Reverse Engineering Breast MRIs: Predicting Acquisition Parameters Directly from Images

Author name(s) withheld EMAIL(S) WITHHELD
Address withheld

Editors: Under Review for MIDL 2023

Abstract

The image acquisition parameters (IAPs) used to create MRI scans are central to defining the appearance of the images. Deep learning models trained on data acquired using certain parameters might not generalize well to images acquired with different parameters. Being able to recover such parameters directly from an image could help determine whether a deep learning model is applicable, and could assist with data harmonization and/or domain adaptation. Here, we introduce a neural network model that can predict many complex IAPs used to generate an MR image with high accuracy solely using the image, with a single forward pass. These predicted parameters include field strength, echo and repetition times, acquisition matrix, scanner model, scan options, and others. Even challenging parameters such as contrast agent type can be predicted with good accuracy. We perform a variety of experiments and analyses of our model's ability to predict IAPs on many MRI scans of new patients, and demonstrate its usage in a realistic application. Predicting IAPs from the images is an important step toward better understanding the relationship between image appearance and IAPs. This in turn will advance the understanding of many concepts related to the generalizability of neural network models on medical images, including domain shift, domain adaptation, and data harmonization.

Keywords: MRI, acquisition parameters, breast, domain shift, inverse problem

1. Introduction

Magnetic resonance images (MRIs) are obtained under many different settings, such as the field strength, repetition and echo time, the type of contrast agent used, the acquisition matrix, as well as simply the scanner manufacturer and model (Chua et al., 2015; Saha et al., 2018). We collectively label these settings as the *image acquisition parameters* (IAPs). MRI scans are often used to train neural networks for medical image analysis tasks (Liu et al., 2018a; Angulakshmi and Deepa, 2021), such as breast cancer detection (Saha et al., 2018). While deep approaches often provide superior performance compared to early methods in medical image analysis (Le et al., 2019), they are also more prone to overfitting and *domain shift*. Domain shift is a problem in deep learning where the distribution of the data used to train a model differs from the distribution of the data used to test the model, and overfitting to the training domain(s) results in poor generalization to the test domain(s) (Guan and Liu, 2021; Wang and Deng, 2018).

In the context of medical imaging and MRIs in particular, domain shift occurs when images are taken according to different IAPs (along with other factors such as patient demographics and disease characteristics) (Glocker et al., 2019; Guan and Liu, 2021). Even if large training datasets are used that cover a variety of IAPs, it can still be common

© 2023 Anonymous.

for datasets to only represent a fraction of all possible MRI IAP settings; *e.g.*, due to an institution using only certain scanner types, which can cause model generalization issues.

Due to the susceptibility of neural networks to domain shift, it is essential to understand the precise domain/IAPs of data before using it to train or test a model. For example, knowing the IAPs of the data where the model will be applied allows developers to determine whether the target scans (1) fall in the range of training data and the model can be applied with more confidence or (2) fall outside of that range and the model should be applied with more caution (or not applied). However, IAPs are not always recorded and available in datasets. If the precise domain/IAPs of data *is* known, the development of image harmonization/standardization or domain adaptation techniques (Sec. 2) for transforming images and/or models to a particular desired domain becomes possible.

In this paper, we present a model (Fig. 1) that solves this *inverse problem* of recovering the IAPs that created an MR image using only the image itself.

Contributions. Our contributions are summarized as follows.

1. We introduce a neural network model for predicting many categorical and continuous IAPs of an MR image in one forward pass, trained via multi-task learning (Sec. 4.1).

2. We show that our model predicts many complex IAPs of MRI scans of new patients to high accuracy, over a test set of about $2,000$ slice images, with a series of experiments (Sec. 5.1). We predict six out of ten categorical IAPs to $> 97\%$ top-1 accuracy on the test set, and all but two with $> 95\%$ top-2 accuracy.

3. We show that our method achieves fair accuracy ($> 84\%$ top-1 accuracy, $> 95\%$ top-2) on IAPs that are more challenging to predict, such as contrast agent type.

4. We demonstrate a realistic application of our model: using it to sort new unlabeled data into different domains to determine which models to apply to the data for a downstream task (Sec. 5.2).

