

Towards Learning Contrast Kinetics with Multi-Condition Latent Diffusion Models

Richard Osuala^{1,2,3}, Daniel M. Lang^{2,3}, Preeti Verma¹, Smriti Joshi¹, Apostolia Tsirikoglou⁴, Grzegorz Skorupko¹, Kaisar Kushibar¹, Lidia Garrucho^{1,4}, Walter H. L. Pinaya⁵, Oliver Diaz^{1,6}, Julia A. Schnabel^{2,3,5}, and Karim Lekadir^{1,7}

¹ Departament de Matemàtiques i Informàtica, Universitat de Barcelona, Spain
richard.osuala@ub.edu

² Helmholtz Center Munich, Munich, Germany

³ Technical University of Munich, Munich, Germany

⁴ Karolinska Institutet, Stockholm, Sweden

⁵ Kings College London, London, UK

⁶ Computer Vision Center, Bellaterra, Spain

⁷ Institució Catalana de Recerca i Estudis Avançats (ICREA), Barcelona, Spain

Abstract. Contrast agents in dynamic contrast enhanced magnetic resonance imaging allow to localize tumors and observe their contrast kinetics, which is essential for cancer characterization and respective treatment decision-making. However, contrast agent administration is not only associated with adverse health risks, but also restricted for patients during pregnancy, and for those with kidney malfunction, or other adverse reactions. With contrast uptake as key biomarker for lesion malignancy, cancer recurrence risk, and treatment response, it becomes pivotal to reduce the dependency on intravenous contrast agent administration. To this end, we propose a multi-conditional latent diffusion model capable of acquisition time-conditioned image synthesis of DCE-MRI temporal sequences. To evaluate medical image synthesis, we additionally propose and validate the Fréchet radiomics distance as an image quality measure based on biomarker variability between synthetic and real imaging data. Our results demonstrate our method’s ability to generate realistic multi-sequence fat-saturated breast DCE-MRI and uncover the emerging potential of deep learning based contrast kinetics simulation. We publicly share our accessible codebase at <https://github.com/RichardObi/ccnet> and provide a user-friendly library for Fréchet radiomics distance calculation at <https://pypi.org/project/frd-score>.

Keywords: Contrast Agent · Synthesis · DCE-MRI · Generative Models

1 Introduction

Magnetic resonance imaging (MRI) and, in particular, dynamic contrast enhanced (DCE)-MRI are remarkably sensitive and effective modalities for tumor detection, localization and characterization and, thus, have become ubiquitous in clinical cancer treatment planning and monitoring. The uptake of contrast

in DCE-MRI sequences plays a pivotal role as biomarker for cancer detection, tumor molecular subtype and malignancy differentiation, as well as cancer recurrence and treatment response prediction [21,1]. However, intravenously-injected gadolinium-based contrast agents (CA) used in DCE-MRI have been associated with a wide range of concerns [11], such as a risk of nephrogenic systemic fibrosis, bioaccumulation in the brain, and its invasiveness causing a non-applicability in patient populations with pregnancy, adverse reactions, kidney malfunction or where consent is missing. Furthermore, due to its multiple temporal acquisitions, DCE-MRI is costly and time-consuming, prone to motion artifacts, and the contrast injection an uncomfortable procedure for patients [25,12,11].

Public health institutions, such as the European Medicines Agency, recommend to restrict gadolinium-based CA [3], which further emphasizes the need to develop alternative methods. A potentially faster, cost-effective, motion artifact-free, and non-invasive alternative is the synthetic generation of DCE-MRI using deep generative models. First studies applied generative adversarial networks (GANs) [4] to generate post-contrast images from pre-contrast images [10,25,12]. Recent works have further used diffusion models [17] and latent diffusion models (LDMs) [15] to synthesize medical images [14,2,7], such as pre-contrast breast MRI conditioned on anatomical segmentation masks [8]. Despite their recent successes in computer vision, diffusion models and LDMs have, to the best of our knowledge, not been applied to pre- to post-contrast DCE-MRI translation.

A largely unaddressed aspect in medical image synthesis is domain-specific image quality evaluation. To date, popular methods, such as the Fréchet inception distance (FID) [5], are based on feature extractor models trained on natural image datasets. Despite a considerable domain gap, these methods are commonly applied to medical imaging data without alteration, thereby failing to capture medical nuances such as abnormalities [2]. Recently, an FID calculation based on a radiology domain-specific feature extractor was proposed [13], which nevertheless showed some volatility and no significant correlation with human judgement [20]. A further limitation of such methods is the pretraining of the feature-extractor on 2D data (with unknown inherent biases), which is not applicable nor readily extendable to 3D medical images.

