

CAVM: Conditional Autoregressive Vision Model for Contrast-Enhanced Brain Tumor MRI Synthesis

Lujun Gui^{1,*}, Chuyang Ye^{2,*}, and Tianyi Yan^{1(✉)}

¹ School of Medical Technology, Beijing Institute of Technology, Beijing, China
yantianyi@bit.edu.cn

² School of Integrated Circuits and Electronics, Beijing Institute of Technology, Beijing, China

Abstract. Contrast-enhanced *magnetic resonance imaging* (MRI) is pivotal in the pipeline of brain tumor segmentation and analysis. Gadolinium-based contrast agents, as the most commonly used contrast agents, are expensive and may have potential side effects, and it is desired to obtain contrast-enhanced brain tumor MRI scans without the actual use of contrast agents. Generic deep learning methods have been applied to synthesize virtual contrast-enhanced MRI scans from non-contrast images. However, as this synthesis problem is inherently ill-posed, these methods fall short in producing high-quality results. In this work, we propose *Conditional Autoregressive Vision Model* (CAVM) for improving the synthesis of contrast-enhanced brain tumor MRI. As the enhancement of image intensity grows with a higher dose of contrast agents, we assume that it is less challenging to synthesize a virtual image with a lower dose, where the difference between the contrast-enhanced and non-contrast images is smaller. Thus, CAVM gradually increases the contrast agent dosage and produces higher-dose images based on previous lower-dose ones until the final desired dose is achieved. Inspired by the resemblance between the gradual dose increase and the Chain-of-Thought approach in natural language processing, CAVM uses an autoregressive strategy with a decomposition tokenizer and a decoder. Specifically, the tokenizer is applied to obtain a more compact image representation for computational efficiency, and it decomposes the image into dose-variant and dose-invariant tokens. Then, a masked self-attention mechanism is developed for autoregression that gradually increases the dose of the virtual image based on the dose-variant tokens. Finally, the updated dose-variant tokens corresponding to the desired dose are decoded together with dose-invariant tokens to produce the final contrast-enhanced MRI. CAVM was validated on the publicly available BraSyn-2023 dataset with brain tumor MRI scans, where it outperforms several state-of-the-art methods for medical image synthesis. The code is available at <https://github.com/Luc4Gui/CAVM>.

Keywords: Medical image synthesis · Contrast-enhanced MRI · Autoregressive model.

*L. Gui and C. Ye: Co-first authors, equal contribution.

1 Introduction

Gadolinium-based contrast agents significantly improve the diagnostic accuracy based on brain *magnetic resonance imaging* (MRI) by enhancing the contrast of vascular structures and pathological conditions [25]. The contrast-enhanced brain MRI scan is crucial in detecting and characterizing brain tumors [1]. However, the use of gadolinium-based contrast agents requires careful consideration of potential risks, such as nephrogenic systemic fibrosis in patients with renal impairment and gadolinium deposition in the brain [19]. To avoid the use of contrast agents, cross-modality MRI synthesis models are developed to estimate contrast-enhanced brain tumor MRI scans from available non-contrast modalities [3].

Most current cross-modality MRI synthesis models [11,16] are based on deep learning, and they are trained with paired input source and output target modality images. For example, ResViT [2] is a *generative adversarial network* (GAN) approach for cross-modality MRI synthesis, which harnesses the contextual sensitivity of vision Transformers [4], combined with the precision of convolutional operators and the authenticity encouraged by adversarial learning; SynDiff [17] is based on the diffusion model [9], where a conditional diffusion process is used and large diffusion steps are replaced by adversarial projections for computational efficiency; the task-specific sequence-to-sequence network TSF-Seq2Seq [6] transforms cross-modality MRI synthesis from an image-to-image problem into a sequence-to-sequence problem by fusing multi-modal information; MT-Net [15] is a cross-modality MRI synthesis framework with edge-aware self-supervised pre-training, which can alleviate the challenge of limited paired data. However, the performance of these methods is still unsatisfactory for contrast-enhanced brain tumor MRI synthesis from non-contrast MRI, as the problem is inherently ill-posed due to the lack of contrast in non-contrast MRI.

