



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

LIP-CAR: a learned inverse problem approach for medical imaging with contrast agent reduction

Davide Evangelista¹, Elena Morotti², Sonia Colombo Serra³, Pengpeng Luo⁴, Giovanni Valbusa³, and Davide Bianchi^{5*}

¹ Department of Computer Science and Engineering, University of Bologna, Bologna, 40126, Italy, davide.evangelista5@unibo.it

² Department of Political and Social Sciences, University of Bologna, Bologna, 40126, Italy, elena.morotti4@unibo.it

³ Centro Ricerche Bracco, Bracco Imaging SpA, Colliereito Giacosa, 10010, Italy, {sonia.colombo, giovanni.valbusa}@bracco.com

⁴ Bracco Imaging Medical Technologies Co. Ltd, Shanghai, 200000, China, ethe.luo@bracco.com

⁵ School of Mathematics (Zhuhai), Sun Yat-sen University, Zhuhai, 519082, China, bianchid@mail.sysu.edu.cn

Abstract. The adoption of contrast agents in medical imaging is essential for accurate diagnosis. While highly effective and characterized by an excellent safety profile, the use of contrast agents has its limitation, including rare risk of allergic reactions, potential environmental impact and economic burdens on patients and healthcare systems. This work addresses the contrast agent reduction (CAR) problem, aiming to minimize the administered dosage while preserving image quality. Unlike existing deep learning methods that simulate high-dose images from low-dose inputs via end-to-end models, we propose a learned inverse problem (LIP) approach. By learning an operator that maps high-dose to low-dose images, we reformulate CAR as an inverse problem, solved through regularized optimization to enhance data consistency. Numerical experiments on pre-clinical images demonstrate improved accuracy compared to traditional methods.

Keywords: Deep Learning · Neural Network Operator · Regularized Inverse Problem · Contrast Agent Imaging · Medical Image Computing

1 Introduction

Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) are essential in healthcare, but detecting small or low-contrast structures remains challenging. Contrast Agents (CAs), like iodine for CT and gadolinium for MRI, enhance image clarity by differentiating lesions from healthy tissue. While generally safe, minimizing CA dosage without compromising image quality is a technical challenge that could also reduce medical costs. In the following, we refer to the

* Corresponding author

Contrast Agent Reduction (CAR) problem as the task aiming to lower CA levels while maintaining diagnostic accuracy.

The outstanding advancements in deep learning have led to new digital imaging techniques for addressing the CAR problem. However, these applications are very recent, starting with [7] in 2018 and followed by studies from 2021 onward [1, 3, 4, 10, 11, 14–18, 21, 23]. This highlights that CAR is a modern challenge in medical imaging, with a limited body of existing literature. Interestingly, most of the cited methods use deep Neural Networks (NNs) to simulate high-dose images from low-dose inputs directly, i.e., through end-to-end processing. This approach can be expressed as:

$$\text{Given } \mathbf{x}_L, \text{ compute } \mathbf{x}_H^{\text{sim}} := \Psi(\mathbf{x}_L) \quad (1)$$

where $\mathbf{x}_H^{\text{sim}}$ is the simulated high-dose image computed by the neural network Ψ from the low-dose input \mathbf{x}_L . This scheme assumes the existence of a phenomenon linking \mathbf{x}_L images to the \mathbf{x}_H ones, whose formal closed form is not known exactly; thus, it is approximated by a NN. Despite the impressive performance of state-of-the-art neural networks in end-to-end imaging tasks, this accuracy often comes at the cost of robustness [2, 5, 6, 9] and the instability of deep learning models poses a significant concern in the medical field.

