




Neural Medical Image Recovery and Synthesis


S Kevin Zhou

School of Biomedical Engineering, USTC

Fellow of AIMBE | IAMBE | IEEE | MICCAI | NAI

2025.11.27

Image generation using AI



$$P(I | C)$$

Natural Image

Map, text, etc.



《 Théâtre D'opéra Spatia 》

山海奇境

山海奇境





饕餮，兽名，身如牛，人面，
目在腋下，食人。

——《山海经·西荒经》

號食殄食

文鳐鱼

西流注于流沙。是多文鳐鱼，
状如鲤鱼，鱼身而鸟翼，
苍文而白首赤喙，常行西海，
游于东海，以夜飞。

——《山海经·西山经》



毕方

有鸟焉，其状如鹤，一足，
赤文青质而白喙，名曰毕方，
其鸣自叫也，见则其邑有讹火。
——《山海经·西山经》





魏大海 飞虎图



李娜 飞虎图



刘国栋 闹天宫



王. 闹天宫




张. 闹天宫



马. 闹天宫














Medical image generation using AI




Medical Image



$$P(I | C)$$

Diagram illustrating the probability of generating a Medical Image (I) given a condition (C). The expression is $P(I | C)$. An arrow points from the text "Medical Image" to the variable I . Another arrow points from a question mark (?) to the variable C .

Medical imaging pipeline




CT imaging and reconstruction





Filtered Back Projection (FBP)

MRI imaging and reconstruction




k-Space




MRI image

AI-generated medical image (AIGMI): 2 forms




① Image Recovery



② Image Synthesis

① Neural medical image recovery



Medical image recovery: an inverse problem

□ Forward problem :
from object to observations

$$Y = A X$$

□ Inverse problem :
from observations to object

$$\arg \min_X \text{Loss}(Y, AX) + \text{Reg}(X)$$

Medical image recovery: a learning perspective

Recovery task:

- From an acquired signal X , estimate a 'better' image $Y = f(X)$

Assumption:

- There exists a set of paired training examples $\{(X_i, Y_i)\}$ or unpaired sets $\{X_i\} + \{Y_i\}$

(Supervised) Learning:

$$\arg \min_W \sum_i \text{Loss}(Y_i, f(X_i; W)) + \text{Reg}(W)$$


Pros:

- Data-driven or data as regularizer
- Fast estimation
- Leverage the ML strength

Cons:

- Need for training data
- Data bias
- Knowledge-agnostic

Neural medical image recovery = Deep learning + imaging knowledge




(Supervised) Learning:


$$\arg \min_W \sum_i \text{Loss}(Y_i, f(X_i; W)) + \text{Reg}(W)$$

Key
Challenges


Which knowledge? How to integrate?
Better network architecture
Less supervised




Neural medical image recovery



Neural medical image recovery




Dual-domain network (DuDoNet)




DuDoNet: Dual-domain network for metal artifact reduction


Radon Inverse Layer




CT metal artifact reduction





CBCT artifact reduction




DuDoRNet: Dual-domain recurrent network for MR recon



PSNR: **2-7dB** better than SOTA



DSA video denoising



Xray rib suppression




Li et al. High-resolution chest X-ray bone suppression using unpaired CT structural priors TMI 2020.


Han et al. GAN-based disentanglement learning for chest X-ray rib suppression. Medical Image Analysis 2022.



Xray spine enhancement




Xray pelvis extraction





CT metal artifact reduction using NeRF

$$\begin{aligned}
 p(\mathbf{r}) &= -\ln \int_{\mathcal{E}} \eta(E) \cdot \exp \left(- \int_{\mathbf{r}} \mathcal{M}(\mathbf{x}) \cdot \mu_E^*(\mathbf{x}) + [1 - \mathcal{M}(\mathbf{x})] \cdot \mu(\mathbf{x}) d\mathbf{x} \right) dE \\
 &= \underbrace{-\ln \int_{\mathcal{E}} \eta(E) \cdot \exp \left(- \int_{\mathbf{r}} \mathcal{M}(\mathbf{x}) \cdot \mu_E^*(\mathbf{x}) d\mathbf{x} \right) dE}_{\text{Nonlinear Metal Effect}} + \underbrace{\int_{\mathbf{r}} [1 - \mathcal{M}(\mathbf{x})] \cdot \mu(\mathbf{x}) d\mathbf{x}}_{\text{Linear Integral}}.
 \end{aligned}$$