By showing that it is possible to learn a mapping between an image and it's domain-defining acquisition parameters, we take an important first step in precisely defining the relationship between medical images and their domain, laying the groundwork for future work on domain adaptation and image harmonization methods.

2. Related Works

Inverse problems and recovering imaging parameters from images. The task of recovering the parameters that generated an image directly from the image can be thought of as a type of inverse problem (Ren et al., 2020), where the goal is to recover the input to a function from the output of that function. Another inverse problem in MRI is that of recovering the physical phenomena that generated an image, such as the spin proton density (Benning and Ehrhardt, 2016), which is different from our task of recovering the scanner settings and characteristics that generated the image. This was attempted once for natural photographs (Laurence and Murphy, 2018), but was not successful.

The domain shift problem and domain adaptation approaches. The domain shift problem is well-established in both general computer vision (Wang and Deng, 2018) and

medical image analysis, the latter of which is due to differing IAPs and other factors (Glocker et al., 2019). Common *domain adaptation* techniques for mitigating this phenomena in medical imaging include adapting a trained model for the target domain (or preventing drastic overfitting to the source domain) (Zakazov et al., 2021), or transforming test images to a common domain prior to some downstream task (Koch et al., 2022; Diao et al., 2022; Zhang et al., 2018; Modanwal et al., 2021). Other methods seek to encode the intrinsic information content of medical images (Konz et al., 2022) in a way that is invariant to domain shift (Wolleb et al., 2022; Yang et al., 2019; Cao et al., 2022; Sun et al., 2022).

3. Dataset and MRI Acquisition Parameters

We use the Duke Breast Cancer (DBC) MRI dataset (Saha et al., 2018) for all experiments, which contains dynamic contrast-enhanced MRIs of 922 biopsy-confirmed breast cancer patients from over a decade. A large amount of clinical, demographic, pathological, genomic, and other data is provided for each scan, including 12 various MRI acquisition parameters (IAPs) which we study in this paper. These IAPs can each take different values (either categorical or continuous), and are summarized in Table 1. The categorical IAPs range from having 2 to 27 categories/classes each.

For all experiments, we use a subset of 14,000 2D slices randomly sampled from the original dataset's 3D fat saturated scan volumes, split randomly by patient into training, validation and testing subsets according to a ratio of 5:1:1, respectively, ensuring that slices from a given patient only appeared in one subset. We excluded 5 of the 922 patients from our data due to missing IAP values.

4. Methods

4.1. IAP Prediction Model

Our goal is to train a neural network to take new breast MRI slices as input and predict the values of the IAPs that generated these images. We thought it more practical and scalable to use a single network/forward pass to predict all IAPs at once, rather than train a seperate model for each IAP. To do so, we used a ResNet-18 (He et al., 2016) modified so that the final fully-connected layer is used to predict all IAPs at once, described as follows.

Each of the categorical studied IAPs (1-10 in Table 1) is converted into a one-hot encoding problem with C_k ($k = 1, \ldots, K$, $K = 10$) possible values/categories (Table 1), *e.g.*, $C_2 = 8$ for *Scanner Model*. Each kth categorical IAP is classified according to the maximum of the C_k corresponding units in the final layer. For the $M = 2$ continuous-valued IAPs TE and TR (in milliseconds), we use a single unit for each of these in our network output layer to predict the IAP value directly. As such, the final layer of our model has $\sum_{i=1}^{K} C_i + 2$ units, where $\sum_{i=1}^{K} C_i$ is the total of the number of possible values for each of the 10 categorical IAPs. We summarize our model in Fig. 1.

To train our model, we take a multi-task learning approach similar to that of (Liu et al., 2018b). We use a loss function that combines cross-entropy losses for the categorical IAPs

with mean squared error (MSE) losses for the continuous IAPs, of

$$\mathcal{L}_{\text{IAP}} = \lambda \sum_{k=1}^{K} \mathcal{L}_{\text{CE}}(\hat{y}_k, y_k) + \eta \sum_{m=1}^{M} \mathcal{L}_{\text{MSE}}(\hat{y}_m, y_m). \tag{1}$$

The first term in the loss is the sum of the classification losses for the categorical IAPs, where $\mathcal{L}_{\text{CE}}$ is the cross-entropy loss for the model's predicted class $\hat{y}_k$ and true class label y_k for the kth categorical IAP. Similarly, the second term is the sum of the regression losses for the continuous IAPs, where $\mathcal{L}_{\text{MSE}}$ is the MSE loss for the mth continuous IAP with predicted value $\hat{y}_m$ and true value y_m. λ and η are hyperparameters that control the relative importance of the categorical and continuous losses, respectively. We use $\lambda = 1$ and $\eta = 1$ for all experiments following experimentation on the validation set.