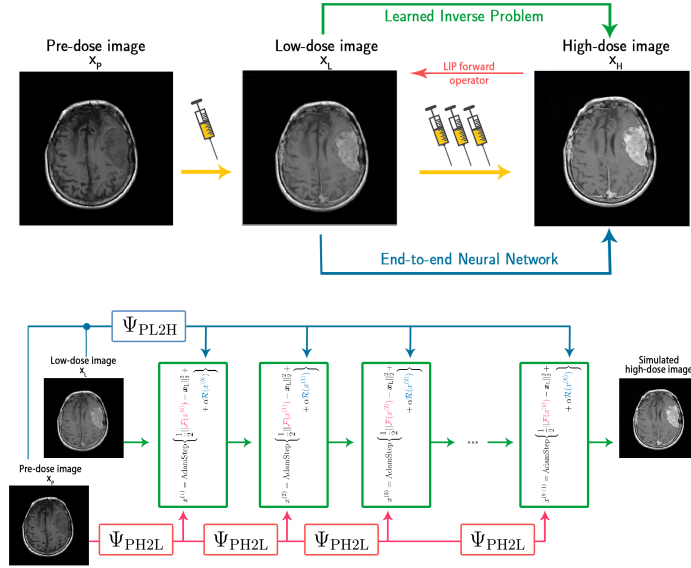


Fig. 1: Visual representation of the LIP-CAR scheme. On the top, the paradigm shift from an end-to-end approach (blue) to a Learned Inverse Problem framework (green); on the bottom, the iterative scheme of the proposed implementation based on two neural networks, Ψ_{PL2H} and Ψ_{PH2L} .

In this work, rather than using a deep network as a direct end-to-end solver, we train a neural network to implement the forward operator \mathcal{F} , mapping high-dose images to their low-dose counterparts. Indeed, we assume the existence of a second phenomenon, approximated by \mathcal{F} , which has never been explored in the literature for the CAR problem.

Our formulation allows us to re-design the digital CA-enhancement as an imaging inverse problem which can be described as:

$$\text{Given } \mathbf{x}_L, \text{ compute } \mathbf{x}_H^{\text{sim}} \text{ solving } \mathcal{F}(\mathbf{x}_H^{\text{sim}}) = \mathbf{x}_L. \quad (2)$$

To solve this inverse problem, we exploit the wide literature of model-based optimization techniques, and reformulate problem (2) into a regularized least square minimization problem. In this way, we can incorporate well-established regularization tools acting as robust image priors \mathcal{R} , which further ensure robustness and accuracy in digitally simulated high-dose images. The proposed formulation to address CAR imaging problem therefore reads:

$$\text{Given } \mathbf{x}_L, \text{ solve } \mathbf{x}_H^{\text{sim}} \in \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ \frac{1}{2} \|\mathcal{F}(\mathbf{x}) - \mathbf{x}_L\|_2^2 + \alpha \mathcal{R}(\mathbf{x}) \right\}, \quad (3)$$

where $\alpha > 0$ balances the trade-off between data consistency and the regularization effect. As we will detail in Section 2, in our proposal we select a regularizer $\mathcal{R}(\mathbf{x})$ that exploits the ability of low-dose to high-dose neural networks (described in (1)) to guide the optimization algorithm toward an accurate solution, while preserving the robustness guaranteed by the inverse problem literature. We denote the entire framework as LIP-CAR (Learned Inverse Problem for Contrast Agent Reduction), and it is shown in Figure 1.

Contributions. We introduce three key innovations in our approach:

- We assign a novel role to a neural network in the context of the CAR imaging by training it to approximate the unknown forward mapping between high-dose to low-dose images, thereby defining the forward operator of a blind non-linear imaging problem.
- Leveraging the learned forward operator introduced in the previous point, we propose a learned inverse problem formulation for the CAR imaging task, where high-dose images are reconstructed through a regularized inverse problem approach, ensuring data consistency and stability. This represents a paradigm shift from the literature on CAR, and it introduces well-established mathematical techniques for the first time.
- We validate LIP-CAR on a pre-clinical MRI dataset constituted by real acquisitions provided by an industry partner, and the results demonstrate that LIP-CAR outperforms state-of-the-art methods in reconstruction accuracy and robustness. This proves the factual effectiveness and practical applicability of the proposed scheme.

2 Method

In this section, we first set the notations and contextualize our research into the state-of-the-art, then we detail our proposal.

Notations. As previously introduced, we denote the digitally-simulated high-contrasted image as $\mathbf{x}_H^{\text{sim}}$, representing the solution to the CAR imaging problem. In addition, we label as \mathbf{x}_H each high-dose image, corresponding to the injection of a standard reference level of CA, as established either by a regulatory body or by laboratory protocols. Conversely, \mathbf{x}_L refers to low-dose images, acquired when a reduced amount of CA is administered to the patient. In some clinical trials, it is also possible to obtain images before any CA injection. We refer to this type of data as "pre-dose" images, denoted by \mathbf{x}_P .