In this work, we aim to address the aforementioned gaps, resulting in the following three contributions:

- Design, implementation and validation of a multi-conditional latent diffusion model for pre- to post-contrast MRI synthesis.
- We present the first work that simulates time-dependent contrast uptake on imaging data using diffusion models.
- We propose and validate the Fréchet radiomics distance (FRD), a novel radiology-specific quality evaluation method of 3D and 2D synthetic images based on biomarker variability.




Fig. 1: Overview of our proposed methods, including ContrastControlNet (CCNet) and the Fréchet radiomics distance (FRD). CCNet trains the denoising U-Net and the ControlNet in consecutive stages under contrast enhancement-specific conditioning (pre-contrast image, text, acquisition time). During inference, E is discarded (in violet) and, based on a random latent z_T and w -weighted ControlNet guidance, the U-Net generates the post-contrast image latent z_{T0} . z_{T0} is divided by factor S and decoded via D into image space. Finally, FRD compares extracted real and synthetic imaging biomarker distributions.

2 Methodology

2.1 Multi-Condition Latent Diffusion Model

Recently, diffusion models [17, 6] have emerged as a promising new class of generative models due to their exceptional ability to model complex distributions. Such diffusion models consist of two processes, namely, a forward diffusion process and a reverse denoising process. The forward process is a Gaussian transition that gradually destroys the structure of a real data point $x_0 \sim p(x_0)$ by adding noise with different scales to obtain a series of noisy variables $x_1, x_2, \dots, x_T$:

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t | \sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}) \mathbf{I}). \quad (1)$$

The reverse process is parameterized by another Gaussian transition which gradually denoises x_T in T timesteps resulting in restoration of initial data point x_0 :

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1} | \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)). \quad (2)$$

Latent diffusion models (LDM) [15] are designed to iteratively denoise a learned latent representation z_T of image x_T rather than operating directly on image space, which allows for more memory-efficient training and improved conditioning. The latter includes high-quality image synthesis based on textual input

[15,2]. However, text descriptions have been shown to be either inefficient or insufficient to accurately convey detailed controls upon the final generated images (e.g., to control the fine-grained semantic layout).

ControlNet [23] introduces plug-and-play condition encoders tailored for pre-trained diffusion models. We adopt ControlNet as auxiliary encoder to integrate the conditioning pre-contrast image while preserving the integrity of the original post-contrast LDM generator. As shown in Fig. 1, the ControlNet encoder outputs are reintegrated into the diffusion model through zero-convolutional layers. We enhance this approach by extracting and propagating relevant textual meta-data. We parse a clinical tabular dataset to assemble free-text prompts based on factors influencing DCE-MRI contrast manifestation. The final text is input to the denoising U-Net of both the LDM and the ControlNet via cross-attention and contains manufacturer, scanner, field strength, contrast agent, and bolus volume information. Furthermore, we integrate the time passed since pre-contrast acquisition into the model. Here, we extract pre- and post-contrast acquisition times from respective DICOM headers and input them as continuous variables into two dense layers before concatenating the resulting output with the timestep embedding of ControlNet and the denoising U-Net. Through optimization of the mean squared error between predicted and added noise per timestep, our ContrastControlNet (CC-Net) method learns to (i) extract meaningful patterns from pre-contrast, (ii) interpret conditions, (iii) reconstruct the corresponding post-contrast image, (iv) localize the lesion, and (v) predict realistic hyper- and hypo-intense lesion contrast uptake patterns.

2.2 Biomarker Variability as Image Quality Measure

Despite the significant domain gap stemming from its feature extractor being trained on natural RGB images, FID [5] is a commonly used metric in medical imaging, measuring the diversity and fidelity of synthetic images via real-synthetic distribution comparison. However, high-fidelity (and high-diversity) synthetic data can still be of low clinical utility and vice versa [22]. For analyzing synthetic data, we note the need to ideally encompass (a) image utility indicators, (b) medical imaging domain-specific measurement, (c) scalability from 2D to 3D settings, and (d) explainable results based on interpretable features. Balancing these requirements, we propose measuring synthetic data quality based on imaging biomarker variability. In particular, we measure the feature distribution distance of multiple normalized biomarker values extracted from real and synthetic images. In radiology settings, 2D and 3D radiomics features can be extracted as non-invasive imaging biomarkers that quantify phenotypic characteristics [9,19]. Further noting the capabilities of radiomics to capture tumor heterogeneity [9], or as predictor of treatment response [1] and tumor subtype [16] in DCE-MRI, we introduce the Fréchet radiomics distance (FRD) as synthetic data quality measure. While the selection of radiomics features included in FRD is flexible, we compute FRD based on the common set of features suggested by [19] including first-order statistics ($n=19$), co-occurrence gray level matrix (GLM) ($n=24$), run length GLM ($n=16$), size zone GLM ($n=16$), neighbouring