Unlike general cross-modality MRI synthesis, the problem of contrast-enhanced MRI synthesis features a unique continuous variable, the contrast agent dose, where the difference between contrast-enhanced MRI and non-contrast MRI is greater with an increasing dose [25]. Yet this property has not been explored by previous works on the contrast-enhanced MRI synthesis problem. Based on this relationship between image contrast and the contrast agent dose, we assume that it is less challenging to synthesize a virtual contrast-enhanced image with a lower dose and seek to accordingly further explore the challenging problem of contrast-enhanced brain tumor MRI synthesis from non-contrast MRI.

Specifically, we propose the *Conditional Autoregressive Vision Model* (CAVM), which synthesizes contrast-enhanced MRI by gradually increasing the dose until the final desired dose is achieved. Motivated by the resemblance between the gradual dose increase and the Chain-of-Thought approach [24] in natural language processing, CAVM uses an autoregressive strategy with a decomposition tokenizer and a decoder. In CAVM, the tokenizer is first applied to the input images to obtain a more compact image representation for computational efficiency, and it decomposes the image into dose-variant and dose-invariant tokens. Then, a masked self-attention mechanism is developed based on LLaMA-

style [21] Transformer blocks for autoregression, which gradually produces images of increased doses with the dose-variant tokens of non-contrast images and intermediate lower-dose images that are already produced. Finally, the updated dose-variant tokens associated with the desired dose are decoded together with dose-invariant tokens to produce the final contrast-enhanced MRI. CAVM was validated on the publicly available BraSyn-2023 dataset [14] and compared with state-of-the-art medical image synthesis methods. The results show that our method improves the synthesis quality by noticeable margins.

2 Methodology

2.1 Problem Formulation and Method Overview

Given non-contrast brain tumor MRI scans of different modalities, including T1-weighted (T1w), T2-weighted (T2w), and *fluid-attenuated inversion recovery* (FLAIR) images, our work aims to synthesize a corresponding T1 gadolinium-based contrast-enhanced (T1Gd) image, so that contrast-enhanced brain tumor MRI can be obtained without actually injecting contrast agents. We also assume that the brain tumor mask is available for assistance, where a model can be trained to reliably perform brain tumor segmentation on the T1w, T2w, and FLAIR images. For convenience, we denote the non-contrast images by x_{NC} , the tumor mask by x_{TM} , and the standard-dose T1Gd image by y_{SD} . By concatenating x_{NC} and x_{TM} , we obtain x , which is the overall input to the synthesis model.

To reduce the difficulty in synthesizing contrast-enhanced brain tumor MRI from non-contrast MRI, we transform the synthesis of contrast-enhanced information into three sequential steps, where contrast-enhanced information is gradually recovered, first for a lower dose, then for a higher dose, and finally for the standard dose. We propose an autoregressive model CAVM for the gradual dose increase, the design of which is illustrated in Fig. 1. The input images are tokenized first and decomposed into dose-variant and dose-invariant tokens. The dose-variant tokens are fed into an autoregression module that gradually increases the dose. Then, the updated dose-variant tokens are decoded with the dose-invariant tokens to produce the target contrast-enhance brain tumor MRI. The detailed design of the model architecture is described below.

2.2 Model Architecture

Decomposition Tokenizer Since contrast agents do not alter the intensities of all image regions, we decompose the input images into a dose-variant component and a dose-invariant component, so that the synthesis problem can focus on the relevant part and its difficulty can be reduced. Moreover, as the dimension of the original image is high, for computational efficiency, the dose-variant and dose-invariant components are represented by lower-dimensional tokens. The dose-variant and dose-invariant tokens are obtained with a decomposition tokenizer,




Fig. 1. The overall architecture of CAVM. Decomposition Tokenizer comprises two encoders located on the left. The Dose-variant Autoregression is implemented by LLaMA-style Transformer and two encoders situated at the bottom right. The decoder, positioned at the top right of the diagram, decodes all output images from image tokens during the autoregressive process.

which is built upon the encoder of Swin UNETR [8], a backbone proven powerful in both medical image segmentation and generation, to ensure the quality of the tokens. The decomposition tokenizer comprises two distinct units, f_{DV} for dose-variant tokenization and f_{DI} for dose-invariant tokenization. They comprise two and four stages of the Swin UNETR encoder, respectively, and extract the outputs of the last two stages as tokens.