All images involved in this study can be acquired through tomographic procedures or MRI techniques and are subsequently reconstructed as grayscale images. For this work, we focus on 2D slices of size $H \times W$, although the framework can be extended to 3D imaging with minor modifications. To simplify the processing, we normalize all images to the range $[0, 1]$. We define \mathcal{X} as the set of pixel-wise non-negative images, i.e., $\mathcal{X} = \mathbb{R}_+^{H \times W}$ and we assume that $\mathbf{x}_P, \mathbf{x}_L, \mathbf{x}_H \in \mathcal{X}$.

2.1 Background of CAR imaging

Since the 1990s, mathematical models have been developed to describe the distribution-excretion of contrast agents (CAs), notably in Tofts et al. [20]. However, the exact mechanisms remain unclear due to the complexity of modeling diverse physical and chemical interactions in the body. The advent of deep learning has reignited interest in CA-related biomedical imaging, as neural networks, being universal approximators [12], show promise in capturing these complex dynamics. Specifically relevant to our work, [3] leverages a V-Net to generate high-dose MRI brain images from low-dose scans using just 25% of the standard gadolinium-based CA dose. Further examples of NNs usages for CAR are provided in [1, 4, 7, 10, 14, 15, 17, 23].

These methods strongly rely on NNs, as they propose to compute the digitally enhanced CA image through an inference on one end-to-end NN, hereafter called Ψ , as stated in Equation (1). Precisely, the Ψ operator is trained to approximate the \mathbf{x}_H images, based on the availability of registered data (both for training and for inference). When only the \mathbf{x}_L and \mathbf{x}_H images are provided, the end-to-end network reads:

$$\Psi_{L2H}: \mathbf{x}_L \mapsto \mathbf{x}_H \quad (4)$$

as it directly maps the low-dose image to the desired one. When also the pre-dose images are provided, the comparison between \mathbf{x}_L and \mathbf{x}_P offers important hints for the contrast enhancement of the CAR imaging task. Consequently, the Ψ operator becomes:

$$\Psi_{PL2H}: (\mathbf{x}_P, \mathbf{x}_L) \mapsto \mathbf{x}_H, \quad (5)$$

where the notation remarks that the network expects both the images in input. More recently, deep generative models (and, in particular, GANs [8]) have been successfully applied to the task of CAR, obtaining remarkable results [11, 16]. However, due to the inherent mechanism driven by generator/discriminator duality, these models are prone to hallucinations [22] and their use in the field of medical imaging is still limited.

2.2 The learned image-to-image operator

To implement our LIP-CAR approach, the first task is to approximate the forward operator $\mathcal{F}: \mathcal{X} \rightarrow \mathcal{X}$, introduced in Equation (2). In this work, we propose learning \mathcal{F} through a neural network that is used in inference mode at each iteration, as denoted in Figure 1. When the pre-dose images are not available, we can approximate it as:

$$\Psi_{\text{H2L}}: \mathbf{x}_\text{H} \mapsto \mathbf{x}_\text{L}, \quad (6)$$

whereas

$$\Psi_{\text{PH2L}}: (\mathbf{x}_\text{P}, \mathbf{x}_\text{H}) \mapsto \mathbf{x}_\text{L} \quad (7)$$

is preferable whenever \mathbf{x}_P samples are provided.

We remark that the comparison between \mathbf{x}_P and \mathbf{x}_H , i.e., between images having minimal and maximum contrast agent, is expected to be very informative for simulating the target low-dose image, as the two input elements of Ψ_{PH2L} already contain much of the structural information. This may give Ψ_{PH2L} an inherent advantage, with respect to the end-to-end aforementioned approaches.

