Generative AI for medical image reconstruction in positron emission tomography (PET)

PET Image Reconstruction with Diffusion Models

Andrew J. Reader

Overview

- Quick review of PET, image reconstruction, ML to MAP
- Justification for deep learning

- Diffusion models / score-based generative models

- **Unsupervised** diffusion models for reconstruction: a likelihood-scheduling method
- **Steerable** diffusion models (*for when our acquired data are out of distribution*)
- **Personalisation** of diffusion models (*for improving quality of training data*)
- **Supervised** diffusion models (*exploiting measured data directly paired with high-quality reference images*)

Positron emission tomography (PET)

Used to image the brain, heart and body (e.g. dementia, heart disease, cancer)

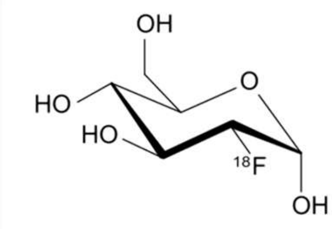

Radiotracer is injected

$^{18}F \Rightarrow {}^{18}O + e^{+}$ e^{-}

511 keV

511 keV

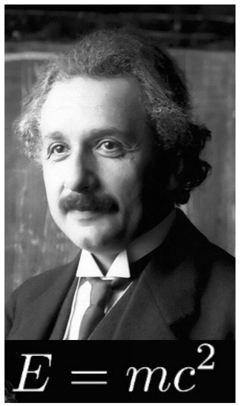

PET scanner uses high-density crystals to detect the pairs of annihilation photons

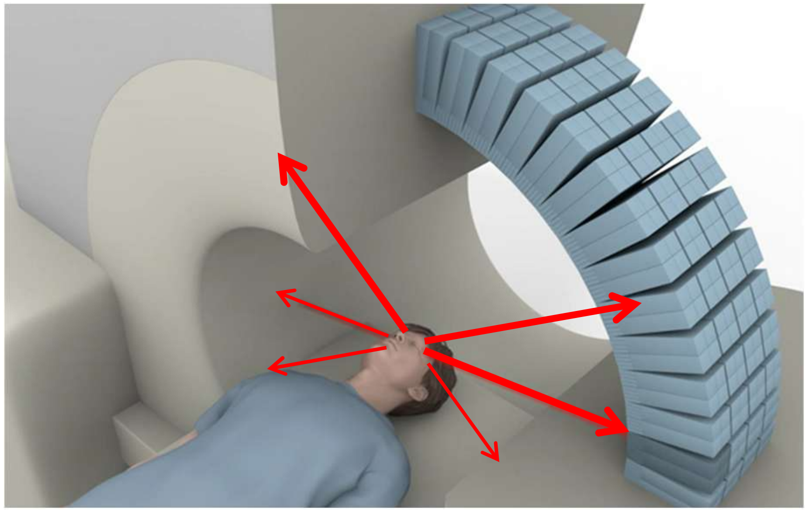

From data acquisition to PET image reconstruction

Radiotracer distribution

$f(\boldsymbol{r})$

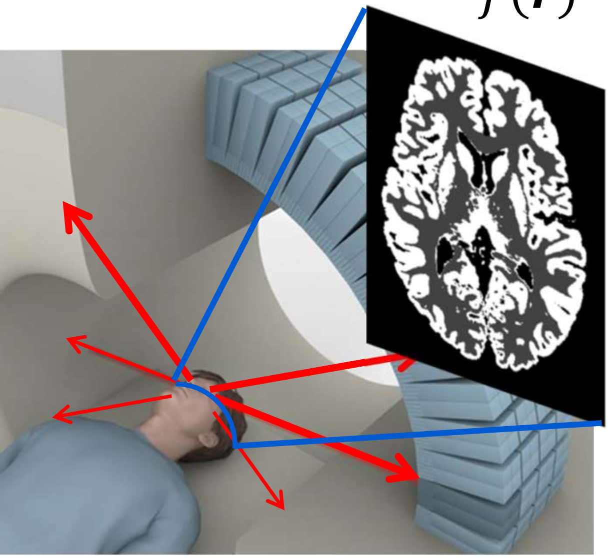

PET scanner

$\boldsymbol{A}$

Measured data

$\boldsymbol{m}$

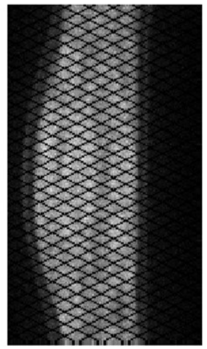

Reconstruction

$\hat{\boldsymbol{x}}$

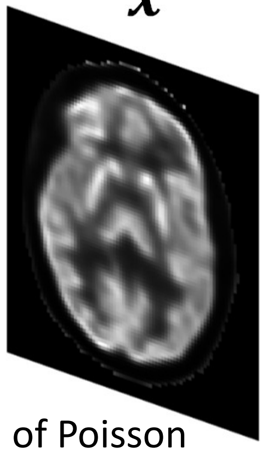

Mean of Poisson process in each voxel

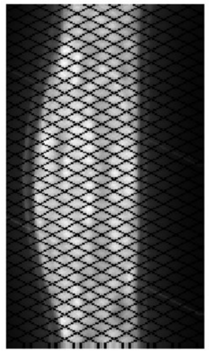

$$\boldsymbol{q} = \boldsymbol{A}\hat{\boldsymbol{x}} + \boldsymbol{b}$$

Maximise: $PLL(\boldsymbol{x}|\boldsymbol{m})$

$$\hat{\boldsymbol{x}} = \arg\max_{\boldsymbol{x}} \sum_{i=1}^{I} (m_i \ln[\boldsymbol{A}\boldsymbol{x}]_i - [\boldsymbol{A}\boldsymbol{x}]_i)$$

Maximum *a posteriori* (MAP) reconstruction

$$\overset{\text{POSTERIOR}}{\mathrm{p}(\boldsymbol{x}|\boldsymbol{m})} \propto \overset{\text{LIKELIHOOD}}{\mathrm{p}(\boldsymbol{m}|\boldsymbol{x})}\overset{\text{PRIOR}}{\mathrm{p}(\boldsymbol{x})} \qquad \mathrm{p}(\boldsymbol{x}) \propto \exp[-\beta U(\boldsymbol{x})]$$

$$\hat{\boldsymbol{x}} = \arg\max_{\boldsymbol{x}} \underbrace{\underbrace{\sum_{i=1}^{I} \left(m_i \ln[\boldsymbol{A}\boldsymbol{x}]_i - [\boldsymbol{A}\boldsymbol{x}]_i\right)}_{\text{LOG LIKELIHOOD}} \underbrace{- \beta U(\boldsymbol{x})}_{\text{LOG PRIOR}}}_{\text{LOG POSTERIOR}}$$

Model-based reconstruction: maths, stats, physics, images

1980s – 1990s (filtered backprojection (FBP))

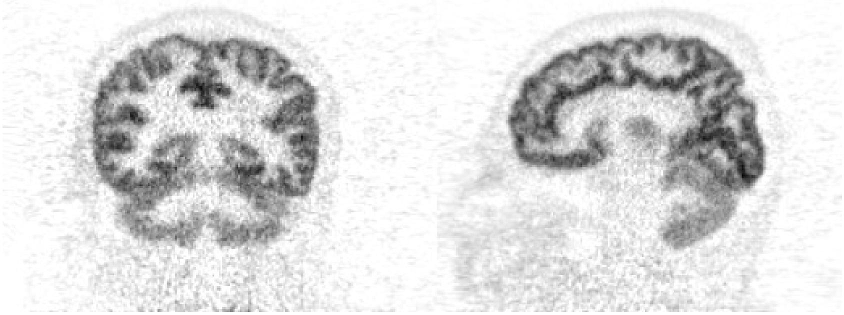

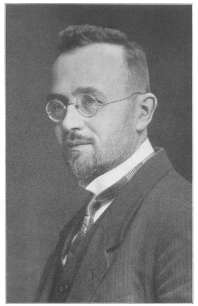

Radon

1990s (iterative reconstruction, OSEM, MLEM)

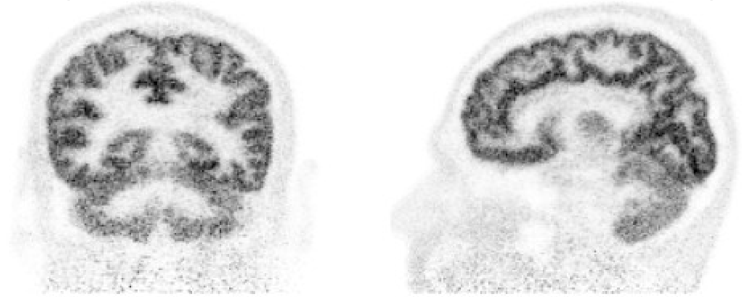

Poisson

2000s (OSEM+PSF, MLEM+PSF)

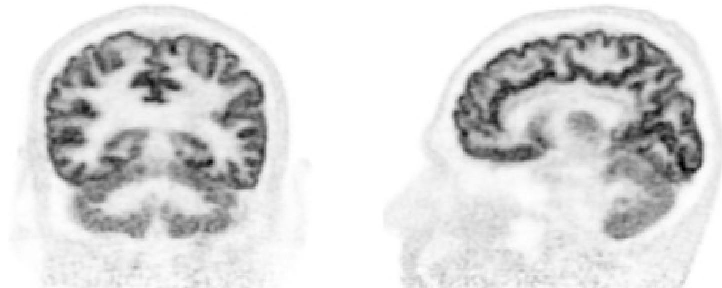

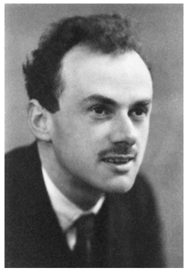

Dirac

Colsher 1980, Kinahan & Rogers 1989, Shepp & Vardi 1982, Hudson & Larkin 1994

5 min

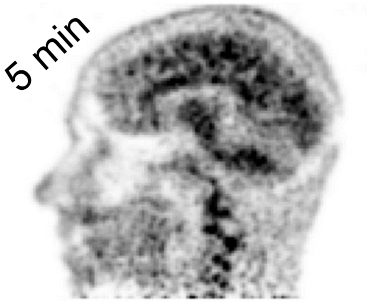

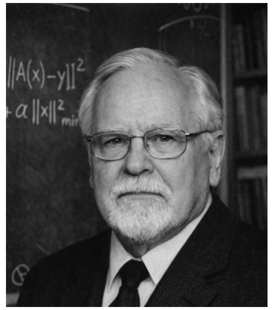

Tikhonov

Regularise (MAPEM)

MRI guidance

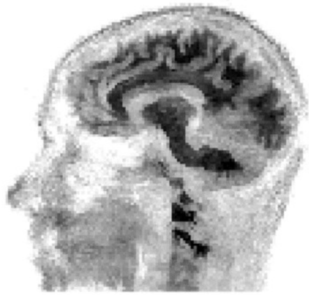

Limitations of ML and MAP Image Reconstruction

Images courtesy
G. Corda–D'Incan

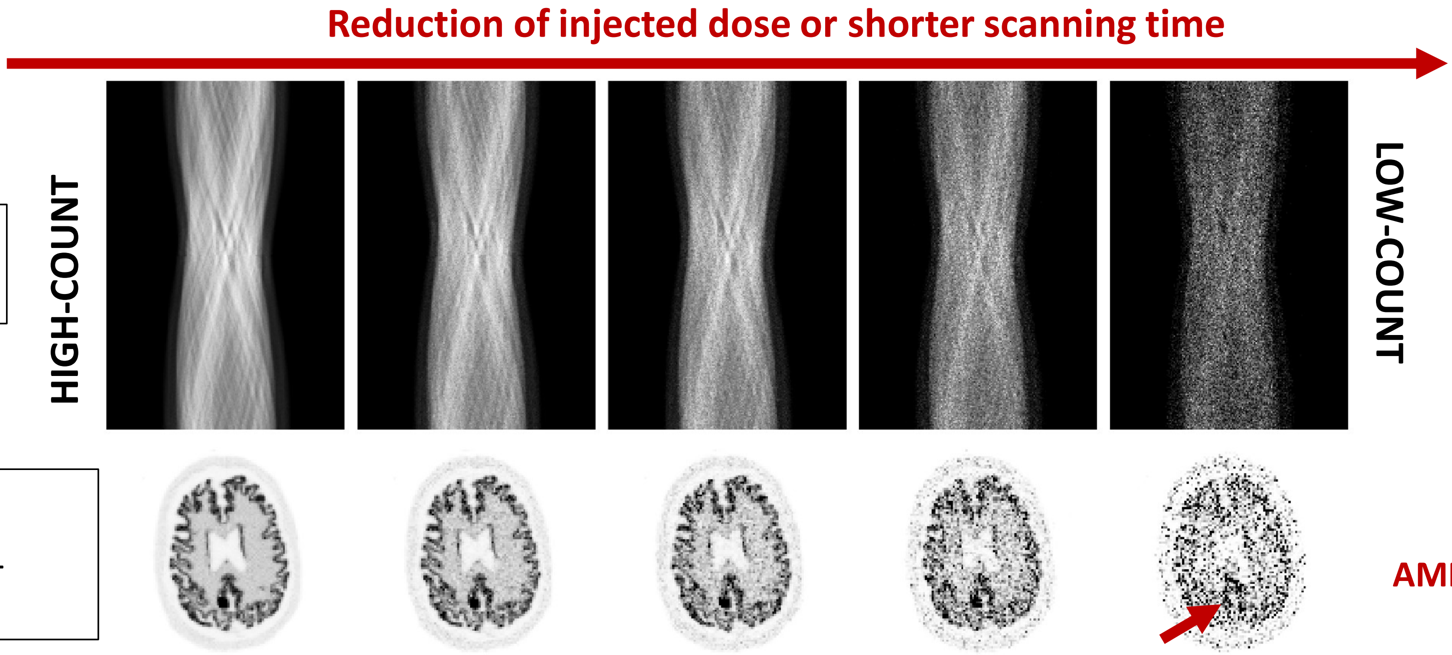

Measured data

OSEM
(4 iterations, 21 subsets)

