei Wang, Ruiming Guo, Yapeng Tian, and Wenming Yang. CFSNet: Toward a Controllable Feature Space for Image Restoration. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4139–4148, 2019.
- [44] Heran Yang, Jian Sun, Liwei Yang, and Zongben Xu. A unified hyper-gan model for unpaired multi-contrast mr image translation. In Marleen de Bruijne, Philippe C. Cattin, Stéphane Cotin, Nicolas Padoy, Stefanie Speidel, Yefeng Zheng, and Caroline Essert, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, pages 127–137, Cham, 2021. Springer International Publishing.
- [45] George Yiasemis, Jan-Jakob Sonke, Clarisa Sánchez, and Jonas Teuwen. 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*, pages 732–741, 2022.