Network architecture and training. In the experiments, to allow for a fair comparison with the state-of-the-art, we employ a VNet-like architecture for our Ψ_{H2L} model, as in [3]. The VNet is a variant of the popular UNet [19] where strided convolutions are used for downscaling instead of maxpooling. When employed, the pre-dose data is fed into the network by concatenating it with the high-dose image. Note that, since \mathbf{x}_P and \mathbf{x}_H have the same shape, this operation does not require any extra process. The resulting model is then trained for a total of 80 epochs and a batch size of 16, employing the Adam optimization algorithm [13] with a fixed step size of 10^{-3} . We consider an SSIM-based loss function [24] for training. Therefore, we consider:

$$\mathcal{L}(\Psi_{\text{PH2L}}) = \frac{1}{N} \sum_{i=1}^N (1 - \text{SSIM}(\Psi_{\text{PH2L}}(\mathbf{x}_\text{P}^i, \mathbf{x}_\text{H}^i), \mathbf{x}_\text{L}^i)), \quad (8)$$

where the index i enumerates the samples in our training set, for the Ψ_{PH2L} network, and a similar loss for Ψ_{H2L} .

2.3 The reformulation into an optimization problem

The presence of the forward operator allows the CAR reformulation as the inverse problem stated in (2), which can be solved through the problem statement

as in (3). Indeed, the neural network structure of Ψ_{H2L} and Ψ_{PH2L} guarantees continuity and, under suitable assumptions on the activation functions, differentiability, making them suitable for first-order optimization algorithms. We highlight that such minimization problem must be solved with an iterative procedure, to be set according to the mathematical properties of the functional to minimize.

It is known in the literature that solving an inverse problem through an optimization reformulation can offer advantages in terms of efficiency, flexibility, and robustness compared to direct inversion methods. In expression (3), in fact, the least square term forces the data fidelity with the simulated high-dose image, whereas the component $\mathcal{R}(\mathbf{x})$ serves as a regularization term that incorporates desired a-priori information on the solution and penalizes all the candidate solutions that do not present the features we aim to recover.

The choice of \mathcal{R} is of crucial importance. To set it, we exploit the well-established regularization techniques that have already been extensively used in medical image processing to get a reliable solution. We thus select the Generalized ℓ_1 -Total Variation (GenTV) regularizer, which ensures that the solution image closely resembles a possibly good guess of $\mathbf{x}_{\text{H}}^{\text{sim}}$, in terms of gradient transforms. Denoting that guess as $\overline{\mathbf{x}}_{\text{H}}$, GenTV reads:

$$\mathcal{R}(\mathbf{x}) = \|\nabla(\mathbf{x} - \overline{\mathbf{x}}_{\text{H}})\|_1, \quad (9)$$

where ∇ denotes the gradient operator. In our scenario, GenTV can easily take advantage of the end-to-end networks, and $\overline{\mathbf{x}}_{\text{H}}$ can be computed as $\Psi_{\text{L2H}}(\mathbf{x}_{\text{L}})$ or $\Psi_{\text{PL2H}}(\mathbf{x}_{\text{L}})$. As depicted in Figure 1, in our iterative resolution, Ψ_{PL2H} undergoes a single forward pass to plug its output in the regularization term.

At last, we remark that our selection of \mathcal{R} is merely illustrative since the LIP-CAR approach lets us accommodate models other than (3), both for the data-fidelity operator, the regularizer and the feasible set \mathcal{X} .

3 Numerical Experiments

3.1 Dataset

We conduct our study on real CA images from a pre-clinical trial. Procedures were conducted according to the national and international laws on experimental animal research (L.D. 26/2014; Directive 2010/63/EU) and under a specific Italian Ministerial Authorization (project research number 215/2020-PR), by CRB/Test Facility of Bracco Imaging S.p.A. The original dataset is composed of 61 cranial MRI examinations of lab rats affected by an orthotopically induced C6 glioma. Each lab rat underwent an MRI session composed of three T1-weighted spin echo sequences, respectively acquired before any administration (pre-dose images), after administration of 0.01 mmol Gd/kg of a non-commercial high relaxivity dimeric gadolinium based CA (low-dose images), and after administration of additional 0.04 mmol Gd/kg of the same agent resulting in a full-dose injection (high-dose images). We split the original 3D volumes into 2D images

of 256×256 pixels. After cleaning the dataset, we randomly selected 840 images for training and used the remaining 240 as testing subsets.

Some exemplar \mathbf{x}_H , \mathbf{x}_P , and \mathbf{x}_L images from the test set are reported in Figure 2. In the first two sequences, a glioma is present and visible only on the images with injected CA. The last two cases do not report tumors and are therefore useful for investigating whether the approaches add false positives. Additionally, the third case shares similar shapes with the previous samples, whereas the last one shows a completely diverse head morphology including the olfactory bulb.