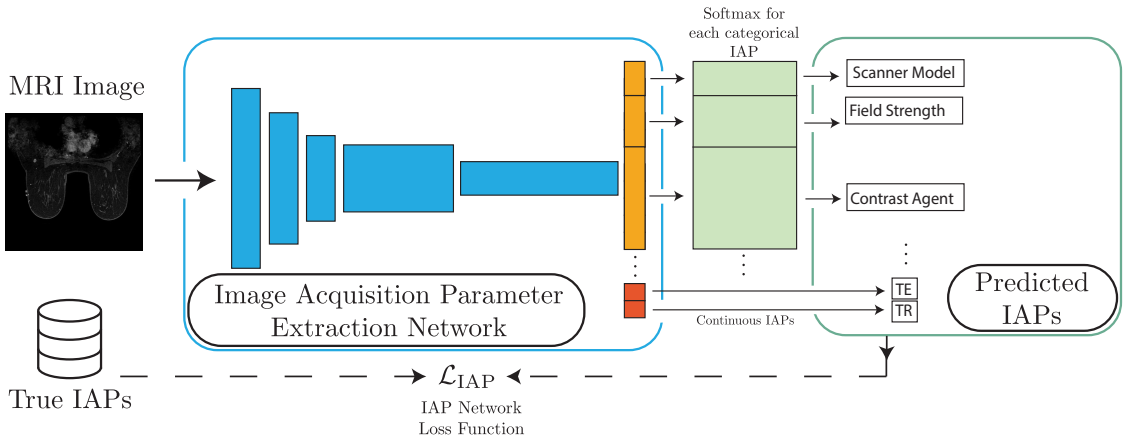

Figure 1: **Our MR image Acquisition Parameter (IAP) Extraction Model.** The model (Sec. 4.1) takes a breast MRI slice image as input, and predicts groups of class probabilities for each categorical IAP, and values for each continuous IAP (Table 1). The training pipeline is shown with dashed lines.

4.2. Experimental Settings

Our model was trained with a batch size of 512 for 100 epochs, using the Adam optimizer (Kingma and Ba, 2014) with a learning rate of 0.001 and weight decay strength of 0.0001. A model was saved for test set inference according to best performance on the validation set. All images were resized to 224×224 and normalized to $[0, 255]$, and all experiments were performed on a 48 GB NVIDIA A6000. The network took about 40 minutes to train.

5. Experiments and Results

We will first describe our experiments of using our model to predict the values of 12 IAPs (Table 1) of breast MR images of the test set of new patients (Sec. 5.1). Next, in Sec.

5.2 we will demonstrate an application of our IAP prediction model: using it to sort new unlabeled data into different domains to determine which trained models to apply to the data for a cancer classification task.

5.1. Predicting the acquisition parameters of MR images

In this section, we will describe our experiments of using our model to predict the values of 12 IAPs (Table 1) of breast MR images of the test set of new patients. We evaluated the performance of the network using mean squared error (MSE) for the continuous IAPs, and top-1 and top-2 accuracy for the categorical IAPs.

All IAP prediction results are shown in Table 1. We note again that all IAPs are predicted at the same time in a single forward pass of our model given an image, rather than separate passes/models for each IAP. We see that our model can predict most of the categorical IAPs to high accuracy, with $> 97\%$ top-1 accuracy on the test set, and all but two IAPs to greater than $> 95\%$ top-2 accuracy. Our model predicted the continuous IAPs (TE and TR) with relative errors of $< 1\%$ on average, as the TR values in the dataset range from 3.54–7.40ms, while TE ranges from 1.25–2.76ms. We also provide example predictions on specific diverse cases in Fig. 2. We see that our model is able to generalize to new patients taken with a variety of IAPs, able to predict the correct values for most of the IAPs.