NOISE AMPLIFICATION

Limitations of ML and MAP Image Reconstruction

Images courtesy
G. Corda–D'Incan

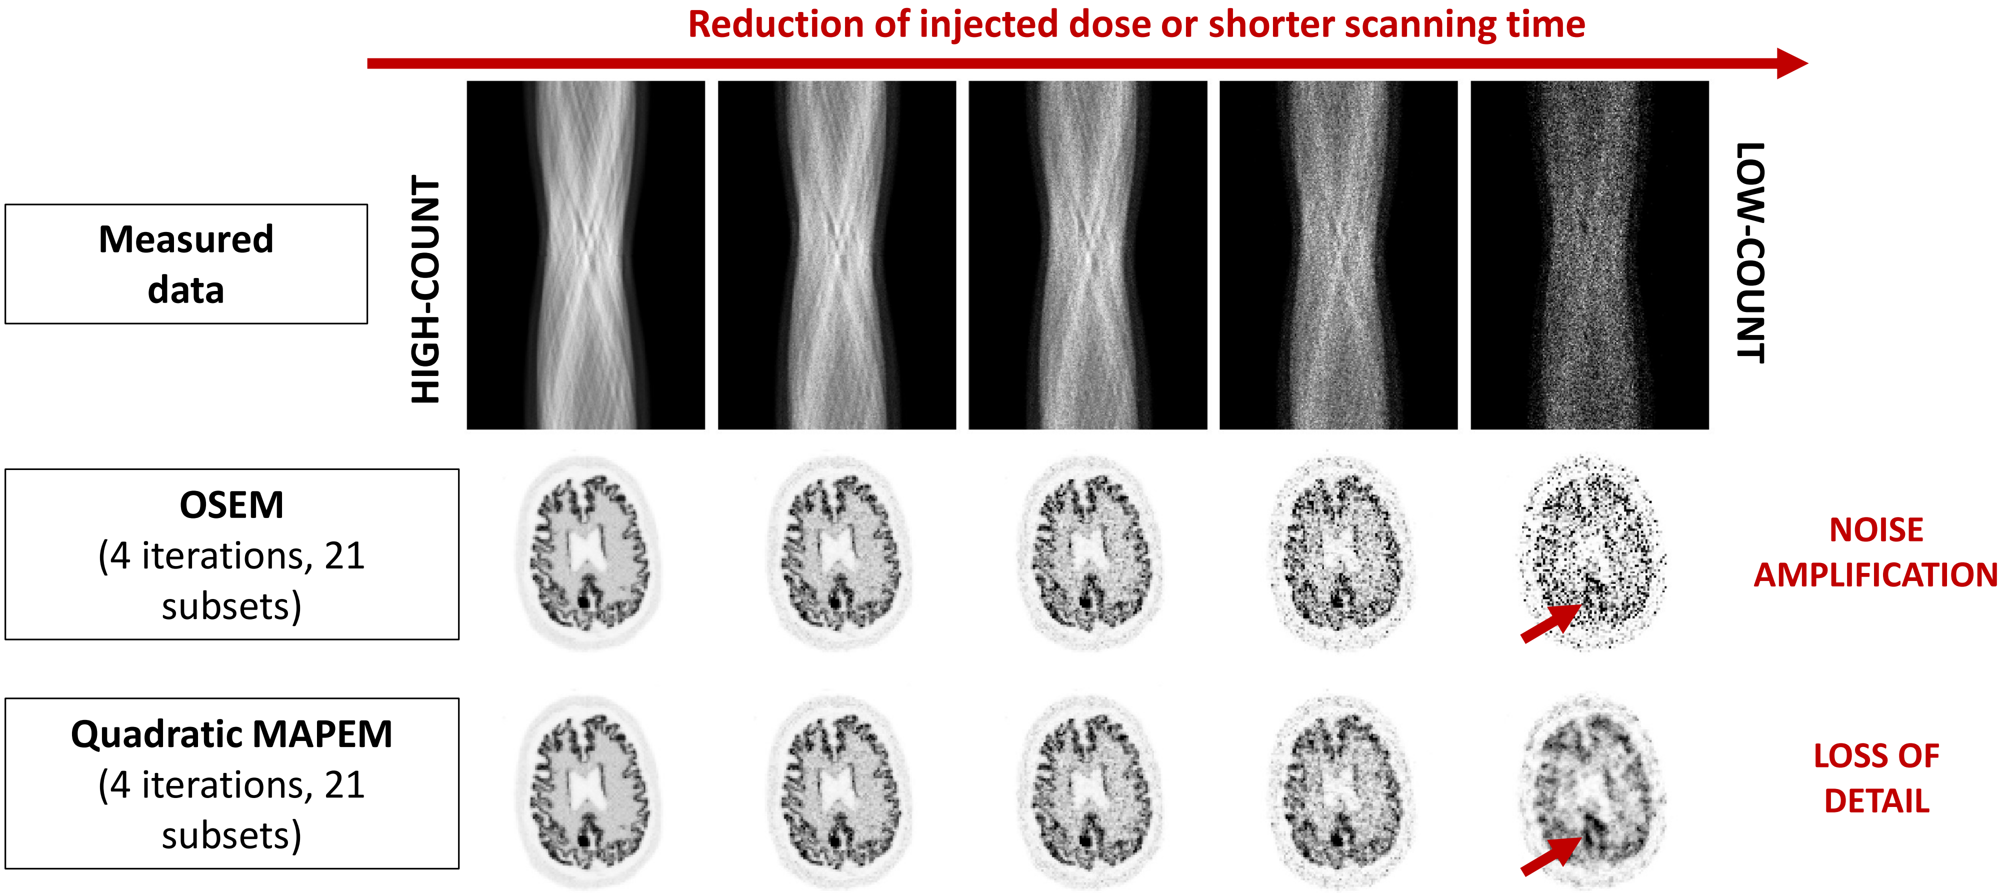

Limitations of ML and MAP Image Reconstruction

Images courtesy
G. Corda–D'Incan

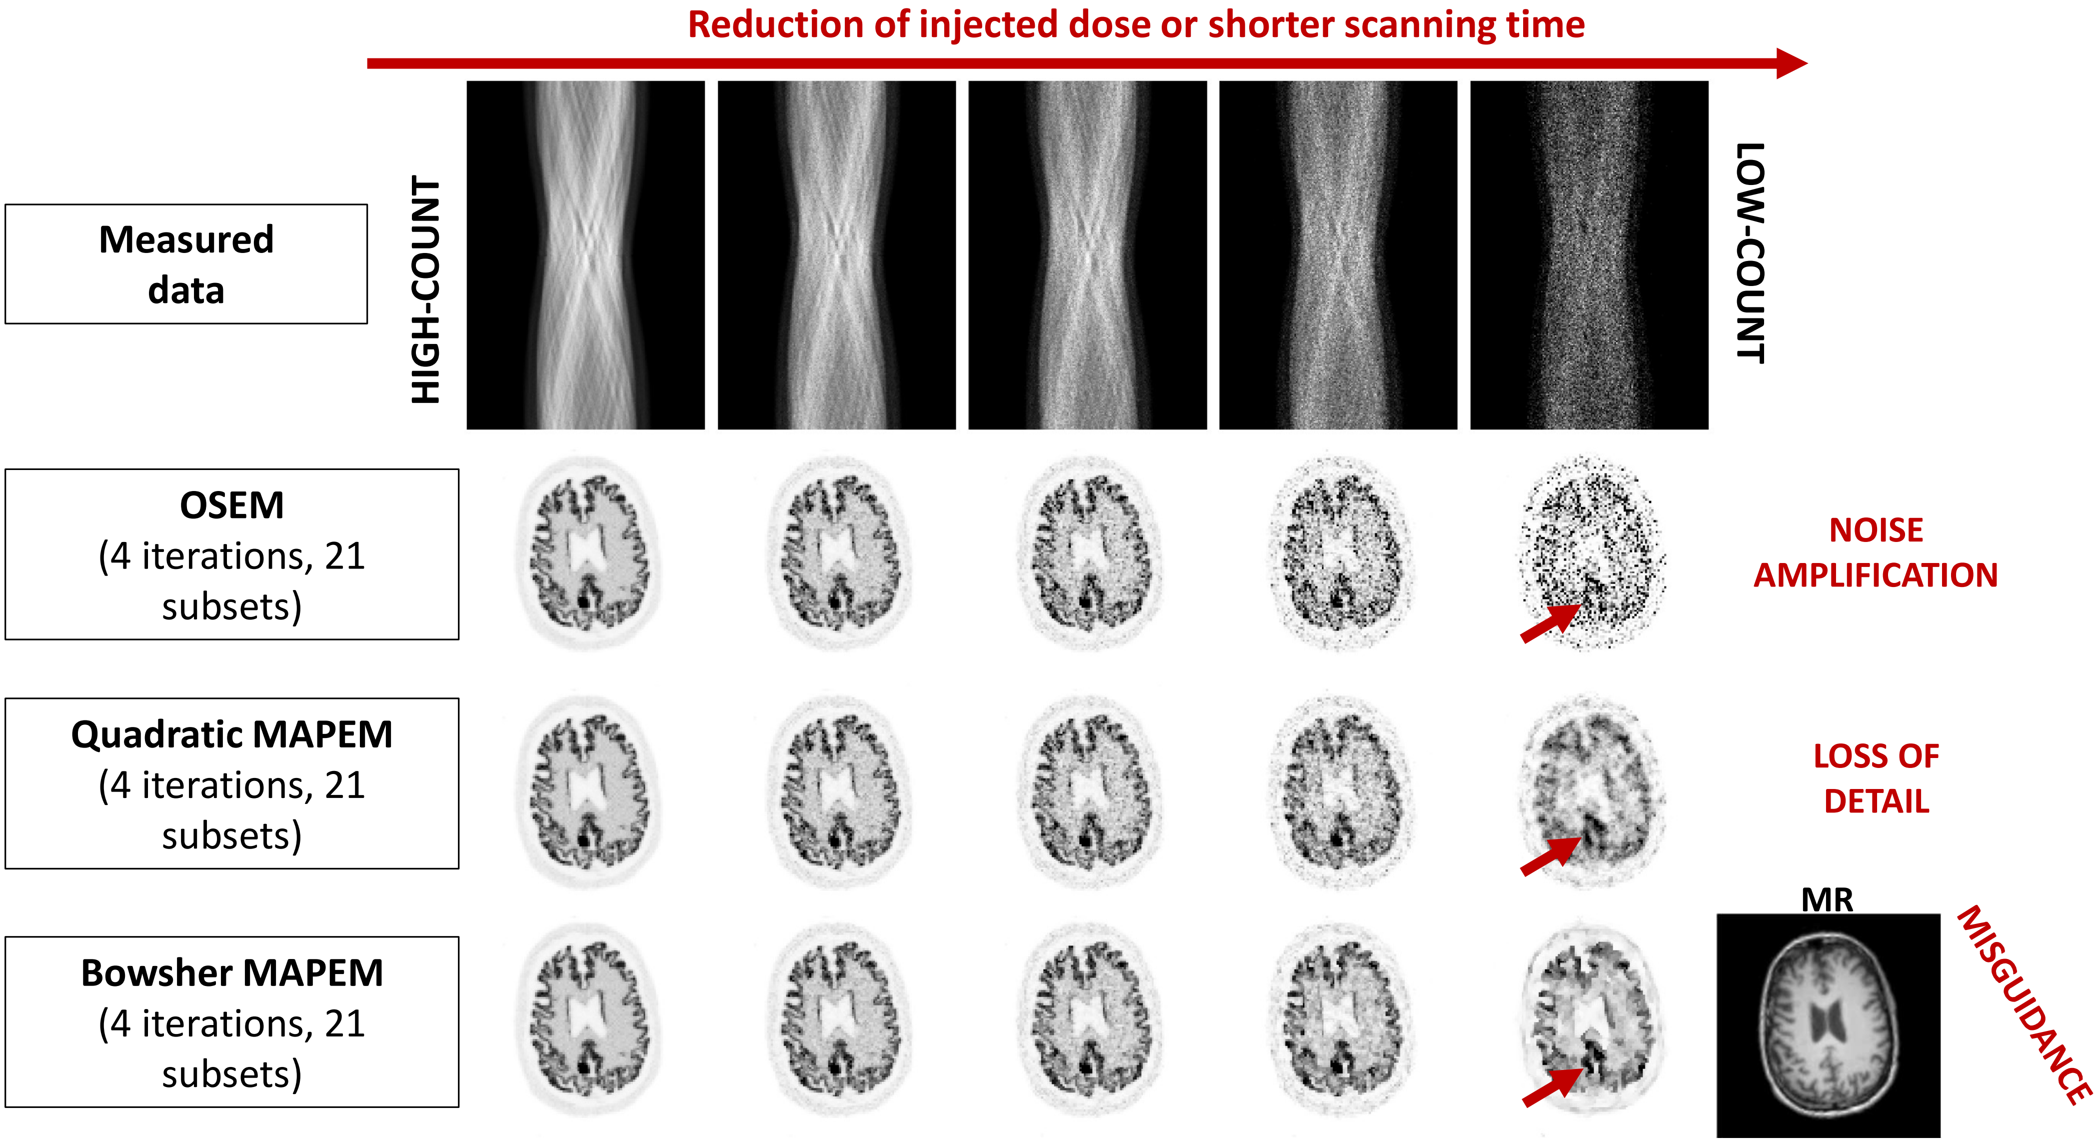

The motivation for deep learning

- Conventional MLEM and MAPEM
  - Noisy, low-resolution data: ML estimate is noisy, with Gibbs ringing
  - MAPEM: noise compensation (*regularisation*) can be too simple (quadratic, TV, RDP)
    … or too imposing (e.g. MRI guidance can be wrong!)
- Assumes
  - Imaging system model
  - Data noise distribution
  - How to regularise (i.e. a simple model of $p(\boldsymbol{x})$ - how images should likely appear)

… but do we really know all these things?