3.2 Results and comparisons

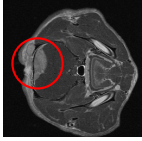
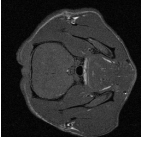
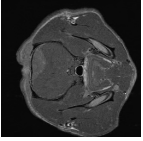
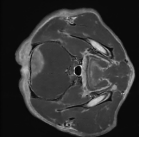
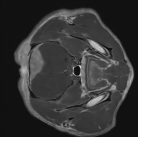
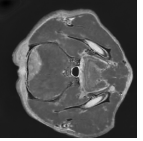
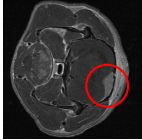
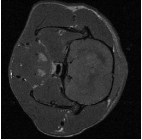
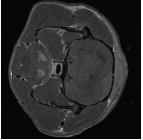
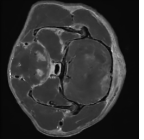
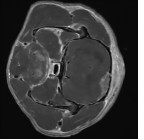
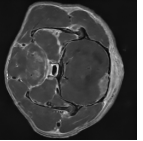
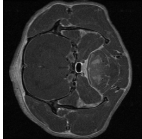
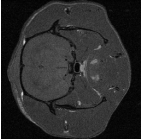
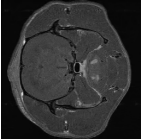
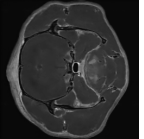
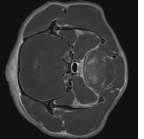
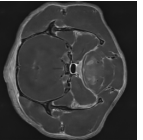
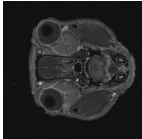
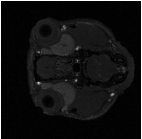
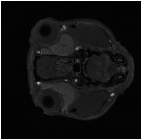
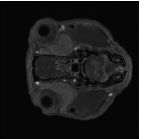
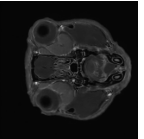
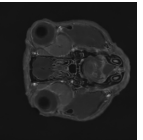
	\mathbf{x}_H	\mathbf{x}_P	\mathbf{x}_L	NN-L2H	NN-PL2H	LIP-CAR
Image 39	-	0.4118	0.6627	0.8646	0.8756	0.9007
						
Image 177	-	0.4284	0.6764	0.8534	0.8722	0.8989
						
Image 226	-	0.4512	0.6821	0.8680	0.8804	0.9020
						
Image 141	-	0.3697	0.6330	0.8346	0.8488	0.8883
						

Fig. 2: Results on test samples. From left to right: the \mathbf{x}_H , \mathbf{x}_P and \mathbf{x}_L images as references, the simulated high-dose computed by **NN-L2H** and **NN-PL2H**, and the LIP-CAR solutions. The number above each image indicates its SSIM value relative to the target \mathbf{x}_H . The red circles highlight tumoral masses, expected to be more discernible in the simulated high-dose images than in \mathbf{x}_L .

We run the proposed LIP-CAR framework on test samples and compare its high-dose simulated images to those generated by some state-of-the-art networks.




Fig. 3: Comparison of metrics computed with respect to \mathbf{x}_H on the test samples, among the original input images and the solutions by the considered methods. From left to right: SSIM and PSNR (the higher, the better), and RMSE (the lower, the better).

Some results are reported in Figure 2 for visual inspection, whereas Figure 3 shows boxplots to summarize the metrics evaluated on all the test set.

Here, the label **LIP-CAR** refers to the version where the \mathcal{F} operator is played by Ψ_{PH2L} , and the GenTV regularizer defined in Equation (9) relies on the $\overline{\mathbf{x}}_H$ image computed by the Ψ_{PL2H} operator. We observe that this implementation heavily exploits all the provided items and tools at best, by using the pre-dose images both for training the LIP operator and for the regularization term. To solve the regularized inverse problem we have exploited the widely used **Adam** solver [13]. In particular, the optimizer is executed for 150 iterations, with \mathbf{x}_L as starting iterate, and with default Pytorch parameters. We selected $\alpha = 6 \cdot 10^{-3}$ as a regularization parameter; it has been chosen to better balance the visual appearance of the output and the metrics over all the test samples, heuristically. Further details about the implementation and parameter choice can be found on Github, at <https://github.com/devangelista2/LIP-CAR/>.