Dose-variant Autoregression The autoregression module then applies gradual dose increase by autoregressing the dose-variant tokens, where tokens associated with higher doses are obtained sequentially. Multiple steps are used in the autoregression, where at each step, a higher-dose token is obtained based on the non-contrast dose-variant tokens and all lower-dose tokens that have been obtained, until the desired dose is reached. Here, we choose to use three autoregression steps, which balance the computation complexity and reduction in regression difficulty. This leads to the sequential synthesis of a lower-dose image y_{LD} , a higher-dose image y_{HD} , and the standard-dose image y_{SD} .

We construct the autoregression module based on LLaMA-style Transformer blocks with RMSNorm normalization and rotary position embedding [21], as they have achieved excellent performance among Transformer-based architectures in both computer vision and natural language processing [13,22]. Mathematically, suppose the dose-variant tokens of x given by f_{DV} is t_x , and the dose-variant

tokens of y_{LD} , y_{HD} , and y_{SD} are t_{LD} , t_{HD} , and t_{SD} , respectively; then the autoregression is formulated as

$$t_{LD}, t_0, t_0 = \text{MMHSA}(\text{RoPE}(W \times (t_x, t_0, t_0))) + (t_x, t_0, t_0) \quad (1)$$

$$t_{LD}, t_{HD}, t_0 = \text{MMHSA}(\text{RoPE}(W \times (t_x, \tilde{t}_{LD}, t_0))) + (t_x, \tilde{t}_{LD}, t_0) \quad (2)$$

$$t_{LD}, t_{HD}, t_{SD} = \text{MMHSA}(\text{RoPE}(W \times (t_x, \tilde{t}_{LD}, \tilde{t}_{HD}))) + (t_x, \tilde{t}_{LD}, \tilde{t}_{HD}) \quad (3)$$

Here, t_0 represents placeholder tokens, W is a linear weighting layer, $\text{RoPE}(\cdot)$ denotes rotary position embedding [20], and $\text{MMHSA}(\cdot)$ denotes the masked multi-head self-attention [23]; moreover, instead of using the autoregressed t_{LD} and t_{HD} directly for subsequently autoregression steps, we use their updated versions, $\tilde{t}_{LD}$ and $\tilde{t}_{HD}$ respectively. $\tilde{t}_{LD}$ and $\tilde{t}_{HD}$ are obtained by feeding t_{LD} and t_{HD} into a decoder for image reconstruction (to be introduced next) and then using a shared tokenizer f_{CE} to encode the reconstructed low-dose and high-dose contrast-enhanced images, respectively. This update allows dose-invariant tokens to correct the potential errors in autoregressed dose-variant tokens. f_{CE} comprises four stages of the Swin UNETR encoder and gives the outputs of the two last layers.

Note that in CAVM, we revise the original design of the self-attention mask in MMHSA, as the default lower triangular self-attention mask only uses the tokens in the previous state. To fully exploit all tokens of previous states, the self-attention mask is designed with a staircase shape (also shown in Fig. 1). Specifically, the i -th token z_i in (t_{LD}, t_{HD}, t_{SD}) is computed as

$$z_i = \sum_{k=1}^K a_{ik} v_k, \quad (4)$$

where K is the total number of tokens, v_k is the value embedding, and a_{ik} is the attention weight associated with v_k . Suppose n is the number of tokens per image; then the masking scheme is achieved by setting $a_{ik} = 0$ for $\lfloor k/n \rfloor > \lfloor i/n \rfloor$ and $\sum_{k=1}^K a_{ik} = 1$. Another important benefit of the revised masking is the inference speedup. With this staircase masking, the whole image can be processed with a single forward pass instead of one forward pass for each single token, which is computationally impractical for our task.

Decoder for Reconstruction To obtain the contrast-enhanced images from the tokens with gradually increased doses, a single decoder f_D combines the dose-invariant tokens t_x of the input images with t_{LD} , t_{HD} , and t_{SD} and produces the estimated low-dose, high-dose, and standard-dose images $\hat{y}_{LD}$, $\hat{y}_{HD}$, and $\hat{y}_{SD}$, respectively. f_D uses four stages of the decoder of Swin UNETR [8] for image reconstruction.