gray tone difference matrix ($n=5$) and dependence GLM ($n=14$). Features are computed given the image and, optionally, a respective region of interest annotation. As depicted in Fig. 1, to compute the FRD between two imaging datasets, we extract and normalize FRD features per image and dataset and model the resulting two feature sets as Gaussian distributions, allowing us to compute a distance between them. Hence, for each image x_i in a dataset, we extract a value v_{ji} for each FRD feature j . Then, each v_{ji} is min-max normalized based on the values $v_{j1}, v_{j2}, \dots, v_{jn}$ over all images X in the dataset. Next, the resulting values v_j of feature j are scaled to the common range observed for FID latent feature values, i.e. $[0, 7.456]$. This calibration later allows for interpretation of final FRD value and its comparison to FID, considering the intuition the image synthesis field has developed for FID value interpretation [2,13,20,14,22]. The obtained synthetic and real feature sets V are fitted to multivariate Gaussian distributions. These distributions are defined by their means (μ) and covariance matrices (Σ). The FRD value is computed as the dissimilarity between real data X and synthetic data Y via the Fréchet distance defined as:

$$FD(X, Y) = \|\mu_X - \mu_Y\|_2^2 + \text{tr}(\Sigma_X + \Sigma_Y - 2(\Sigma_X \Sigma_Y)^{\frac{1}{2}}). \quad (3)$$

3 Experiments and Results

3.1 Dataset and Implementation

In this study, we use the public Duke-Breast-Cancer-MRI Dataset [16]. The dataset and its imaging metadata encompasses 922 biopsy-confirmed breast cancer cases, each comprising one fat-saturated T1 sequence (pre-contrast) and up to 4 corresponding fat-saturated T1-weighted DCE sequences (post-contrast) with a median of 131 seconds passed between DCE sequences. The 1.5T or 3T MRI scans come in dimensions of either 320^2 , 448^2 or 512^2 in the coronal and sagittal planes, with varying slice numbers in the axial plane. As we extract tumor-containing axial slices, we note considerable changes between pre- and post-contrast in non-tumor related areas (e.g., heart area). Hence, we crop the extracted slices first increasing the width and height of the tumor bounding box to half the width and height of the full image (e.g., 224^2 in case of a 448^2 image) thereby resulting in a tumor-containing single breast image. We split the dataset by patient into train ($n=842$), validation ($n=50$), and test ($n=30$) sets. All models were trained on a single GPU, either Nvidia A100 (80GB RAM) or RTX A6000 (48GB RAM), using Python3.11's pytorch, diffusers and monai-generative libraries [14].

3.2 Fréchet Radiomics Distance as Image Perturbation Correlate

Adopting the validation strategy of the FID in its original publication [5], we observe the correlation between the Fréchet radiomics distance (FRD) and the amount by which the quality of an imaging dataset is reduced. Concretely, we compare FRD feature distributions between an unchanged imaging dataset and




Fig. 2: (a) Image perturbation scales in breast MRI, (b) resulting Fréchet radiomics distance (FRD) values (y-axis) per percentage scale (x-axis) per applied perturbation for 2D axial slices and for 3D volumes based on DCE-MRI post-contrast phase 1 data from 254 patient cases.

its quality-reduced equivalent, where the quality reduction is based on a scaling factor. As visualized in Fig. 2a, we apply this procedure to 254 DCE-MRI (phase 1) cases of the Duke Dataset using Gaussian noise and Gaussian blurring for image perturbation on scales from 1% up to 50%. Unlike the FRD results shown in Table 1, in Fig. 2b we calculate the scores on full axial slices using the tumor mask. We observe FRD monotonically increasing with perturbation scale demonstrating FRD’s capability of capturing the quality-reduction level for both 2D and 3D data.