As state-of-the-art models for comparison, we use the end-to-end Ψ_{L2H} and Ψ_{PL2H} operators, respectively defined in Equations (4) and (5). We refer to these methods as **NN-L2H** and **NN-PL2H**, respectively, to emphasize their reliance on the direct application of NN, in contrast to the LIP formulation. They are based on the same VNet architecture used in our Ψ_{PH2L} and detailed in Section 2.2, for fair comparison.

The results reported in Figure 2 show that all the considered frameworks manage to approximate the high-dose images carefully, emphasizing the contrast of tumoral masses and adding no false enhancements (potentially causing hallucinations and false diagnosis). However, **LIP-CAR** overcomes NN-based approaches, computing sharp edges and always getting the highest SSIM values. Indeed, Figure 3 firmly remarks the LIP-CAR success in terms of both SSIM, PSNR (the higher, the better) and RMSE metrics (the lower, the better). The boxplots referring to **LIP-CAR** (the green ones) consistently denote that we hit the best values, overcoming **NN-L2H** and **NN-PL2H** (respectively represented by the cyan and blue boxplots) and strongly improving the quality of the acquired \mathbf{x}_P and \mathbf{x}_L images (red and orange boxplots).

We thus conclude empirically that our mathematically-grounded regularized approach is highly competitive with state-of-the-art technology.

4 Conclusions

This paper introduces a novel application of neural networks for imaging. LIP-CAR tackles the contrast-agent reduction imaging task as an inverse problem, where a convolutional network has learned the forward imaging operator. The inverse problem is solved with mathematically grounded tools belonging to the well-established class of optimization and regularization techniques. Numerical experiments performed on real MRI preclinical images and comparisons with some state-of-the-art schemes demonstrate the great potential of the LIP-CAR approach. LIP-CAR outperformed the state-of-the-art schemes in terms of the quality of simulated high-dose images.

We emphasize that this work represents an introductory study of the learned inverse problem framework, and that there remains significant room for improvement. While the ResUNet architecture leads to a non-convex objective function that does not guarantee convergence to a global minimum, the exploration of architectures that incorporate convexity or smoother behavior could further improve both the robustness and interpretability of the optimization process. We also plan to investigate alternative optimization algorithms beyond Adam, with the goal of accelerating convergence. Finally, we aim to conduct a structured analysis of the regularizer, with the objective of designing a functional that more effectively captures the global structure of the target image, similar in spirit to the GenTV regularizer, and to develop data-driven strategies for automatically selecting the regularization parameter based on the input.

Acknowledgments. The results presented in this study were funded by Bracco Imaging Medical Technologies Co., Ltd. E. Morotti is supported by the PRIN 2022 project 20225STXSB, funded by the European Commission under the NextGeneration EU program. D. Bianchi is supported by the Startup Fund of Sun Yat-sen University.

Disclosure of Interests. The authors Sonia Colombo Serra, Pengpeng Luo, and Giovanni Valbusa are employees of Bracco Group.

References

1. Ammari, S., Bône, A., Balleyguier, C., Moulton, E., Chouzenoux, É., Volk, A., et al.: Can deep learning replace gadolinium in neuro-oncology?: a reader study. *Investigative Radiology* **57**(2), 99–107 (2022)
2. Antun, V., Renna, F., Poon, C., Adcock, B., Hansen, A.C.: On instabilities of deep learning in image reconstruction and the potential costs of AI. *PNAS* **117**(48), 30088–30095 (2020)
3. Bône, A., Ammari, S., Lamarque, J.P., Elhaik, M., Chouzenoux, É., Nicolas, F., et al.: Contrast-enhanced brain MRI synthesis with deep learning: key input modalities and asymptotic performance. In: 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI). pp. 1159–1163. IEEE (2021)