2.3 Model Training

We first pretrain the tokenizer f_{DI} , f_{DV} , and f_{CE} with two joint tasks. The first task is an autoencoding task for contrast-enhanced images y_{LD} , y_{HD} , and

y_{SD} , where f_{DI} and f_{CE} are involved. Note that although the training data only provides non-contrast images and standard-dose contrast-enhanced brain tumor MRI, it is possible to generate synthetic lower-dose (33% dose) and higher-dose (66% dose) images from them for autoencoding with the approach in [18]. The second task is an image-to-image prediction task from x to y_{SD} , where f_{DI} and f_{DV} are involved. For each task, the L1 loss and adversarial loss are used for training, where we use the discriminator from [2] and extend it to 3D. The overall training loss is the direct sum of the two losses. Then, we train the dose-variant autoregression module with the l_2 regression loss from AIM [5] and MAE [7] based on the predicted lower-dose, higher-dose, standard-dose images.

The hyperparameter settings of the LLaMA-style Transformer blocks are as follows. We set the embedding dimensions to 384 and 768, multi-attention heads to 4 and 8, LLaMA-style Transformer layers to 4 and 8, and maximum sequence lengths to 432 and 864. Model training was performed with Adam ($\beta_1 = 0.9$ and $\beta_2 = 0.99$) [12], where we set the batch size to 1 and peak learning rate to 10^{-4} .

3 Experiments and Results

3.1 Experimental Setups

We evaluated CAVM on the public BraSyn-2023 dataset [14], which comprised aligned multi-contrast brain tumor MRI scans of 1,470 subjects, as well as the tumor annotations. The image modalities included the T1w, T2w, FLAIR, and T1Gd images. We randomly selected 1,043, 219, and 208 subjects as the training, validation, and test set, respectively. All images were center-cropped to $192 \times 192 \times 128$ and normalized by dividing the image intensities by the value at the 95th percentile of the intensity distribution. To provide the tumor mask input to CAVM, a nnU-Net model [10] was trained with the T1w, T2w, and FLAIR images and annotations in the training set.

We computed the *Structural Similarity Index Measure* (SSIM) and *Peak Signal-to-Noise Ratio* (PSNR) to assess image quality. These two metrics were summarized for the tumor region and other healthy brain tissue separately. To assess the utility of CAVM for downstream analysis, brain tumor segmentation was performed with another nnU-Net model trained with the T1w, T2w, FLAIR, and T1Gd images and annotations in the training set. This model was applied to the real T1w, T2w, and FLAIR images and synthesized T1Gd images of the test set, and the Dice coefficient was computed for the segmentation result.

3.2 Performance Evaluation

We compared CAVM with four state-of-the-art models for multi-contrast brain MRI synthesis, including a GAN-based method ResViT [2], a diffusion model SynDiff [17], a sequence-to-sequence model TSF-Seq2Seq [6], and the autoencoder pretrained MT-Net [15]. For a fair comparison, for models (ResViT, SynDiff, and MT-Net) that can incorporate brain tumor masks into the input, their




Fig. 2. Four examples of real T1Gd images and synthesized results. For CAVM, from left to right the image order is lower-dose, higher-dose, and standard-dose. Note the highlighted tumor region for comparison.

inputs were identical to those of CAVM; otherwise, the input only comprised the non-contrast images.

Fig. 2 presents examples of synthesized T1Gd images produced by different models. ResViT and SynDiff failed to reproduce the enhanced intensity in the tumor. TSF-Seq2Seq and MT-Net produced enhanced intensities, but the results are not consistent with the real images. By contrast, CAVM generated accurate T1Gd images with faithful contrast enhancement in the tumor. Also, the intermediate lower-dose and higher-dose results confirm that CAVM gradually increased the dose for the synthesized image.