3.3 Generation of DCE Sequences from Pre-Contrast Images

In CC-Net, we first initialize a pretrained autoencoder (AE) from stable diffusion (SD), and then (b) proceed to train the denoising U-Net (i.e., LDM) before (c) training the ControlNet, after which we finally (d) run inference on the test set. In each step, we observe several hyperparameters to influence the empirical results. In (a), we note visible differences between the generally accurate breast MRI reconstructions of different pretrained SD AEs. We select the AE from SD *2-1-base* over *v1-5* and *xl-base-1.0* [15]. In (b), we notice a high dependence of output quality on scaling factor s with which the AE representation is multiplied before being used in U-Net training. We find $s=0.1$ to improve results upon the recommended $s=0.18215$. We further identify a tendency for exploding gradients which was better addressed by clipping the gradient value rather than its norm (e.g., at 15). Both fine-tuning the SD U-Net, but also training it from scratch produced desirable outputs. We use a DDPM [6] noise scheduler with 1000 timesteps, AdamW as optimizer, and batch sizes of 8, 16, and 32, and learning rates of 2.5^{-5} and 5^{-5} , for which we obtain similar results. The

Table 1: Synthetic image quality evaluation based on FRD, FID, LPIPS, and MSE metrics. Where applicable, results are reported with standard deviation and are based on 30 test cases consisting of 1010 images in each of the post-contrast phases P1, P2, and P3, and 417 images in phase P4. *Any* refers to time-conditioned model training on all post-contrast phases. In *+TXT* textual input is used in training and inference, *+LDM* refers to LDM model training from scratch as opposed to stable diffusion *2-1-base* fine-tuning. *+CG* refers to a ControlNet guidance weight increase, e.g. from 1 to 1.6, during inference. *+LT* stands for a 50 epoch longer training of ControlNet. Best results are in bold.

224x224 Single Breasts with Tumor		Metrics			
Set 1	Set 2	FRD ↓	FID ↓	LPIPS ↓	MSE ↓
Real Pre-Contrast	Real DCE-P1	49.07	68.20	0.223±.102	51.18±17.80
CC-Net _{Any} P1	Real DCE-P1	35.64	41.38	0.192±.070	46.92±15.40
CC-Net _{Any+Txt} P1	Real DCE-P1	39.98	42.85	0.186±.073	47.20±15.76
CC-Net _{Any+Txt+LDM} P1	Real DCE-P1	41.55	62.41	0.200±.074	49.17±14.48
CC-Net _{Any+Txt+LDM+CG1.6} P1	Real DCE-P1	22.50	64.61	0.194±.072	46.12±14.21
CC-Net _{Any+Txt+LDM+LT} P1	Real DCE-P1	45.19	60.58	0.193±.073	48.08±14.44
CC-Net _{Any+Txt+LDM+CG1.6+LT} P1	Real DCE-P1	37.70	62.21	0.192±.072	45.48±13.60
Real Pre-Contrast	Real DCE-P2	74.26	64.90	0.212±.095	50.00±16.85
CC-Net _{Any} P2	Real DCE-P2	58.07	40.36	0.191±.076	46.10±14.19
Real Pre-Contrast	Real DCE-P3	84.13	60.96	0.208±.092	49.23±16.15
CC-Net _{Any} P3	Real DCE-P3	61.17	37.80	0.190±.074	45.75±13.74
Real Pre-Contrast	Real DCE-P4	100.27	77.31	0.199±.078	52.48±12.96
CC-Net _{Any} P4	Real DCE-P4	47.13	60.80	0.198±.075	50.36±14.26

U-Net is trained for 100 epochs and selected from the epoch with lowest validation loss for further use in (c) and (d). In (c), we follow the hyperparameter setup from (b) with half the batch size and no gradient clipping. In (d), we increase inference speed by using a denoising diffusion implicit model (DDIM) [18] scheduler with 200 timesteps without visible performance decrease. We set the text guidance scale to 1 as higher scales did not improve output quality, however, increasing the ControlNet guidance weight from 1 to a value in range (1, 1.6] enhanced perceived image quality. In (a)-(c), following AE pretraining, the 224x224 input images are stacked in 3-channels and normalized in range [-1, 1] before applying small-scale intensity and affine augmentations. As observable in Fig. 3a, the contrast-enhanced tumors account for a large part of the difference between pre- and post-contrast images. Bearing this in mind, we design experiments comparing the image and distribution-wise difference between our synthetic and the real post-contrast images against the difference between corresponding real pre-contrast and real post-contrast images. To quantify this difference, we use our FRD and the common FID [5], LPIPS [24] and mean squared error (MSE) metrics. Obtained results alongside ablations are summarized in Table 1. Overall, CC-Net achieves substantially better results than the pre-contrast baseline across metrics and across post-contrast phases. Surpris-