4. Bône, A., Ammari, S., Menu, Y., Balleyguier, C., Moulton, E., Chouzenoux, E., et al.: From Dose Reduction to Contrast Maximization: Can Deep Learning Amplify the Impact of Contrast Media on Brain Magnetic Resonance Image Quality? A Reader Study. *Investigative Radiology* pp. 10–1097 (2022)
5. Colbrook, M.J., Antun, V., Hansen, A.C.: Can stable and accurate neural networks be computed?—On the barriers of deep learning and Smale’s 18th problem (2021)
6. Evangelista, D., Loli Piccolomini, E., Morotti, E., Nagy, J.G.: To be or not to be stable, that is the question: understanding neural networks for inverse problems. *SIAM Journal on Scientific Computing* **47**(1), C77–C99 (2025)
7. Gong, E., Pauly, J.M., Wintermark, M., Zaharchuk, G.: Deep learning enables reduced gadolinium dose for contrast-enhanced brain MRI. *Journal of magnetic resonance imaging* **48**(2), 330–340 (2018)
8. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al.: Generative adversarial networks. *Communications of the ACM* **63**(11), 139–144 (2020)
9. Gottschling, N.M., Antun, V., Adcock, B., Hansen, A.C.: The troublesome kernel: why deep learning for inverse problems is typically unstable (2020)
10. Haase, R., Pinetz, T., Bendella, Z., Kobler, E., Paech, D., Block, W., et al.: Reduction of gadolinium-based contrast agents in MRI using convolutional neural networks and different input protocols: limited interchangeability of synthesized sequences with original full-dose images despite excellent quantitative performance. *Investigative Radiology* **58**(6), 420–430 (2023)
11. Haubold, J., Hosch, R., Umutlu, L., Wetter, A., Haubold, P., Radbruch, A., et al.: Contrast agent dose reduction in computed tomography with deep learning using a conditional generative adversarial network. *Eur. Radiol.* **31**(8), 6087–6095 (2021)
12. Hornik, K., Stinchcombe, M., White, H.: Multilayer feedforward networks are universal approximators. *Neural networks* **2**(5), 359–366 (1989)
13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
14. Luo, H., Zhang, T., Gong, N.J., Tamir, J., Venkata, S.P., Xu, C., et al.: Deep learning-based methods may minimize GBCA dosage in brain MRI. *European Radiology* **31**(9), 6419–6428 (2021)
15. Montalt-Tordera, J., Quail, M., Steeden, J.A., Muthurangu, V.: Reducing contrast agent dose in cardiovascular MR angiography with deep learning. *Journal of Magnetic Resonance Imaging* **54**(3), 795–805 (2021)
16. Müller-Franzes, G., Huck, L., Tayebi Arasteh, S., Khader, F., Han, T., Schulz, V., et al.: Using machine learning to reduce the need for contrast agents in breast MRI through synthetic images. *Radiology* **307**(3), e222211 (2023)
17. Pasumarthi, S., Tamir, J.I., Christensen, S., Zaharchuk, G., Zhang, T., Gong, E.: A generic deep learning model for reduced gadolinium dose in contrast-enhanced brain MRI. *Magn. Reson. Med.* **86**(3), 1687–1700 (2021)
18. Pinetz, T., Kobler, E., Haase, R., Deike-Hofmann, K., Radbruch, A., Effland, A.: Faithful synthesis of low-dose contrast-enhanced brain mri scans using noise-preserving conditional gans. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 607–617. Springer (2023)
19. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18. pp. 234–241. Springer (2015)

20. Tofts, P.S., Brix, G., Buckley, D.L., Evelhoch, J.L., Henderson, E., Knopp, M.V., et al.: Estimating kinetic parameters from dynamic contrast-enhanced T1-weighted MRI of a diffusable tracer: standardized quantities and symbols. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* **10**(3), 223–232 (1999)
21. Wang, D., Pasumarthi, S., Zaharchuk, G., Chamberlain, R.: Simulation of arbitrary level contrast dose in mri using an iterative global transformer model. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 88–98. Springer (2023)
22. Wang, H., Zhang, X., Li, T., Wan, Y., Chen, T., Sun, J.: Dmplug: A plug-in method for solving inverse problems with diffusion models. *arXiv:2405.16749* (2024)
23. Wang, Y., Wu, W., Yang, Y., Hu, H., Yu, S., Dong, X., et al.: Deep learning-based 3D MRI contrast-enhanced synthesis from a 2D non contrast T2Flair sequence. *Medical Physics* **49**(7), 4478–4493 (2022)
24. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)