3D Gaussian Representation






Neural Medical Image Recovery publications

1. Lin et al. **DuDoNet**: Dual domain network for CT metal artifact reduction. CVPR 2019.
2. Zhou et al. **DuDoDR-Net**: Dual-domain data consistent recurrent network for simultaneous sparse view and MAR in computed tomography. MedIA 2020.
3. Zhou & Zhou. **DuDuRNet**: Learning a dual-domain recurrent network for fast MRI reconstruction with deep T1 prior. CVPR 2020.
4. Lyu et al. **U-DuDoNet**: Unpaired dual domain network for CT metal artifact reduction. MICCAI 2021.
5. Zhou et al. **DuDoUFNet**: Dual-domain under-to-fully-complete progressive restoration network for simultaneous MAR and low-dose CT reconstruction. IEEE TMI 2022.
6. Wang et al. **DuDoTrans**: Dual-domain transformer for sparse view CT reconstruction. MLMIR 2022.
7. H Liao et al. Adversarial sparse-view CBCT artifact reduction. MICCAI 2018. (oral)
8. Liao et al. Artifact disentanglement network for unsupervised metal artifact reduction. MICCAI 2019. (oral)
9. Liao et al. ADN: Artifact disentanglement network for unsupervised metal artifact reduction. IEEE TMI 2020.
10. Lyu et al. Encoding mask projection for metal artifact reduction in CT. MICCAI 2020.
11. Peng et al. Towards multi-sequence MR image recovery from undersampled k-space data, MIDL 2020. (oral)
12. Peng et al. SAINT: Spatially aware interpolation network for medical slice synthesis. CVPR 2020.
13. Wang et al. Generalizable limited-angle CT reconstruction via sinogram extrapolation. MICCAI 2021.
14. Peng et al. DA-VSR: Domain adaptable volumetric super-resolution for medical images. MICCAI 2021.
15. Liu et al. Universal undersampled MRI reconstruction. MICCAI 2021. (oral)
16. Lyu et al. A3DSegNet: Anatomy-aware artifact disentanglement and segmentation network for unpaired segmentation, artifact reduction, and modality translation. IPMI 2021.
17. Zhou et al. Limited view tomographic reconstruction using a cascaded residual dense spatial-channel attention network with projection data fidelity layer. IEEE TMI 2021.
18. Liu et al. Undersampled MRI reconstruction with side information guided normalization. MICCAI 2022.
19. Peng et al. Towards performant and reliable undersampled MR reconstruction via diffusion model sampling. MICCAI 2022. (oral)
20. Liu et al. Stabilize, decompose, and denoise: Self supervised fluoroscopy denoising. MICCAI 2022.
21. Q Wu et al. Unsupervised polychromatic neural representation for CT metal artifact reduction. NeurIPS 2023.
22. Li et al. 3DGR-CT: Sparse-view CT reconstruction with a 3D Gaussian representation. MedIA 2025.
23. Z Xu et al. NeRF-based CBCT reconstruction needs normalization and initialization. MICCAI 2025.

② Neural medical image synthesis




$$P(I_{b,d} | S_{a,c})$$



② Neural medical image synthesis

2.1)

Causal Longitudinal
Image Synthesis (CLIS)



③ Medical Image Generation FM Causal Longitudinal Image Synthesis (CLIS)



$A\beta$ & τ protein accumulation




Dying neurons and loss of neural connection




Degenerate memory and cognition

Alzheimer's Disease (AD)
Causal process


Measurements: Multiple **factors** and **modalities**



Causality between measurements



Causal Image Synthesis



**A causal model for
brain MRI synthesis**


Q1 : How to find the causal edges?

Integrating multiple methods

Self-Causality assumption			
1. $x_5^t \rightarrow x_5^{t+1}$	2. $x_6^t \rightarrow x_6^{t+1}$	3. $x_7^t \rightarrow x_7^{t+1}$	4. $x_8^t \rightarrow x_8^{t+1}$
5. $x_9^t \rightarrow x_9^{t+1}$	6. $x_{10}^t \rightarrow x_{10}^{t+1}$		
Recovered by all methods			
7. $x_2^t \rightarrow x_3^t$	8. $x_1^t \rightarrow x_5^t$	9. $x_4^t \rightarrow x_5^t$	10. $x_1^t \rightarrow x_6^t$
11. $x_4^t \rightarrow x_6^t$	12. $x_1^t \rightarrow x_7^t$	13. $x_4^t \rightarrow x_7^t$	14. $x_2^t \rightarrow x_8^t$
15. $x_1^t \rightarrow x_9^t$	16. $x_{10}^t \rightarrow x_9^t$	17. $x_8^t \rightarrow x_{10}^t$	18. $x_5^t \rightarrow x_{10}^{t+1}$
19. $x_6^t \rightarrow x_{10}^{t+1}$	20. $x_7^t \rightarrow x_{10}^{t+1}$		
Recovered by two methods			
21. $x_1^t \rightarrow x_8^t$	(FCI, GES)	22. $x_5^t \rightarrow x_8^{t+1}$	(FCI, GES)
23. $x_6^t \rightarrow x_8^{t+1}$	(GES, LiNGAM)	24. $x_7^t \rightarrow x_8^{t+1}$	(GES, LiNGAM)
25. $x_8^t \rightarrow x_9^t$	(FCI, GES)		



Q2: How to describe SCM?

Use MLP





* SCM – structural causal model







Q3: How to mitigate dimensional gap? GenAI




3D StyleGAN



Acquired?
Synthesized?






Causal intervention

Brain Image
Time Machine

brain: 1283ml
ventricle: 49ml
GM: 464ml


s = female
a = 81
score = 22

AD subject




Acquired image

control age



difference map




age=70 years




age=90 years

actual: 81 years


Factor Intervention



control ventricle_volume



difference map



30 ml




70 ml

factual: 49 ml


Comparison with ground truth

SSIM	CIS	Mi-GAN[1]	cGAN[2]
1-yr (N=278)	0.973	0.946	0.944
4-yr (N=45)	0.963	0.942	0.939
All (N=490, avg 1.95yr)	0.971		




Cognition & Brain Future Condition

prediction: future ADAS13 and diagnosis



* unpublished data

Causal PETS




	PSNR (dB)	SUVr MAE (e-1)	pMCI / AD classification accuracy (%)
pGAN	23.19	0.23	-
Pix2Pix	21.83	0.21	-
BPGAN	24.09	0.18	-
Causal PETS (ours)	26.74	0.10	
T1 + Tabular			76.02
T1 + Tabular + Causal PETS			85.20

Summary



GenAI

+ DisAI



Causal

vs Correlational



Factor+Image

Multi-omics




Apps


- Time machine
- Factor intervention
- Workflow modeling

⊖ Neural Medical image synthesis


T1




T2



T1Gd



FLAIR




2.2) Unified multi-Modal Image Synthesis

Summary




AIGC for Creativity



Recovery

- Dual-domain



Synthesis

- Causal
- Universal

AIGMI for Realness