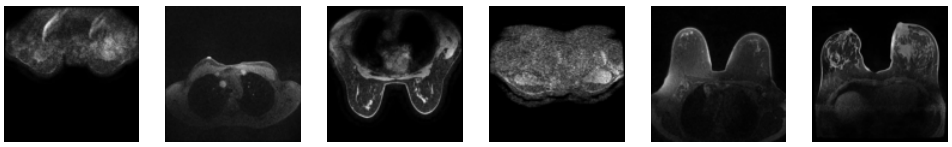

Predicted IAPs						
Manufacturer:	GE ✓	SIEMENS ✓	GE ✓	GE ✓	SIEMENS ✓	SIEMENS ✓
Model:	Signa HDxt ✓	Avanto ✓	Signa HDxt ✓	SIGNA HDx ✓	Avanto ✓	Skyra ✓
Scan Options:	1 ✓	5 ✓	1 ✓	2 ✓	5 ✓	6 ✓
Patient Position:	FFP ✓	FFP ✓	FFP ✓	FFP ✓	FFP ✓	FFP ✓
Field Strength:	3 T ✓	1.5 T ✓	3 T ✓	1.5 T ✓	1.5 T ✓	3 T ✓
Contrast Agent:	MULTIHANCE ✓	MULTIHANCE ✓	MAGNEVIST ✓	MAGNEVIST ✓	MAGNEVIST ✓	MULTIHANCE ✓
Acquisition Matrix:	384 X 384 ✓	448 X 448 ✓	350 X 350 ✓	340 X 340 ✓	448 X 448 ✓	448 X 381 ✓
Slice Thickness:	1.2 mm ✓	2 mm ✓	2 mm ✓	2 mm ✓	1.1 mm ✓	2 mm ✓
Flip Angle:	10 deg ✓	10 deg ✓	10 deg ✓	10 deg ✓	12 deg ✓	10 deg ✓
FOV Computed:	340 cm ✓	340 cm ✓	340 cm ✓	320 cm ✓	360 cm ✗	350 cm ✓
TE (Echo Time):	2.537 ms ✓	1.304 ms ✓	2.382 ms ✓	2.373 ms ✓	1.324 ms ✓	1.445 ms ✓
TR (Rep. Time):	5.963 ms ✓	4.415 ms ✗	5.686 ms ✓	5.171 ms ✗	4.197 ms ✓	3.722 ms ✓

Figure 2: **Example Predictions of Acquisition Parameters for MRIs in the Test Set.** Each image is from a different patient, and below each image are the predicted values for each of its IAPs (listed to the left). The symbols "✓" and "×" indicate correct and incorrect predictions, respectively (TE and TR predictions are treated as "correct" if the relative error is $< 2\%$).

5.1.1. The most challenging acquisition parameters

The two most challenging IAPs were *Slice Thickness* (21 categories) and *FOV Computed* (27 cats.), which had noticeably lower accuracies than the other IAPs (76.66% and 51.21%,

Table 1: **Model Prediction Performance for all IAPs on the Test Set.** Shown are all MR image acquisition parameters (IAPs) that we analyze, the number of possible categories of each in the dataset with examples, and the performance of our model for predicting them on the test set. Top-1 and top-2 prediction accuracy is provided for the categorical IAPs (first 10 rows), and mean squared error is provided for the continuous IAPs (bottom two rows).

	MRI acquisition parameter (IAP)	**No. categories**	**Examples**	**Top-1 pred. acc. (%)**	**Top-2 pred. acc. (%)**	**Pred. MSE**
1	Scanner Manufacturer	2	GE, Siemens	99.74	N/A	N/A
2	Scanner Model	8	Avanto, Signa HDx	97.78	99.29	N/A
3	Scan Options	9	PFP/FS, PFP/SFS	99.40	99.60	N/A
4	Field Strength	5	1.5 T, 3 T	98.19	99.70	N/A
5	Patient Position	2	FFP, HFP	97.73	N/A	N/A
6	Contrast Agent Type	6	Gadavist, MultiHance	84.73	95.46	N/A
7	Acquisition Matrix	10	448×448, 384×360	91.53	99.14	N/A
8	Slice Thickness	21	1.3 mm, 2 mm	76.66	87.05	N/A
9	Flip Angle	4	10°, 12°	99.65	99.75	N/A
10	FOV Computed	27	320 cm, 360 cm	51.21	69.30	N/A
11	Repetition Time (TR)	N/A	4.27 ms, 5.34 ms	N/A	N/A	0.0305 ms
12	Echo Time (TE)	N/A	2.4 ms, 1.5 ms	N/A	N/A	0.0116 ms