Table 1 quantitatively compared CAVM with the competing methods. The results show that CAVM outperforms the other models in both tumor regions and healthy tissue. In the tumor region, CAVM surpasses the other models with a relative improvement of at least 9.61% in PSNR, 5.40% in SSIM, and 5.39% in the Dice coefficient. In healthy tissue, CAVM also has better performance than the competing methods, although the margin is smaller as synthesis of healthy tissue is less challenging. Note that the upper-bound mean Dice coefficient of tumor segmentation achieved on real images is 90.19 (%), and the Dice coefficient achieved with the synthesized image of CAVM is 88.74, which is close to the upper bound.

3.3 Ablation Study

We performed an ablation study to confirm the contribution of each major component of CAVM by individually removing them. We first investigated the ben-

Table 1. The SSIM (%), PSNR (dB), and Dice coefficient (%) results (mean $\pm$ std) of each method on the test set. Boldface marks the top result.

Method	Brain Tumor			Healthy Tissue	
	SSIM	PSNR	Dice	SSIM	PSNR
ResViT	88.01 $\pm$ 0.53	18.16 $\pm$ 2.97	74.25 $\pm$ 9.83	93.97 $\pm$ 0.20	27.24 $\pm$ 2.00
SynDiff	88.12 $\pm$ 0.55	18.21 $\pm$ 2.94	74.14 $\pm$ 9.65	93.62 $\pm$ 0.19	26.91 $\pm$ 1.96
TSF-Seq2Seq	86.03 $\pm$ 0.72	19.71 $\pm$ 3.28	82.25 $\pm$ 9.12	93.10 $\pm$ 0.37	27.35 $\pm$ 2.15
MT-Net	89.64 $\pm$ 0.66	23.82 $\pm$ 3.42	84.20 $\pm$ 9.89	95.04 $\pm$ 0.17	32.38 $\pm$ 1.80
CAVM	94.48$\pm$0.35	26.11$\pm$2.87	88.74$\pm$2.86	95.50$\pm$0.19	32.91$\pm$1.94

Table 2. The SSIM (%) and PSNR (dB) results (mean $\pm$ std) of the ablation study. Boldface marks the top result.

Method	Brain Tumor	
	SSIM	PSNR
One Step	89.98 $\pm$ 0.77	24.22 $\pm$ 3.69
Two Steps	92.53 $\pm$ 0.65	24.92 $\pm$ 3.56
w/o Dose-variant Autoregression	89.70 $\pm$ 0.70	23.97 $\pm$ 3.91
CAVM	94.48$\pm$0.35	26.11$\pm$2.87

efit of each dose increase step in the autoregression, where one and two steps were considered. With one step, CAVM synthesized the T1Gd image without considering lower-dose and higher-dose images, whereas with two steps, the synthesis only considered the lower-dose image. Also, we have completely removed the dose-variant autoregression module, where the T1Gd image was directly obtained by decoding the outputs of the tokenizer. Note that the removal of the decomposition that directly increases the dose based on all image tokens is computationally infeasible, and thus it was not considered. The SSIM and PSNR results in the tumor regions are summarized for the ablation study in Table 2. Fewer steps led to worse image qualities, and without the autoregression the synthesis performance also decreased. These results show that the dose increase steps and the complete autoregression module are all beneficial to the synthesis.

4 Conclusion

In this work, we present CAVM, a conditional autoregressive vision model for contrast-enhanced brain tumor MRI synthesis from non-contrast images. CAVM gradually increases the dosage of the contrast agent and produces higher-dose images from previous lower-dose ones until the final desired dose is achieved. The dose increase is achieved with autoregression based on LLaMA-style Transformer blocks and dose-variant and dose-invariant token decomposition. Experimental results on public data show that CAVM compares favorably with state-of-the-art synthesis models, particularly in tumor regions, and it also improves downstream brain tumor segmentation. Future work can explore further improvement

of CAVM, where the hidden state in the dose-variant autoregression can be quantized for more stable dose update.

Acknowledgments. C. Ye is supported by the Beijing Municipal Natural Science Foundation (7242273) and the Fundamental Research Funds for the Central Universities (2024CX06040). T. Yan is co-funded by Key-Area Research and Development Program of Guangdong Province (2023B0303030002), the STI 2030-Major Projects (grant number 2022ZD0208500), and the National Natural Science Foundation of China (grant numbers U20A20191, 62336002).