- Deep learning can use
  - Real-data examples to learn
    - more accurate imaging and noise models (and their 'inverse')
  - Ground truth or high-quality reference data
    - to learn the probability distribution of high-quality images $p(\boldsymbol{x})$

Data

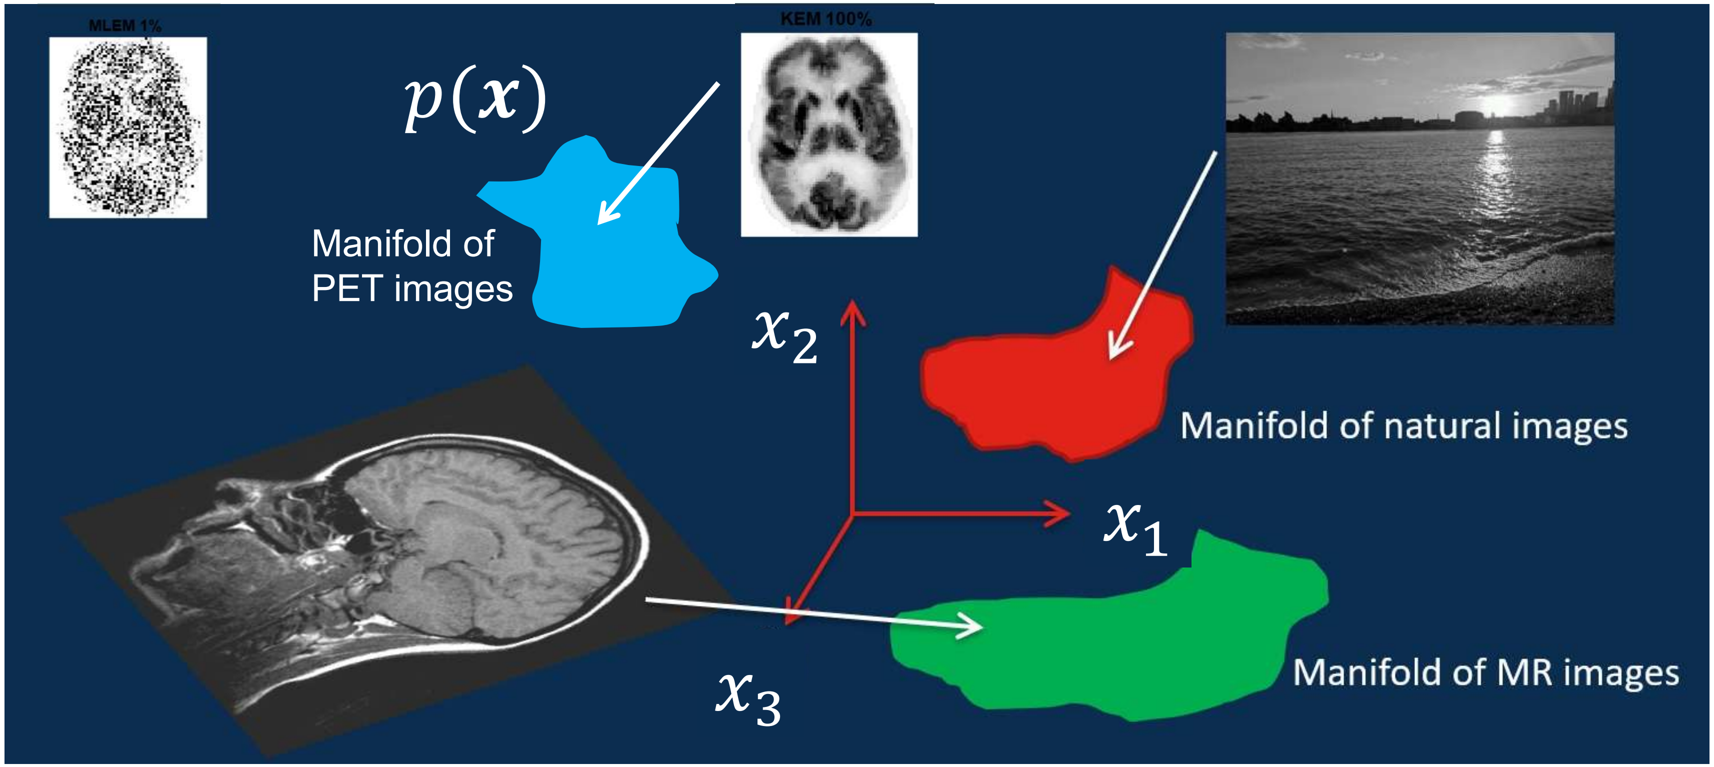

MLEM 1%
KEM 100%
$p(\boldsymbol{x})$
Manifold of
PET images
x_2
x_1
x_3
Manifold of natural images
Manifold of MR images

Modelling $p(x)$ when given very few examples

Method	Training stability	Image realism	Generation time	Mode coverage
Autoencoders (VAE/WAE)	Good	Medium	Fast	Medium
Adversarial (GANs)	Poor	Good	Fast	Poor
Flow-based models	Good	Medium	Medium	Medium
Diffusion models	Good ✓	Very good ✓	Slow	Good ✓

Diffusion models

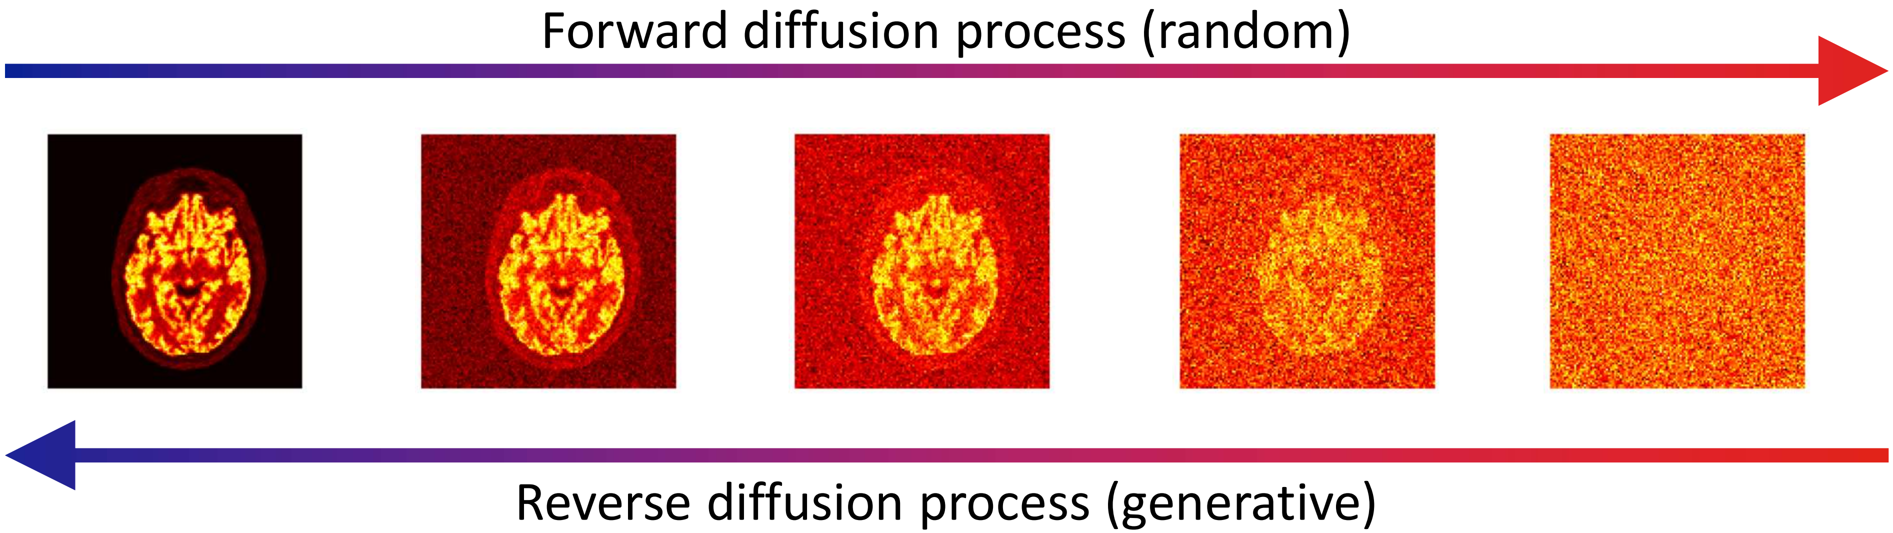

Denoising diffusion probabilistic models (**DDPMs**) - discrete-time step diffusion (Ho et al. 2020, built on the seminal work of Sohl-Dickstein et al. 2015). Reverse process, predicts noise at each step. Fixed number of time steps (100 – 1000).

Score-based generative models (**SGMs**) - continuous-time diffusion via stochastic differential equations, Song et al. 2019

Denoising diffusion implicit models (**DDIMs**) allow steps to be skipped for speed up in the reverse generative process

J. Ho *et al*. *"Denoising diffusion probabilistic models"* NeurIPS 2020

Y. Song and S. Ermon. *"Improved techniques for training score-based generative models"* NeurIPS 2019

J. Song *et al*. *"Denoising diffusion implicit models"* ICLR 2021

DDPM

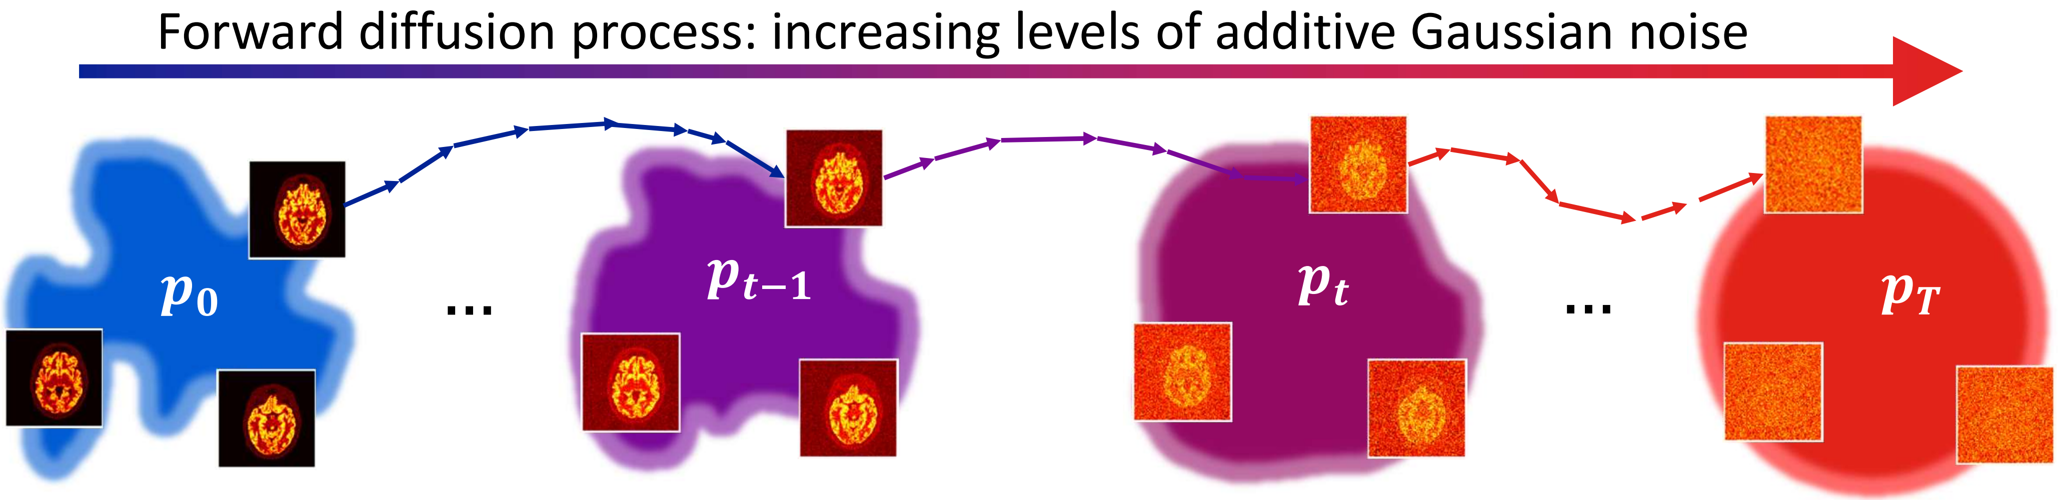

Forward diffusion process: increasing levels of additive Gaussian noise
p_0
...
p_{t-1}
p_t
...
p_T