respectively). Both are physically continuous parameters that were discretized into many categories when the dataset was created, which is why the top-2 accuracies are much higher (87.05% and 69.30%, respectively), as our model was close to the correct values. *Contrast Agent Type* was also more difficult to predict than most other categorical IAPs (84.73% accuracy); this is because predicting this IAP from an image forms a difficult inverse problem, *i.e.*, the mapping between it and the final image is not necessarily one-to-one.

5.2. Application: Sorting Unlabeled Data into Domains for Downstream Task Model Selection

In this section we will demonstrate an application of our model: sorting new MR images without IAP labels into different domains, to determine which cancer classification models that were trained in different domains to apply to the data. This forms a challenge because applying a model trained on images from one scanner manufacturer to an image from another manufacturer can result in lowered performance (a domain shift problem).

We trained a ResNet-18 for a binary cancer classification task on only images taken with a GE scanner from the training set, which we label "**GE Model**". We take slices containing a tumor bounding box in the dataset to be labeled positive, and all other slices at least five slices away from the positives to be negative. The similar Siemens-only trained model is "**Siemens Model**". All other training settings were the same as in Sec. 4.2. We used the same test set as prior experiments, which has images from both GE and Siemens patients (Sec. 3).

Our results are summarized in Table 2. Without our IAP prediction model, the scanner manufacturers (GE or Siemens) of the test set images remain unknown so the only way to proceed is to apply the GE Model to the test set, or the Siemens Model to the test set, which is not ideal, as there is no way to know if the correct model is being used for a given test image. This resulted in cancer classification accuracies of 68.86% and 56.50% for the GE and Siemens models, respectively; a domain shift problem is indeed present.

However, with our trained IAP prediction model (Sec. 5.1), we can determine the correct model to apply to each image according to the predicted manufacturer of that image from the IAP model. This resulted in a cancer classification accuracy of **76.95**%, a significant improvement over the uninformed approach.

Table 2: **Using our IAP prediction model to sort unlabeled data for cancer classification model selection.** See Sec. 5.2; accuracies are on the test set.

	GE Model	**Siemens Model**	**Model chosen according to predicted IAPs**
Cancer classification accuracy	68.82%	56.50%	**76.95**%

6. Discussion

In this paper, we introduced a model that can predict the image acquisition parameters (IAPs) that generated an MR image solely from the image. Trained via multi-task learning on a breast MRI dataset with 12 analyzed IAPs, our model was able to predict most IAPs of new patient scans in the test set to high accuracy. We also demonstrated the usefulness of our model for sorting unlabeled data into domains for cancer classification model selection.

We found very good performance ($> 97\%$ top-1 accuracy) on predicting the categorical IAPs of *Scanner Manufacturer* (2 categories), *Scanner Model* (8 cats.), *Scan Options* (9 cats.), *Field Strength* (5 cats.), *Patient Position* (2 cats.) and *Flip Angle* (4 cats.). *Acquisition Matrix* (10 cats.) was a bit more challenging (91.53% top-1 acc.), but still good. The continuous IAPs of TE and TR were also able to be predicted to high accuracy. Some of these IAPs seem fairly trivial to predict (*Patient Position*), but it is surprising that others such as *Scan Options*, *Field Strength* and TE/TR can be recovered so well, as these are not necessarily clear from the image, even to a radiologist.

For the IAPs mentioned so far, the mapping of IAP value to the final image is mostly predictable, and the domain change of the image due to the IAP is clear. This also confirms that indeed, the predictions of neural networks are clearly affected by the domain of the image, as these IAPs affect the domain and they can be predicted solely from the image.