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

References

1. Baid, U., Ghodasara, S., Mohan, S., Bilello, M., Calabrese, E., Colak, E., Farahani, K., Kalpathy-Cramer, J., Kitamura, F.C., Pati, S., et al.: The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. *arXiv preprint arXiv:2107.02314* (2021)
2. Dalmaz, O., Yurt, M., Çukur, T.: ResViT: Residual vision transformers for multimodal medical image synthesis. *IEEE Transactions on Medical Imaging* **41**(10), 2598–2614 (2022)
3. Dayarathna, S., Islam, K.T., Uribe, S., Yang, G., Hayat, M., Chen, Z.: Deep learning based synthesis of MRI, CT and PET: Review and analysis. *Medical Image Analysis* p. 103046 (2023)
4. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
5. El-Nouby, A., Klein, M., Zhai, S., Bautista, M.A., Toshev, A., Shankar, V., Susskind, J.M., Joulin, A.: Scalable Pre-training of Large Autoregressive Image Models. *arXiv preprint arXiv:2401.08541* (2024)
6. Han, L., Zhang, T., Huang, Y., Dou, H., Wang, X., Gao, Y., Lu, C., Tan, T., Mann, R.: An Explainable Deep Framework: Towards Task-Specific Fusion for Multi-to-One MRI Synthesis. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 45–55. Springer (2023)
7. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 16000–16009 (2022)
8. He, Y., Nath, V., Yang, D., Tang, Y., Myronenko, A., Xu, D.: Swinunetr-v2: Stronger swin transformers with stagewise convolutions for 3d medical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 416–426. Springer (2023)
9. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
10. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods* **18**(2), 203–211 (2021)

11. Jiang, L., Mao, Y., Wang, X., Chen, X., Li, C.: CoLa-Diff: Conditional latent diffusion model for multi-modal MRI synthesis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 398–408. Springer (2023)
12. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
13. Li, C., Wong, C., Zhang, S., Usuyama, N., Liu, H., Yang, J., Naumann, T., Poon, H., Gao, J.: Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems* **36** (2024)
14. Li, H.B., Conte, G.M., Anwar, S.M., Kofler, F., Ezhov, I., van Leemput, K., Piraud, M., Diaz, M., Cole, B., Calabrese, E., et al.: The Brain Tumor Segmentation (BraTS) Challenge 2023: Brain MR Image Synthesis for Tumor Segmentation (BraSyn). ArXiv (2023)
15. Li, Y., Zhou, T., He, K., Zhou, Y., Shen, D.: Multi-scale Transformer Network with Edge-aware Pre-training for Cross-Modality MR Image Synthesis. *IEEE Transactions on Medical Imaging* (2023)
16. Liu, J., Pasumarthi, S., Duffy, B., Gong, E., Datta, K., Zaharchuk, G.: One model to synthesize them all: Multi-contrast multi-scale transformer for missing data imputation. *IEEE Transactions on Medical Imaging* (2023)
17. Özbey, M., Dalmaz, O., Dar, S.U., Bedel, H.A., Öztürk, Ş., Güngör, A., Çukur, T.: Unsupervised medical image translation with adversarial diffusion models. *IEEE Transactions on Medical Imaging* (2023)
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. arXiv preprint arXiv:2306.14678 (2023)
19. Schieda, N., Blachman, J.I., Costa, A.F., Glikstein, R., Hurrell, C., James, M., Jabejdar Maralani, P., Shabana, W., Tang, A., Tsampalieros, A., et al.: Gadolinium-based contrast agents in kidney disease: a comprehensive review and clinical practice guideline issued by the Canadian Association of Radiologists. *Canadian journal of kidney health and disease* **5**, 2054358118778573 (2018)
20. Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., Liu, Y.: Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing* **568**, 127063 (2024)
21. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al.: Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971 (2023)
22. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288 (2023)
23. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
24. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D., et al.: Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems* **35**, 24824–24837 (2022)
25. Zhou, Z., Lu, Z.R.: Gadolinium-based contrast agents for magnetic resonance cancer imaging. *Wiley Interdisciplinary Reviews: Nanomedicine and Nanobiotechnology* **5**(1), 1–18 (2013)