DDPM

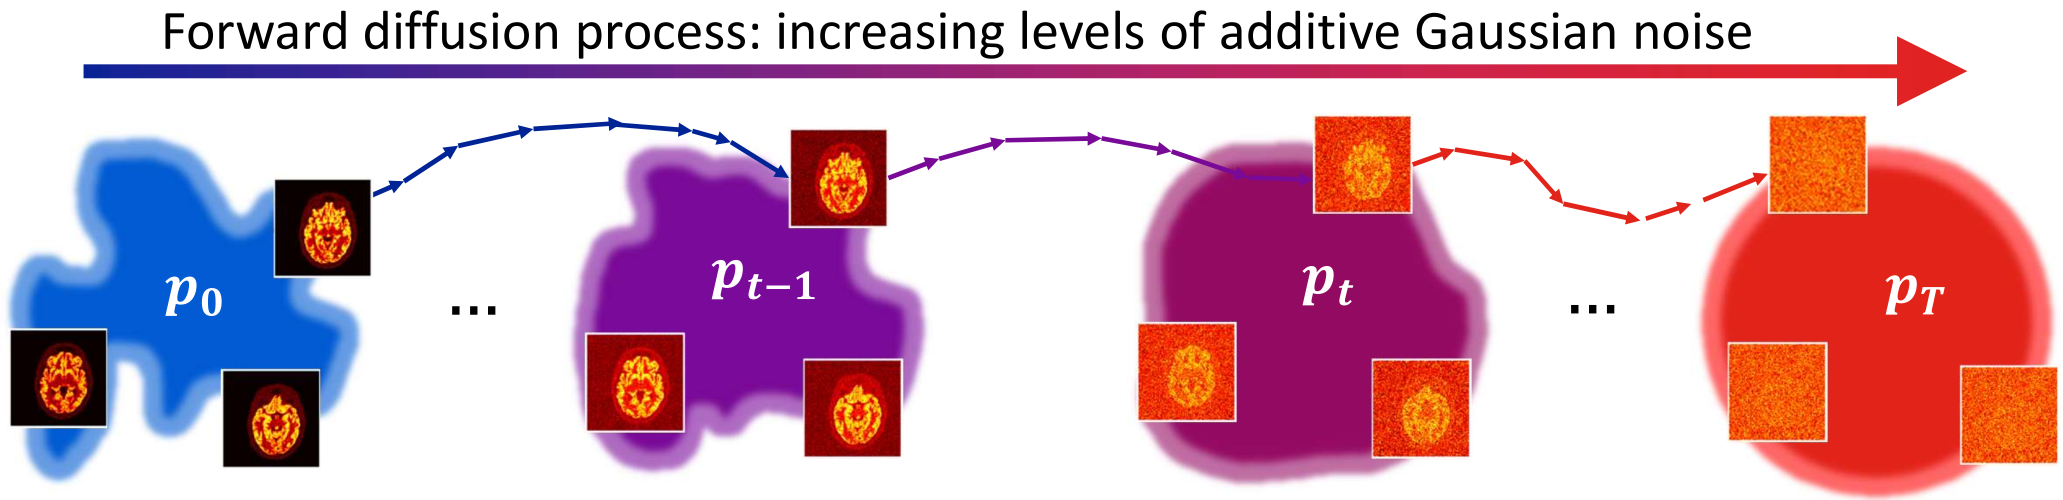

Forward diffusion process: step by step

$$x_t = \sqrt{1-\beta_t}\,x_{t-1} + \sqrt{\beta_t}\,\epsilon \qquad \epsilon \sim \mathcal{N}(0, I)$$

Can generate image at any stage directly by using

$$x_t = \sqrt{\bar{\alpha}_t}\,x_0 + \sqrt{1-\bar{\alpha}_t}\,\epsilon$$

After T time steps

$$x_T = \sqrt{\bar{\alpha}_T}\,x_0 + \sqrt{1-\bar{\alpha}_T}\,\epsilon \qquad x_T \sim \mathcal{N}(0, I)$$

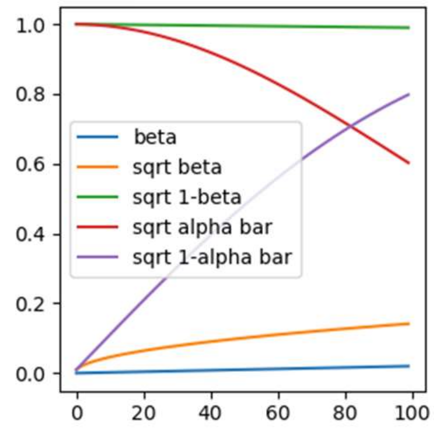

DDPM

Forward diffusion process: increasing levels of additive Gaussian noise

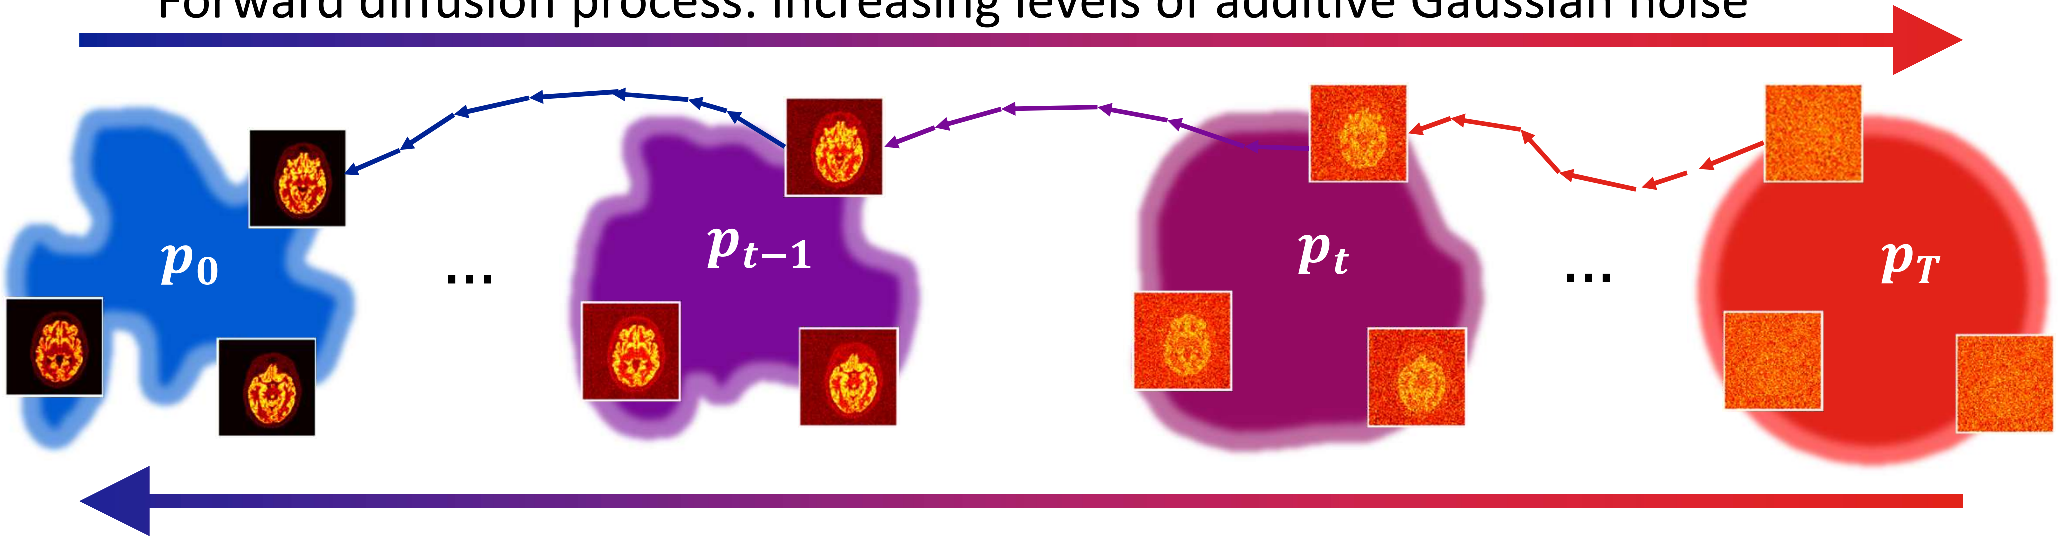

Reverse diffusion (generative) process

Gradually generate each stage, starting from a sample x_T from $p_T(x_T)$

At any given stage, go backwards using

$$x_0 = \frac{1}{\sqrt{\alpha_t}}\left(x_t - \beta_t \frac{1}{\sqrt{1-\bar{\alpha}_t}}\epsilon\right)$$

$$x_{t-1} = x_0 + \sqrt{\beta_t}\epsilon$$

The score function of a probability distribution $p(x)$ is defined by

$$\nabla_{\boldsymbol{x}} \log p(\boldsymbol{x})$$

DDPM

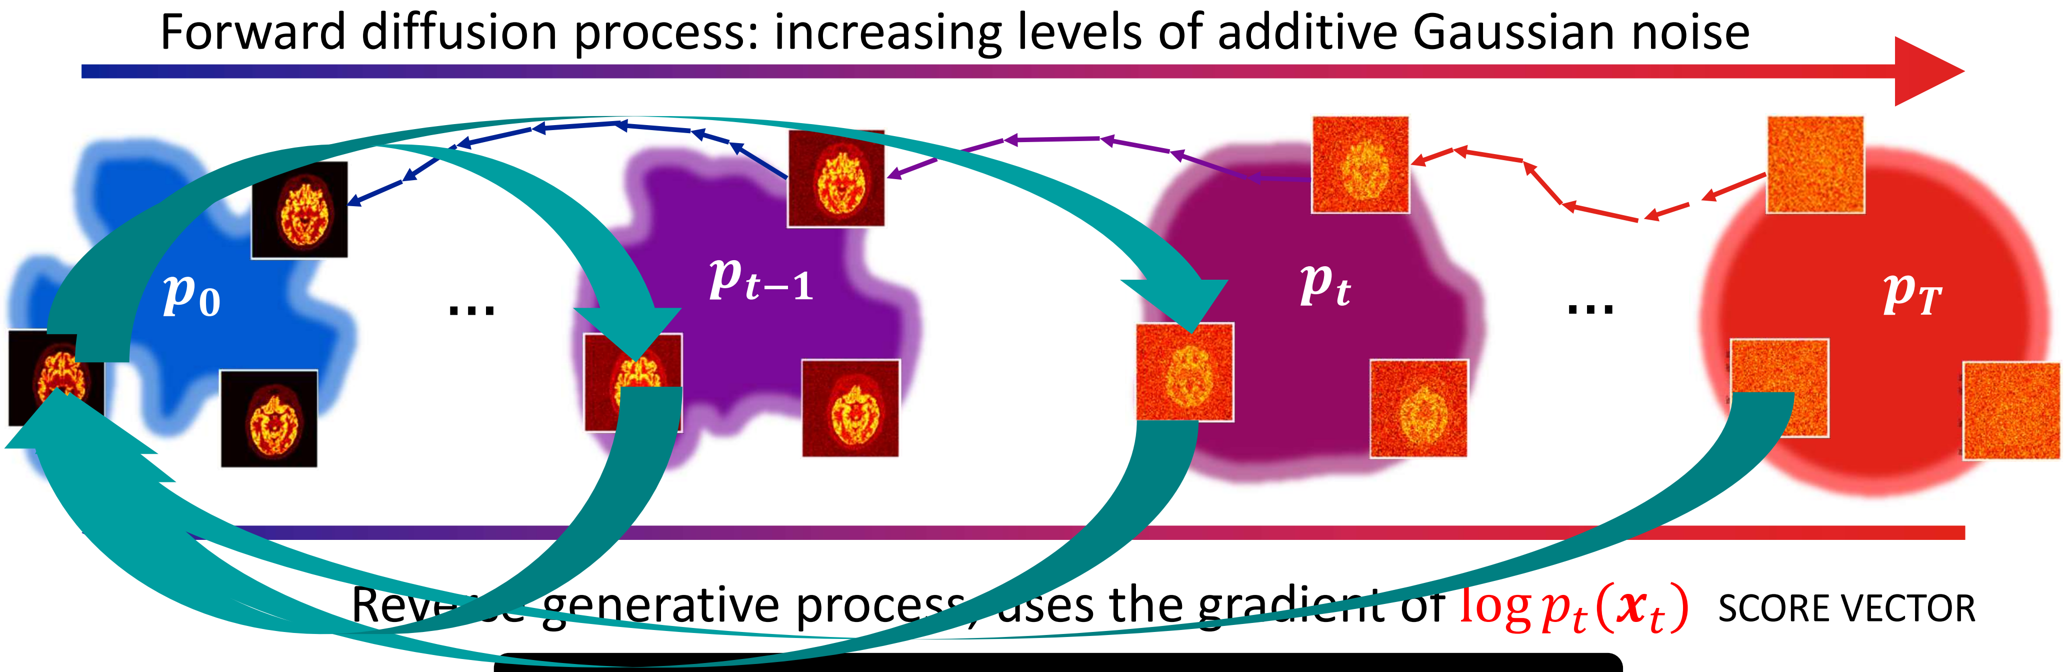

Approximate $\nabla_{\boldsymbol{x}} \log p_t\ (\boldsymbol{x}_t)$ with $s_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t)$

Just need to train a noise-prediction deep network $s_{\boldsymbol{\theta}}$! (i.e. merely a denoiser!).