Fig. 3: Qualitative (a) and quantitative (b) test set contrast enhancement. In (b), marker size represents the standard deviation of the mean intensity within the tumor region averaged across all test cases. When normalized, tumor region mean intensity is divided by mean intensity of the remaining tumor-free pixels.

ingly, providing textual information not necessarily increases the performance of CC-Net_{Any} . Training the LDM from scratch instead of SD fine-tuning rather decreases performance (e.g., FID). The impact of longer ControlNet training is positive although marginal. Increasing the ControlNet guidance weight can be beneficial, and, in line with qualitative visual analysis, improves FRD and LPIPS. In Fig. 3b, using CC-Net_{Any} , we further analyze the pixel intensity distributions in the tumor area across phases aggregated over test cases. Synthetic images noticeably follow the principal trend of real contrast kinetics, despite a scale difference. A similar pattern is observable when taking contrast enhancements outside the tumor area into account by normalizing (dividing) mean tumor intensity by mean intensity of all other (tumor-free) pixels.

4 Discussion and Conclusion

We propose a multi-conditional latent diffusion model to translate pre-contrast into post-contrast images, thereby learning to highlight lesions by simulating their contrast uptake. We further condition the model on textual imaging metadata and continuous time passed since pre-contrast acquisition and demonstrate its synthesis capabilities on multi-sequence breast DCE-MRI data. We further contribute the Fréchet radiomics distance (FRD), a novel radiology-specific im-

age quality metric measuring the distance between real and synthetic distributions of extracted interpretable imaging biomarkers. We validate FRD demonstrating its correlation with image perturbation scales on both 3D and 2D data. In future investigations, we aim to generate images of multiple DCE-MRI time-points jointly and to map from the latent space of 3D autoencoders to the one from 2D-trained latent diffusion models. Further research avenues include confidence metrics to determine cases where synthetic contrast is sufficient, encoding additional clinical metadata using different biomedical text encoders, as well as exploring FRD feature compositions and variability. In conclusion, our work paves the way for practical applications of deep generative models in MRI as a screening modality for unsupervised tumor detection and localization from pre-contrast MRI. It further constitutes a step towards improved treatment of patient populations where invasive contrast agent injection is contraindicated.

Acknowledgments. This study has received funding from the European Union’s Horizon research and innovation programme under grant agreement No 952103 (Eu-CanImage) and No 101057699 (RadioVal). It was further partially supported by the project FUTURE-ES (PID2021-126724OB-I00) and by grant FJC2021-047659-I from the Ministry of Science and Innovation of Spain. RO acknowledges a research stay grant from the Helmholtz Information and Data Science Academy (HIDA).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Caballo, M., Sanderink, W.B., Han, L., Gao, Y., Athanasiou, A., Mann, R.M.: Four-Dimensional Machine Learning Radiomics for the Pretreatment Assessment of Breast Cancer Pathologic Complete Response to Neoadjuvant Chemotherapy in Dynamic Contrast-Enhanced MRI. *Journal of Magnetic Resonance Imaging* **57**(1), 97–110 (2023)
2. Chambon, P., Bluethgen, C., Langlotz, C.P., Chaudhari, A.: Adapting pretrained vision-language foundational models to medical imaging domains. *arXiv preprint arXiv:2210.04133* (2022)
3. European Medicines Agency (EMA): EMA’s final opinion confirms restrictions on use of linear gadolinium agents in body scans. <https://www.ema.europa.eu/> (2023), online; accessed 06 August 2023
4. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative Adversarial Nets. In: *Advances in Neural Information Processing Systems*. pp. 2672–2680 (2014)
5. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. *Advances in Neural Information Processing Systems* **30** (2017)
6. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
7. Khader, F., Müller-Franzes, G., Tayebi Arasteh, S., Han, T., Haarbuerger, C., Schulze-Hagen, M., Schad, P., Engelhardt, S., Baekler, B., Foersch, S., et al.: Denoising diffusion probabilistic models for 3D medical image generation. *Scientific Reports* **13**(1), 7303 (2023)