Contrast Agent Type (84.73% top-1, 6 cats.), *Slice Thickness* (76.77% top-1, 21 cats.) and *FOV Computed* (51.21% top-1, 27 cats.) were the more challenging IAPs to predict, likely for the reasons discussed in Sec. 5.1.1. However, our model was still able to perform significantly better than random guessing, especially when considering the much higher top-2 accuracies (95.46%, 87.05% and 69.30%, respectively). This means that the mapping of these IAPs to the final image is uncertain and not as predictable as the other IAPs, but still somewhat learnable.

There are certain limitations to our study. One is that the dataset has an imbalanced distribution of images with respect to certain IAP values (*e.g.*, some scanner model is less common). This could result in evaluation metrics for our model biased towards common IAPs. We mitigated this by using large training and testing datasets to cover a wide range of IAP settings, and our model seems strong for predicting a wide variety of IAP values (Fig. 2), but this effect should be considered.

Possible future work could include improving our model using other multi-task learning methods, or using 2.5D or 3D convolutional nets on the full 3D scan volumes, which was beyond the scope of this "proof of concept" work. We would also like to test our model's ability to generalize to other MR datasets and/or modalities that include IAPs, such as brain MRI. Finally, applications of our model's ability to predict domain characteristics (IAPs) from images for domain adaptation or data harmonization methods should be explored.

7. Conclusion

We presented a neural network model for inferring the image acquisition parameters (IAPs) that generated an MR image solely from the image, and obtained highly accurate predictions for most IAPs on a breast MRI dataset. This lays the groundwork for a better understanding of the relationship between medical images and their domain, which could be useful for data harmonization methods, domain adaptation approaches, and other applications.

References

M Angulakshmi and M Deepa. A review on deep learning architecture and methods for mri brain tumour segmentation. *Current Medical Imaging*, 17(6):695–706, 2021.

Martin Benning and Matthias J Ehrhardt. Inverse problems in imaging. *Lecture Notes. Nov*, 2016.

Shixing Cao, Nicholas Konz, James Duncan, and Maciej A. Mazurowski. Deep learning for breast mri style transfer with limited training data. *Journal of Digital Imaging*, pages 1–13, 2022.

Alicia S Chua, Svetlana Egorova, Mark C Anderson, Mariann Polgar-Turcsanyi, Tanuja Chitnis, Howard L Weiner, Charles RG Guttmann, Rohit Bakshi, and Brian C Healy. Handling changes in mri acquisition parameters in modeling whole brain lesion volume and atrophy data in multiple sclerosis subjects: comparison of linear mixed-effect models. *NeuroImage: Clinical*, 8:606–610, 2015.

Pengfei Diao, Akshay Pai, Christian Igel, and Christian Hedeager Krag. Histogram-Based Unsupervised Domain Adaptation for Medical Image Classification. In Linwei Wang, Qi Dou, P. Thomas Fletcher, Stefanie Speidel, and Shuo Li, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, Lecture Notes in Computer Science, pages 755–764, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-16449-1. doi: 10.1007/978-3-031-16449-1_72.

Ben Glocker, Robert Robinson, Daniel C Castro, Qi Dou, and Ender Konukoglu. Machine learning with multi-site imaging data: an empirical study on the impact of scanner effects. *arXiv preprint arXiv:1910.04597*, 2019.

Hao Guan and Mingxia Liu. Domain adaptation for medical image analysis: a survey. *IEEE Transactions on Biomedical Engineering*, 69(3):1173–1185, 2021.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Valentin Koch, Olle Holmberg, Hannah Spitzer, Johannes Schiefelbein, Ben Asani, Michael Hafner, and Fabian J. Theis. Noise Transfer for Unsupervised Domain Adaptation of Retinal OCT Images. In Linwei Wang, Qi Dou, P. Thomas Fletcher, Stefanie Speidel, and Shuo Li, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, volume 13432, pages 699–708. Springer Nature Switzerland, Cham, 2022. ISBN 978-3-031-16433-0 978-3-031-16434-7. doi: 10.1007/978-3-031-16434-7_67. URL https://link.springer.com/10.1007/978-3-031-16434-7_67. Series Title: Lecture Notes in Computer Science.