Pick a random noise level t, and pick an image x_0 from the training set:

make the image noisy at level t, to get x_t : train a network to predict the noise added to the image x_0

Generative reverse process:

start with noise, denoise, then renoise the denoised image, to be slightly less noisy than before $(t-1)$, and enter back into the generative process at that new time step (noise level)

JOURNAL ARTICLE

Diffusion models for medical image reconstruction

George Webber, MMathCompSci, Andrew J Reader, PhD

BJR|Artificial Intelligence, Volume 1, Issue 1, January 2024, ubae013, https://doi.org/10.1093/bjrai/ubae013

Published: 29 August 2024 **Article history** ▾

Review article

Diffusion models for unconditional image generation

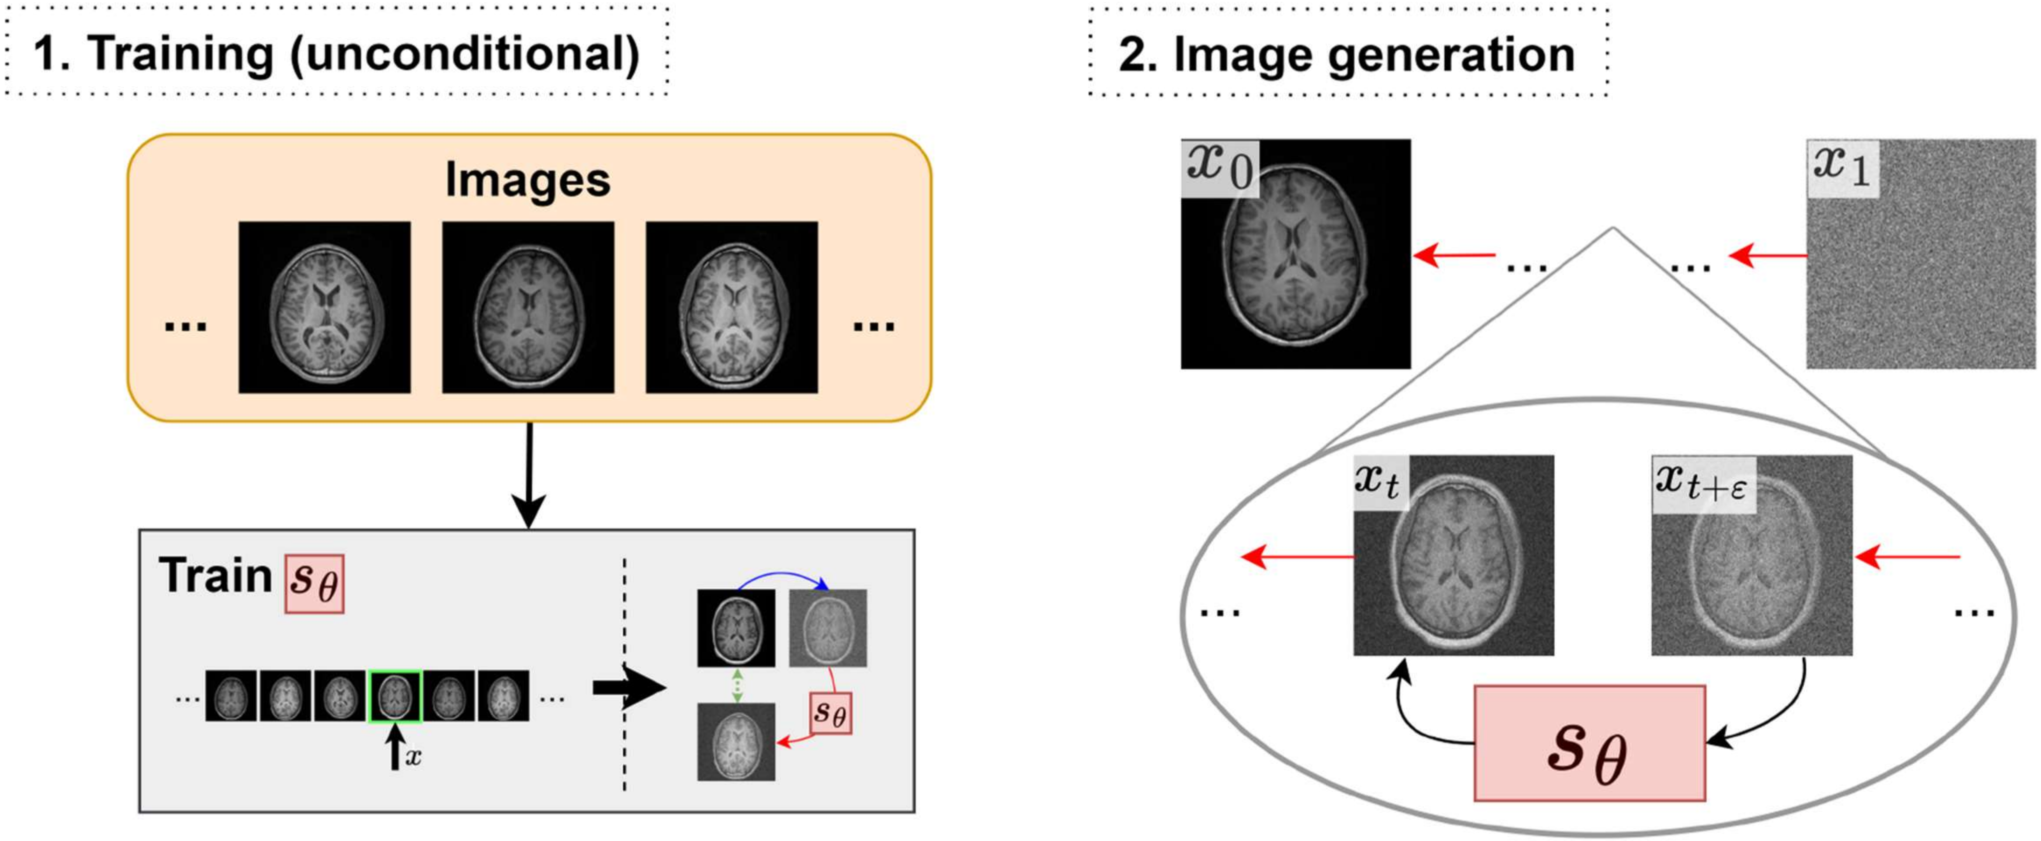

Unsupervised image reconstruction

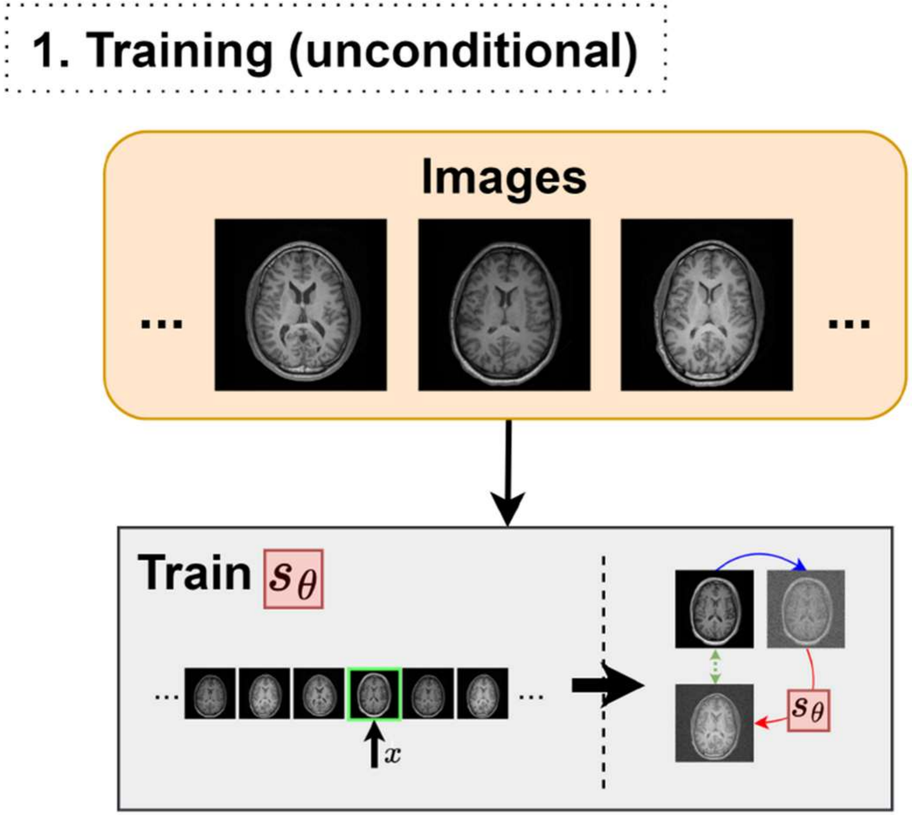

Training is the same!

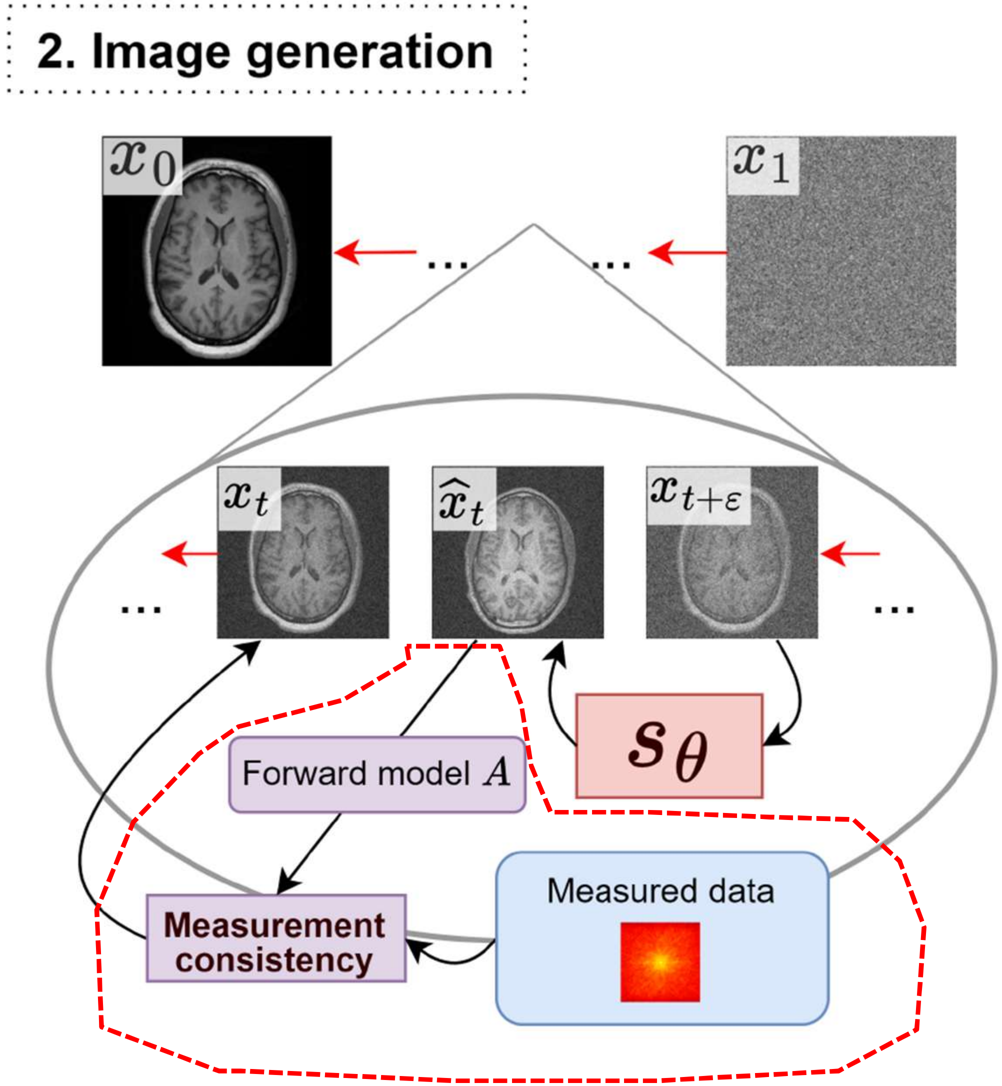

Unsupervised reconstruction

- No need for paired data, just example images to model $p(\boldsymbol{x})$
- Therefore adaptable to different noise levels / scanners / acquisition protocols
- But lower performance than supervised case (high-quality images specifically paired with noisy measured data)

Likelihood-Scheduled Score-Based Generative Modelling for Fully 3D PET Image Reconstruction

G. Webber et al. IEEE Transactions on Medical Imaging 2025

George Webber[1]**,** Yuya Mizuno[2], Oliver D Howes[2], Alexander Hammers[3], Andrew P King[1], Andrew J Reader[1]

1. School of Biomedical Engineering & Imaging Sciences, King's College London, UK
2. Institute of Psychiatry, Psychology & Neuroscience, King's College London, UK
3. Guy's and St Thomas' PET Centre & King's College London, UK

EPSRC Centre for Doctoral Training
Smart Medical Imaging
UKRI
GSK

PET-DDS (Decomposed Diffusion Sampling)

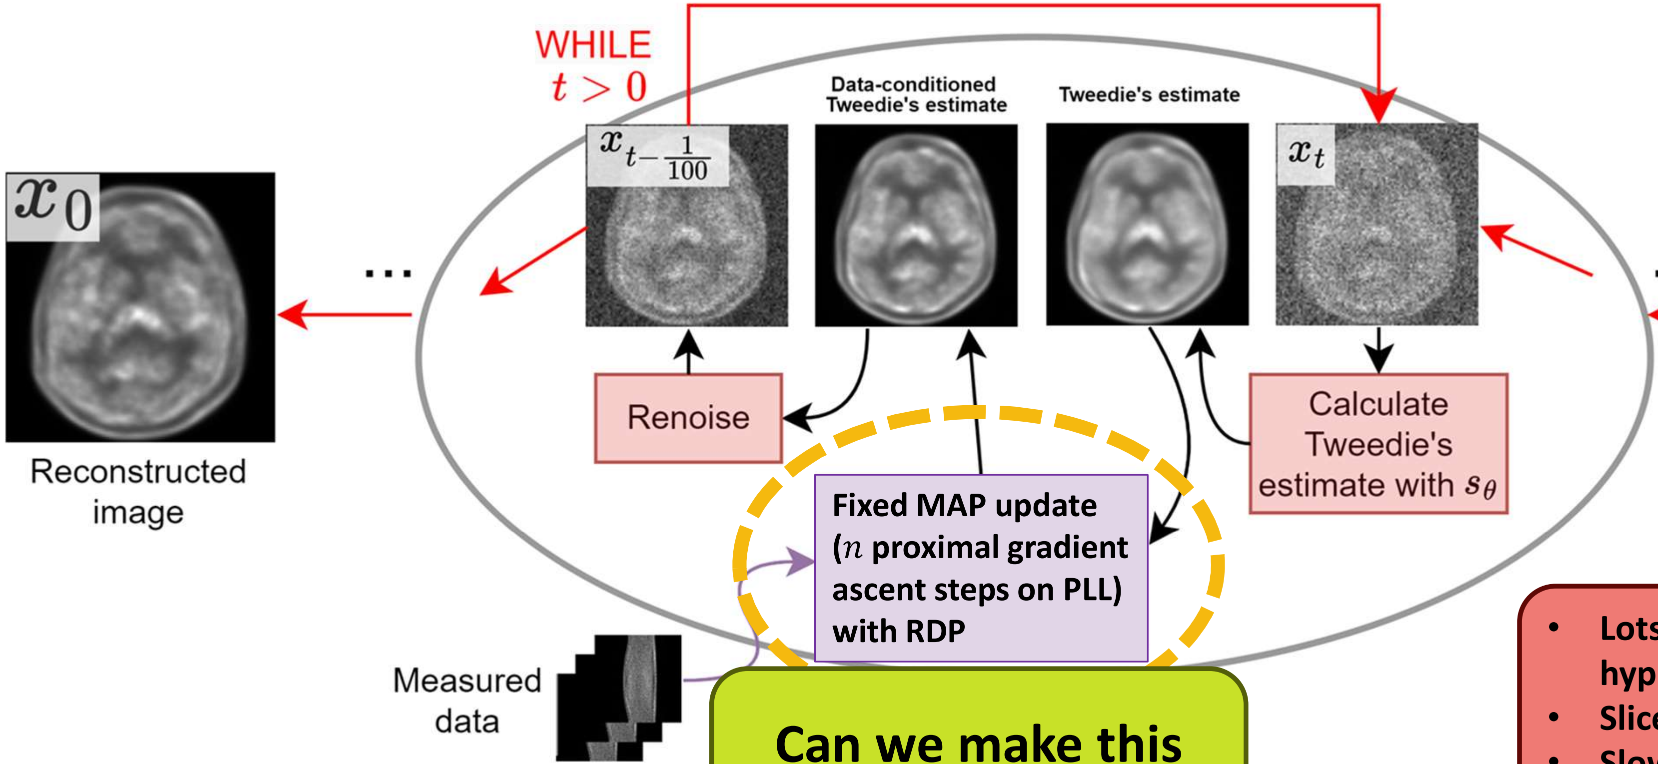

Can we make this data adaptive?