8. Konz, N., Chen, Y., Dong, H., Mazurowski, M.A.: Anatomically-controllable medical image generation with segmentation-guided diffusion models. arXiv preprint arXiv:2402.05210 (2024)
9. Lambin, P., Rios-Velazquez, E., Leijenaar, R., Carvalho, S., Van Stiphout, R.G., Granton, P., Zegers, C.M., Gillies, R., Boellard, R., Dekker, A., et al.: Radiomics: extracting more information from medical images using advanced feature analysis. *European Journal of Cancer* **48**(4), 441–446 (2012)
10. Müller-Franzes, G., Huck, L., Tayebi Arasteh, S., Khader, F., Han, T., Schulz, V., Dethlefsen, E., Kather, J.N., Nebelung, S., Nolte, T., et al.: Using Machine Learning to Reduce the Need for Contrast Agents in Breast MRI through Synthetic Images. *Radiology* **307**(3), e222211 (2023)
11. Olchowy, C., Cebulski, K., Łasecki, M., Chaber, R., Olchowy, A., Kalwak, K., Zaleska-Dorobisz, U.: The presence of the gadolinium-based contrast agent depositions in the brain and symptoms of gadolinium neurotoxicity-a systematic review. *PloS one* **12**(2), e0171704 (2017)
12. Osuala, R., Joshi, S., Tsirikoglou, A., Garrucho, L., Pinaya, W.H., Diaz, O., Lekadir, K.: Pre-to Post-Contrast Breast MRI Synthesis for Enhanced Tumour Segmentation. arXiv preprint arXiv:2311.10879 (2023)
13. Osuala, R., Skorupko, G., Lazrak, N., Garrucho, L., García, E., Joshi, S., Jouide, S., Rutherford, M., Prior, F., Kushibar, K., et al.: medigan: a Python library of pretrained generative models for medical image synthesis. *Journal of Medical Imaging* **10**(6), 061403–061403 (2023)
14. Pinaya, W.H., Graham, M.S., Kerfoot, E., Tudosiu, P.D., Dafflon, J., Fernandez, V., Sanchez, P., Wolleb, J., da Costa, P.F., Patel, A., et al.: Generative AI for Medical Imaging: extending the MONAI Framework. arXiv preprint arXiv:2307.15208 (2023)
15. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10684–10695 (2022)
16. Saha, A., Harowicz, M.R., Grimm, L.J., Kim, C.E., Ghate, S.V., Walsh, R., Mazurowski, M.A.: 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)
17. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *International Conference on Machine Learning*. pp. 2256–2265. PMLR (2015)
18. Song, J., Meng, C., Ermon, S.: Denoising Diffusion Implicit Models. In: *International Conference on Learning Representations* (2021), <https://openreview.net/forum?id=StlgiaarCHLP>
19. Van Griethuysen, J.J., Fedorov, A., Parmar, C., Hosny, A., Aucoin, N., Narayan, V., Beets-Tan, R.G., Fillion-Robin, J.C., Pieper, S., Aerts, H.J.: Computational radiomics system to decode the radiographic phenotype. *Cancer Research* **77**(21), e104–e107 (2017)
20. Woodland, M., Taie, M.A., Silva, J.A.M., Eltaher, M., Mohn, F., Shieh, A., Castelo, A., Kundu, S., Yung, J.P., Patel, A.B., et al.: Importance of Feature Extraction in the Calculation of Fréchet Distance for Medical Imaging. arXiv preprint arXiv:2311.13717 (2023)
21. Wu, S., Berg, W.A., Zuley, M.L., Kurland, B.F., Jankowitz, R.C., Nishikawa, R., Gur, D., Sumkin, J.H.: Breast mri contrast enhancement kinetics of normal parenchyma correlate with presence of breast cancer. *Breast Cancer Research* **18**, 1–10 (2016)

22. Xing, X., Felder, F., Nan, Y., Papanastasiou, G., Walsh, S., Yang, G.: You don't have to be perfect to be amazing: Unveil the utility of synthetic images. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 13–22. Springer Nature Switzerland, Cham (2023)
23. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3836–3847 (2023)
24. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 586–595 (2018)
25. Zhang, T., Han, L., D'Angelo, A., Wang, X., Gao, Y., Lu, C., Teuwen, J., Beets-Tan, R., Tan, T., Mann, R.: Synthesis of Contrast-Enhanced Breast MRI Using Multi-b-Value DWI-based Hierarchical Fusion Network with Attention Mechanism. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 79–88. Springer Nature Switzerland, Cham (2023)