Nicholas Konz, Hanxue Gu, Haoyu Dong, and Maciej A. Mazurowski. The intrinsic manifolds of radiological images and their role in deep learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 684–694. Springer, 2022.

Edward Laurence and Charles Murphy. Camera settings prediction using deep convolutional neural networks. *D'epartement de Physique, de G'enie Physique, et d'Optique, Universit'e Laval, Quebec (Quebec), Canada G1V 0A6*, 2018.

EPV Le, Y Wang, Yuan Huang, Sarah Hickman, and FJ Gilbert. Artificial intelligence in breast imaging. *Clinical radiology*, 74(5):357–366, 2019.

Jin Liu, Yi Pan, Min Li, Ziyue Chen, Lu Tang, Chengqian Lu, and Jianxin Wang. Applications of deep learning to mri images: A survey. *Big Data Mining and Analytics*, 1(1): 1–18, 2018a.

Mingxia Liu, Jun Zhang, Ehsan Adeli, and Dinggang Shen. Joint classification and regression via deep multi-task multi-channel learning for alzheimer's disease diagnosis. *IEEE Transactions on Biomedical Engineering*, 66(5):1195–1206, 2018b.

Gourav Modanwal, Adithya Vellal, and Maciej A. Mazurowski. Normalization of breast MRIs using cycle-consistent generative adversarial networks. *Computer Methods and Programs in Biomedicine*, 208:106225, September 2021. ISSN 0169-2607. doi: 10.1016/j.cmpb.2021.106225. URL https://www.sciencedirect.com/science/article/pii/S0169260721002996.

Simiao Ren, Willie Padilla, and Jordan Malof. Benchmarking deep inverse models over time, and the neural-adjoint method. *Advances in Neural Information Processing Systems*, 33: 38–48, 2020.

Ashirbani Saha, Michael R Harowicz, Lars J Grimm, Connie E Kim, Sujata V Ghate, Ruth Walsh, and Maciej A Mazurowski. A machine learning approach to radiogenomics of breast cancer: a study of 922 subjects and 529 dce-mri features. *British journal of cancer*, 119(4):508–516, 2018.

Xiaoyi Sun, Zhizhe Liu, Shuai Zheng, Chen Lin, Zhenfeng Zhu, and Yao Zhao. Attention-Enhanced Disentangled Representation Learning for Unsupervised Domain Adaptation in Cardiac Segmentation. In Linwei Wang, Qi Dou, P. Thomas Fletcher, Stefanie Speidel, and Shuo Li, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, Lecture Notes in Computer Science, pages 745–754, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-16449-1. doi: 10.1007/978-3-031-16449-1_71.

Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018.

Julia Wolleb, Robin Sandkühler, Florentin Bieder, Muhamed Barakovic, Nouchine Hadjikhani, Athina Papadopoulou, Özgür Yaldizli, Jens Kuhle, Cristina Granziera, and Philippe C. Cattin. Learn to Ignore: Domain Adaptation for Multi-site MRI Analysis. In Linwei Wang, Qi Dou, P. Thomas Fletcher, Stefanie Speidel, and Shuo Li, editors,

Medical Image Computing and Computer Assisted Intervention – MICCAI 2022, Lecture Notes in Computer Science, pages 725–735, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-16449-1. doi: 10.1007/978-3-031-16449-1_69.

Junlin Yang, Nicha C Dvornek, Fan Zhang, Julius Chapiro, MingDe Lin, and James S Duncan. Unsupervised domain adaptation via disentangled representations: Application to cross-modality liver segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 255–263. Springer, 2019.

Ivan Zakazov, Boris Shirokikh, Alexey Chernyavskiy, and Mikhail Belyaev. Anatomy of domain shift impact on u-net layers in mri segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 211–220. Springer, 2021.

Jun Zhang, Ashirbani Saha, Brian J. Soher, and Maciej A. Mazurowski. Automatic deep learning-based normalization of breast dynamic contrast-enhanced magnetic resonance images, July 2018. URL http://arxiv.org/abs/1807.02152. arXiv:1807.02152 [cs].