- **Lots of inter-dependent hyperparameters**
- **Slice inconsistencies in 3D**
- **Slow reconstruction**

H. Chung & J. C. Ye, Score-based diffusion models for accelerated MRI. *Med Image Anal*. 2022;80:102479

I. R. D. Singh *et al.* Score-based generative models for PET image reconstruction. *MELBA*. 2024;2(Generative Models):547-585.

Contributions to improve upon DDS

- Fewer hyperparameters (4 down to 1)
- Clinically-motivated selection of that one remaining hyperparameter
- Faster
- Better 3D slice handling

Our *likelihood-scheduled* approach

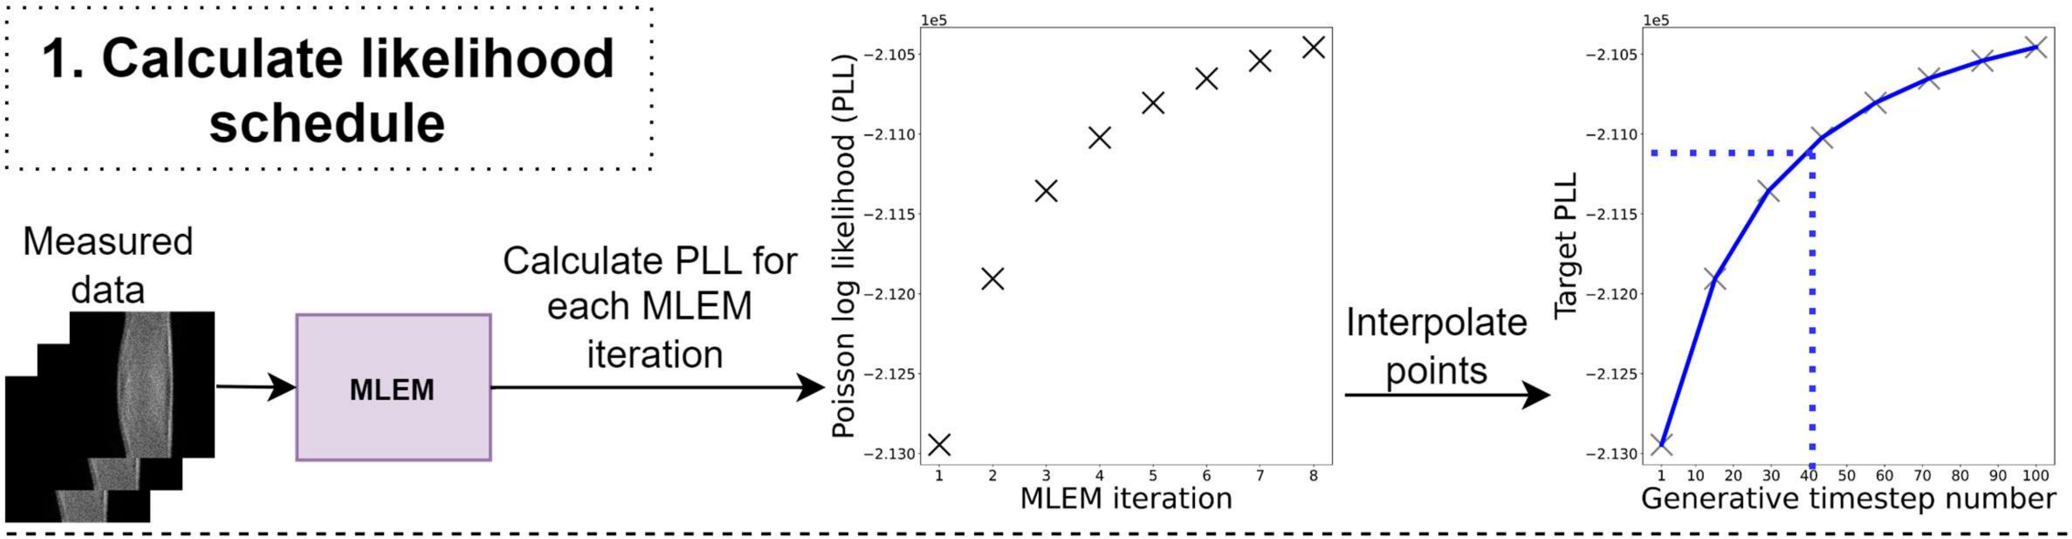

Standard clinical image
reconstruction uses
~50 EM iterations

Our *likelihood-scheduled* approach

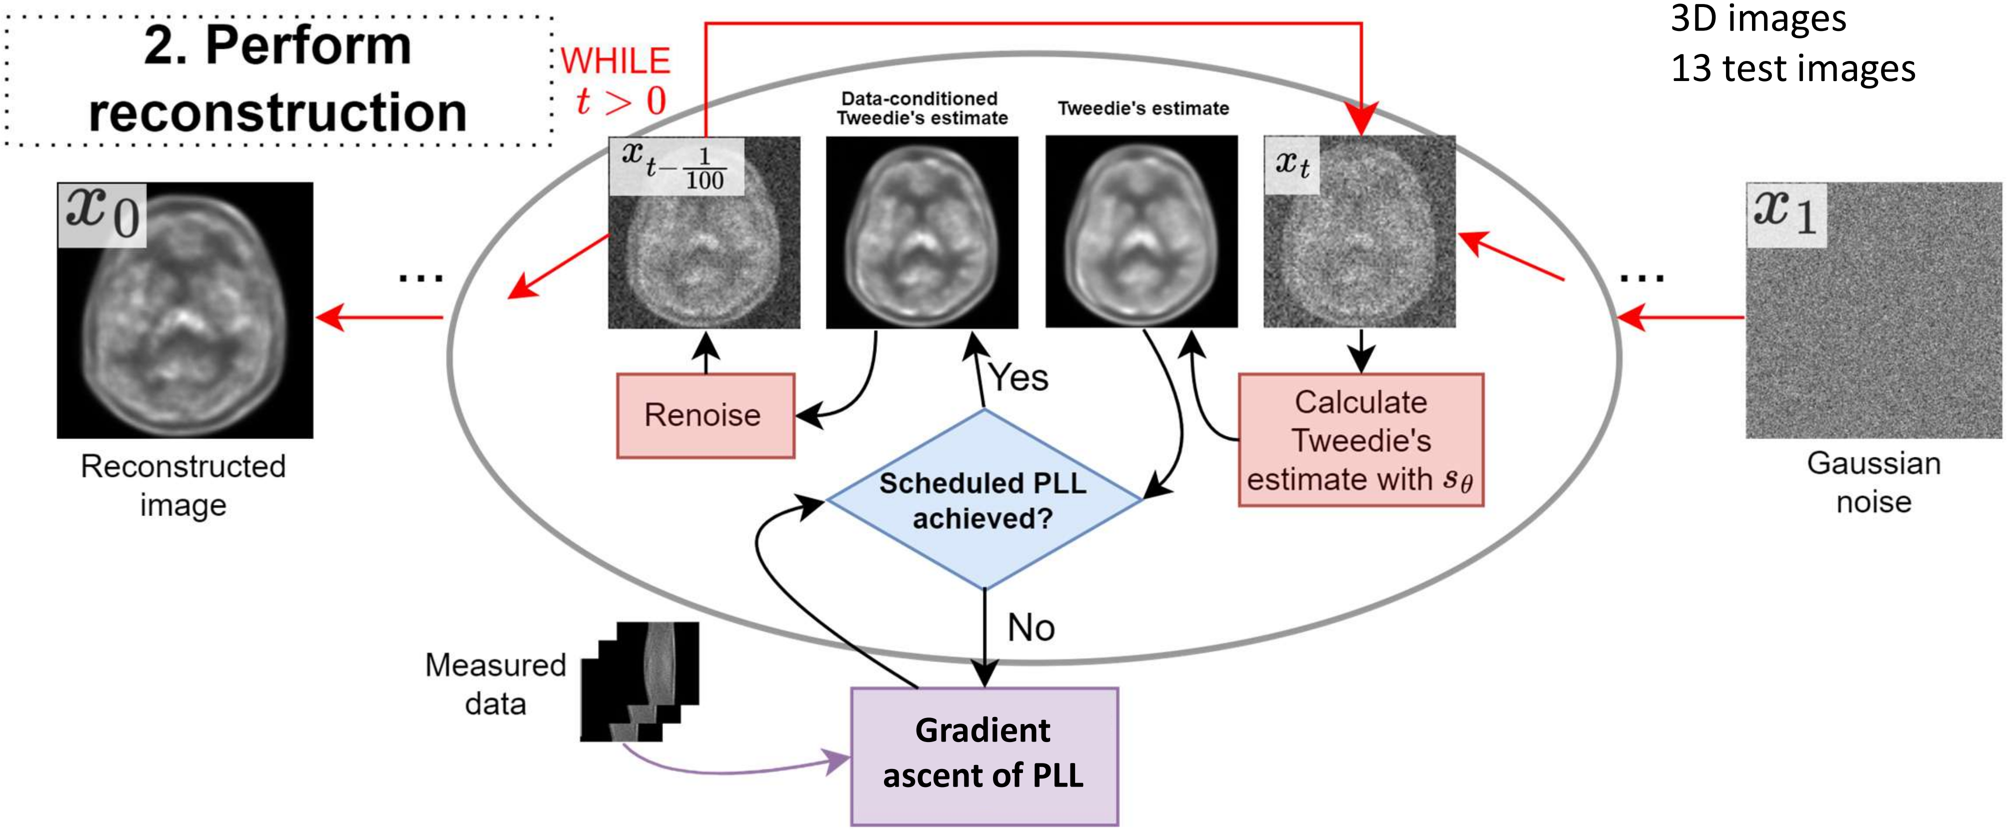

3D real data:
Network trained with 55 3D images
13 test images

Typical penalised reconstruction

Maximize:

$$PLL(\boldsymbol{x}|\boldsymbol{m}) - \lambda U(\boldsymbol{x})$$

Vary λ to vary regularisation

Typical penalised reconstruction

Maximize:
$PLL(\boldsymbol{x}|\boldsymbol{m}) - \lambda U(\boldsymbol{x})$

Vary λ to vary regularisation

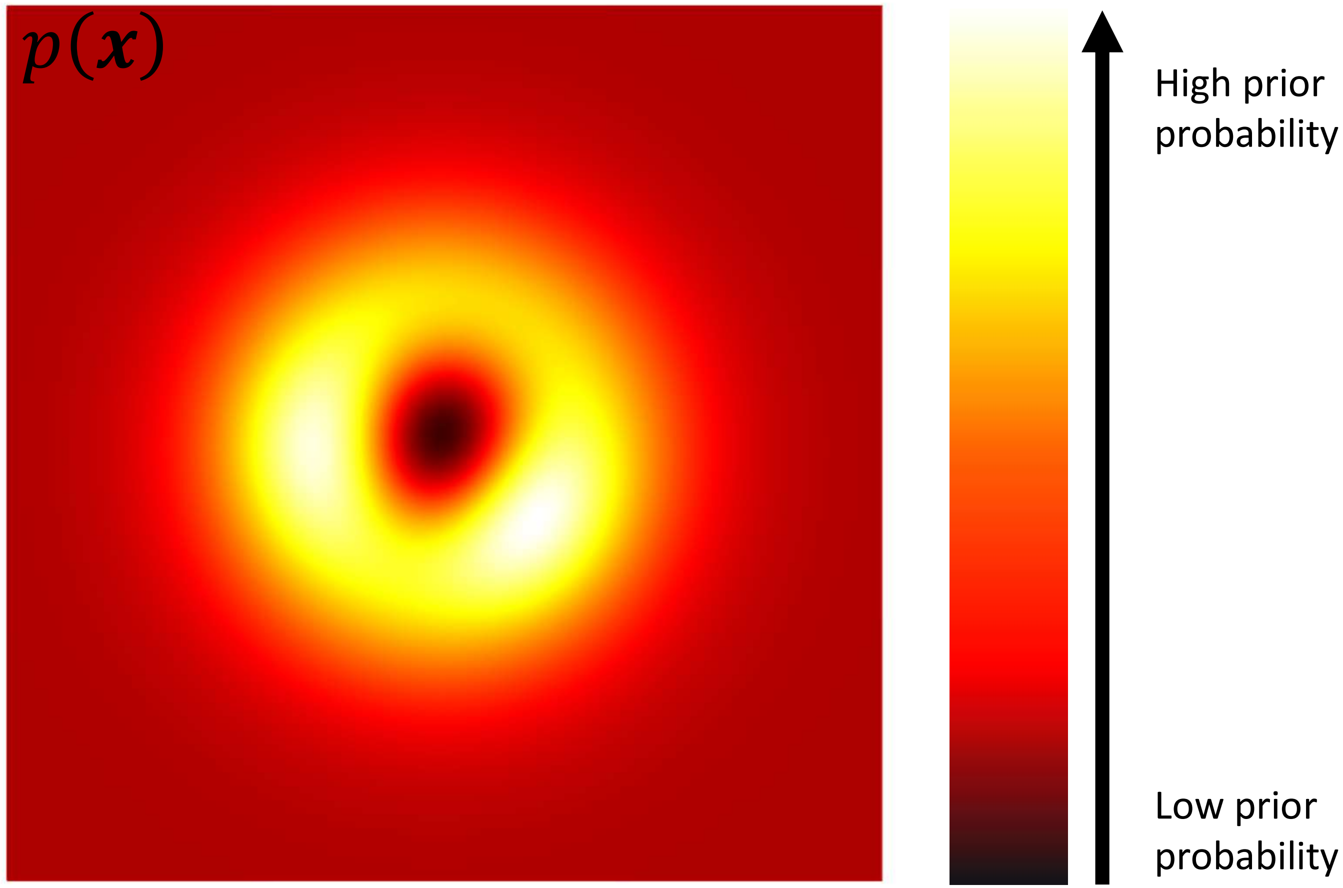

Typical penalised reconstruction

Maximize:

$$PLL(\boldsymbol{x}|\boldsymbol{m}) - \lambda U(\boldsymbol{x})$$

Vary λ to vary regularisation

Iso-likelihood sampling

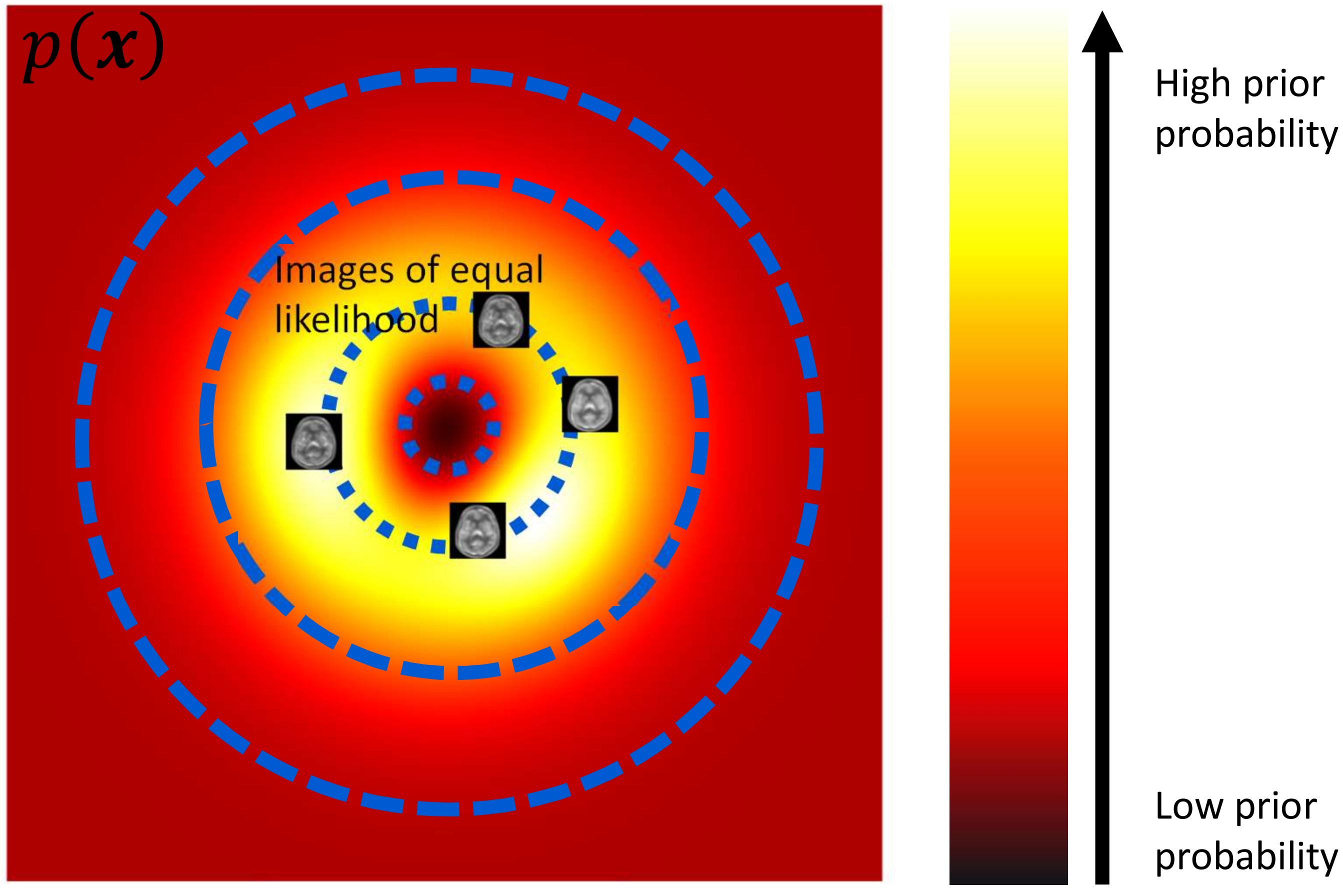

Sample $\boldsymbol{x}$ such that $\boldsymbol{x} \sim p(\boldsymbol{x})$ and $PLL(\boldsymbol{x}|\boldsymbol{m}) = c$

$$c = PLL(\boldsymbol{x}_{MLEM}|\boldsymbol{m})$$

As consistent with $\boldsymbol{m}$ as an MLEM image, but no issue of early-terminated MLEM

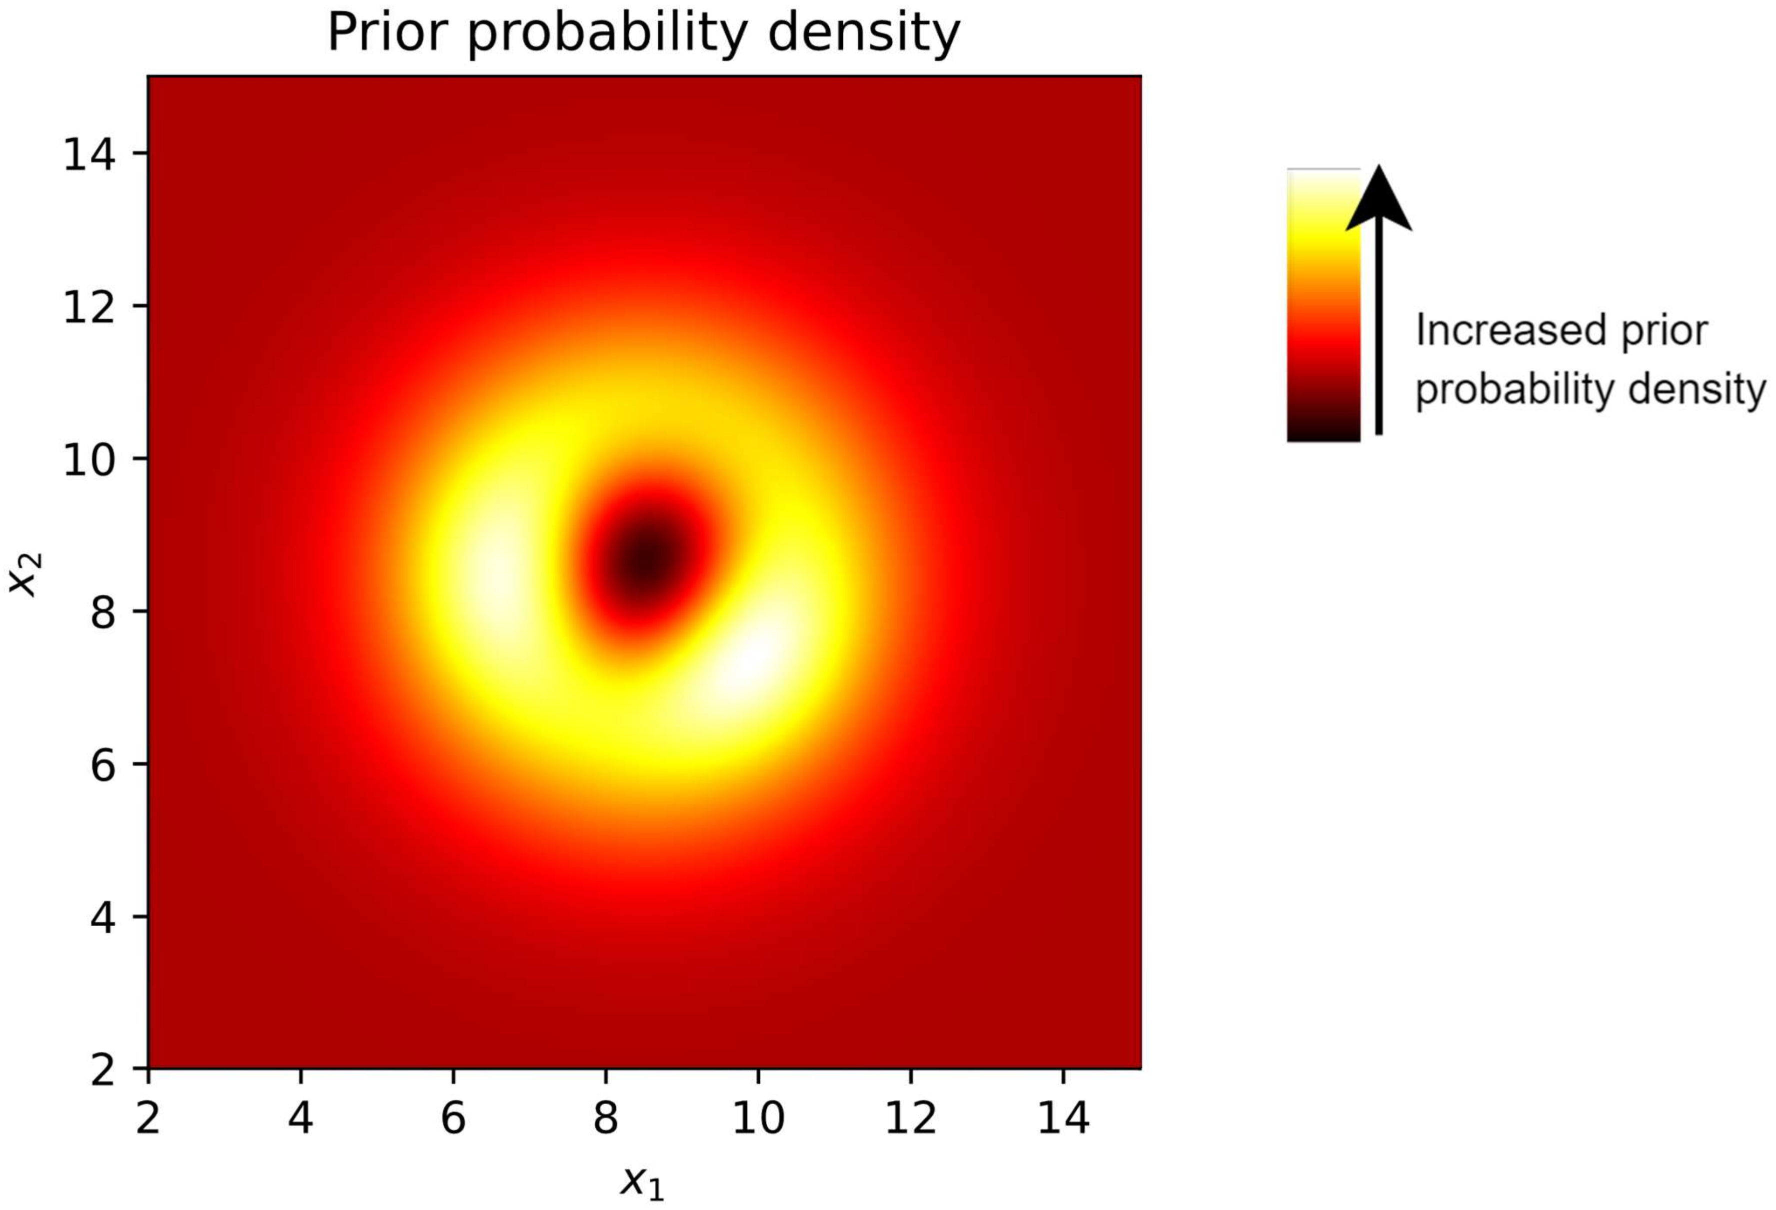

Prior probability density
14
12
10
8
6
4
2
2
4
6
8
10
12
14
x_1
x_2
Increased prior
probability density

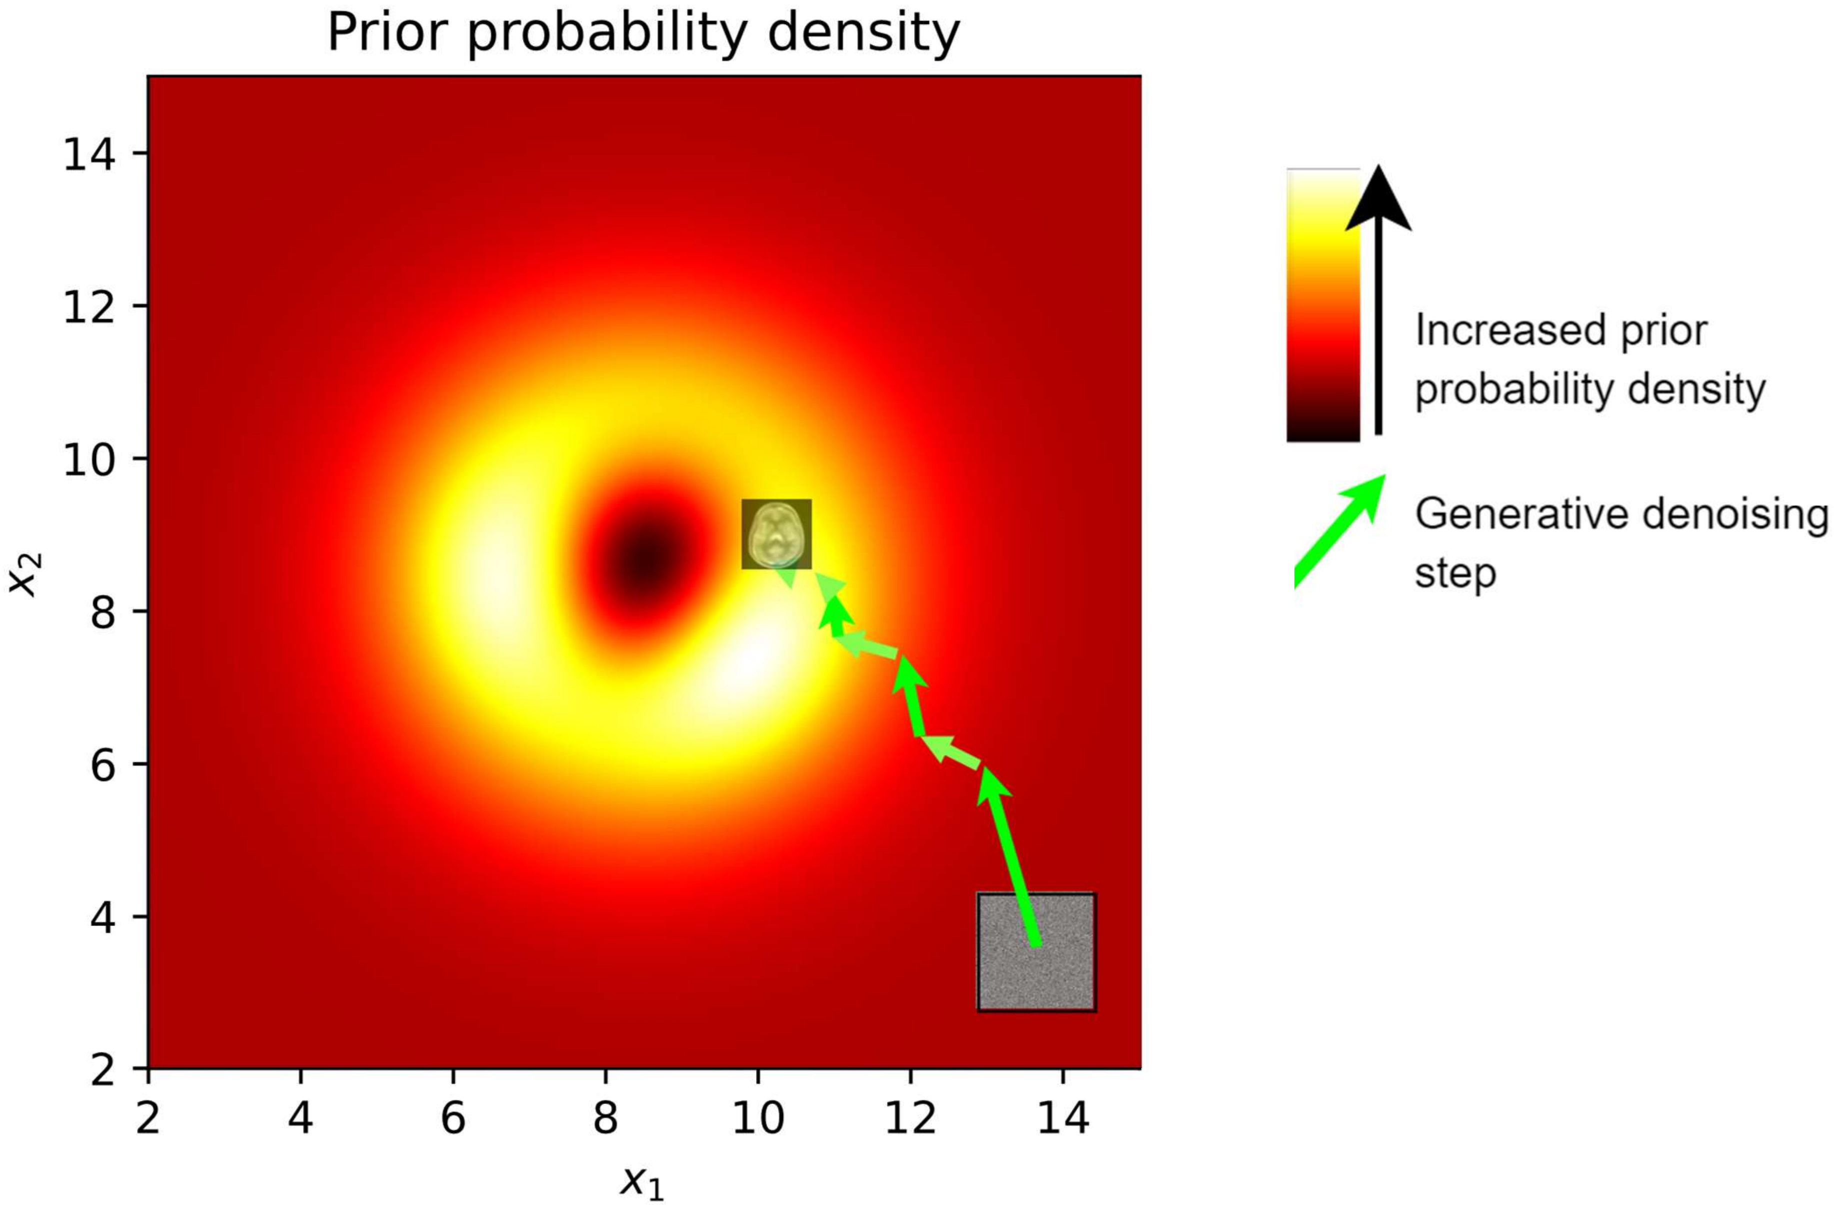
Prior probability density
14
12
10
8
6
4
2
2
4
6
8
10
12
14
x_1
x_2
Increased prior probability density
Generative denoising step

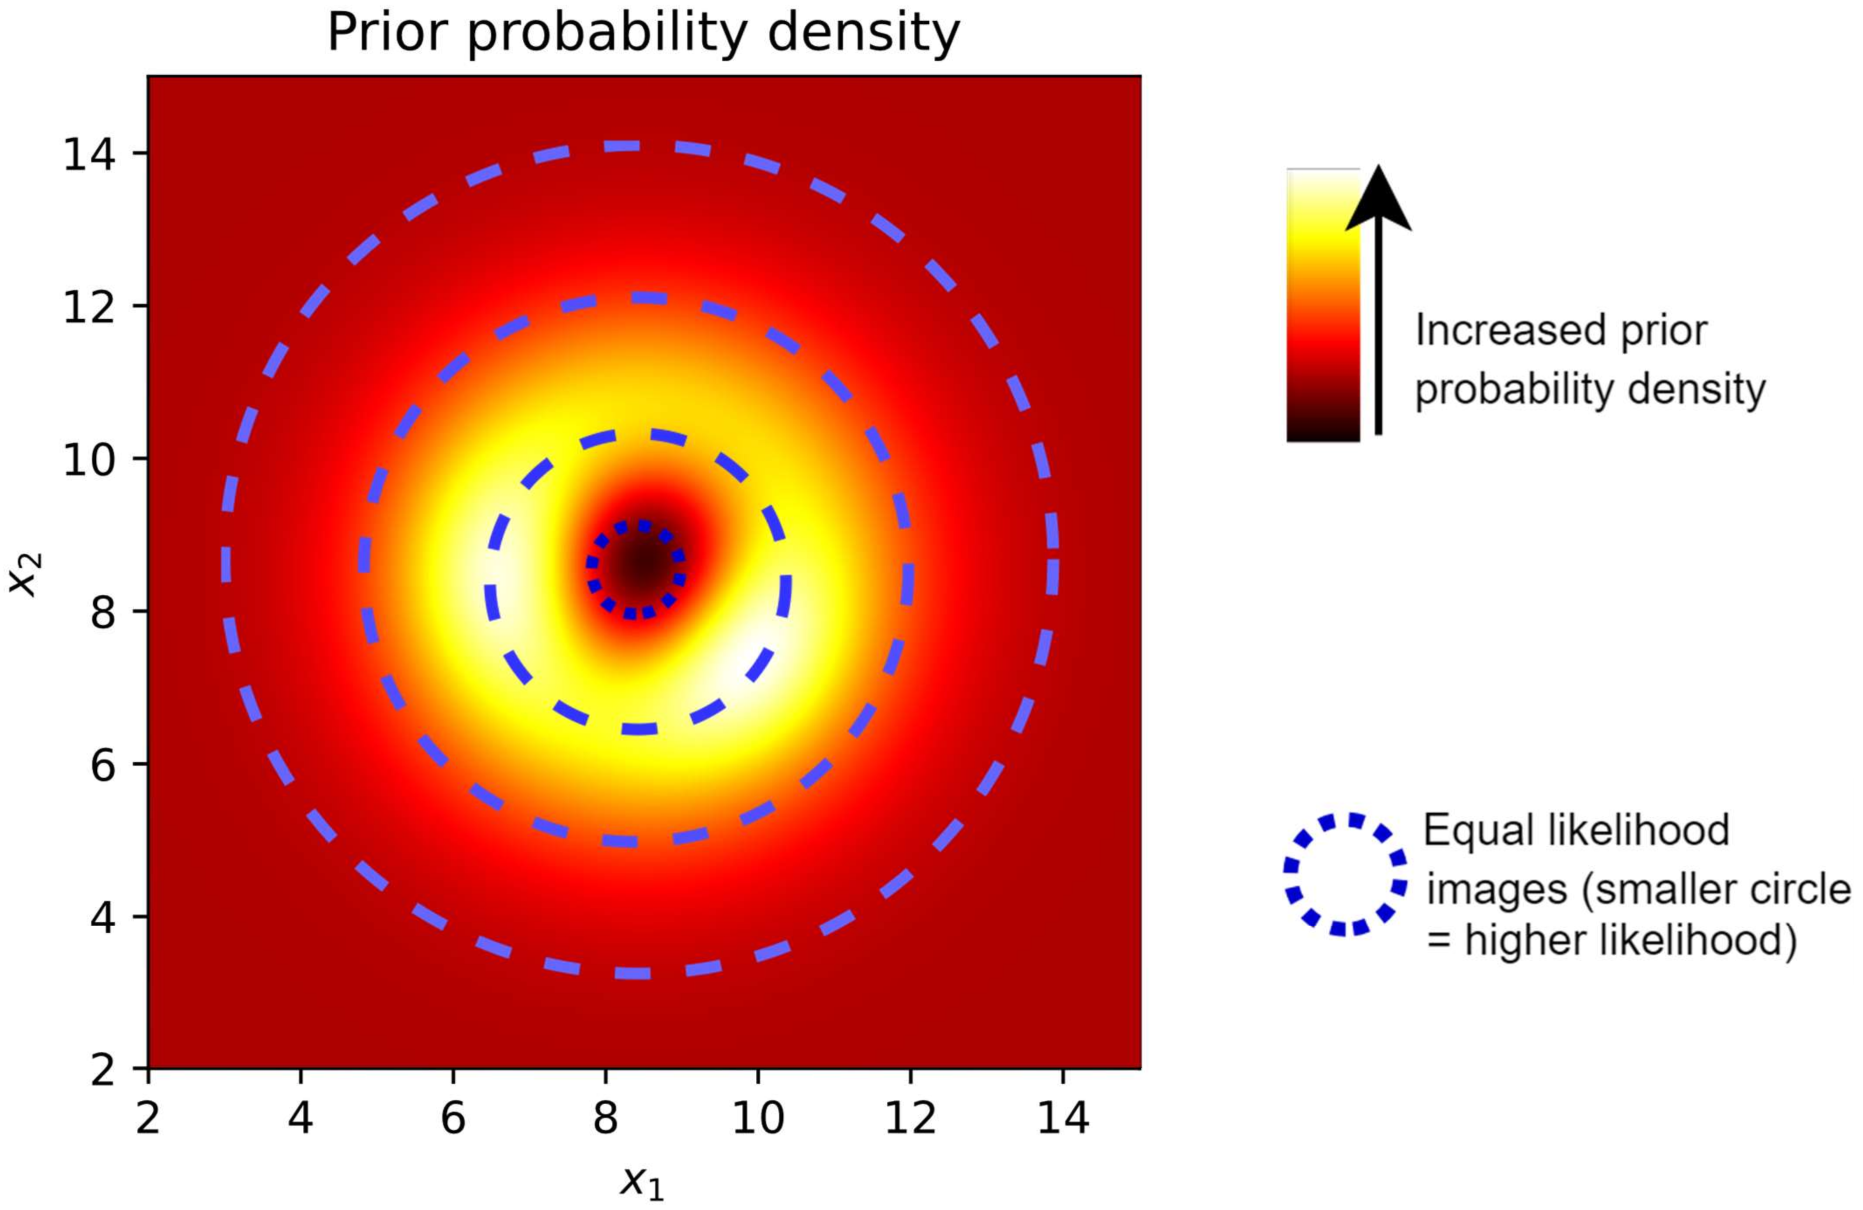
Prior probability density
14
12
10
8
6
4
2
2
4
6
8
10
12
14
x_1
x_2
Increased prior probability density
Equal likelihood images (smaller circle = higher likelihood)

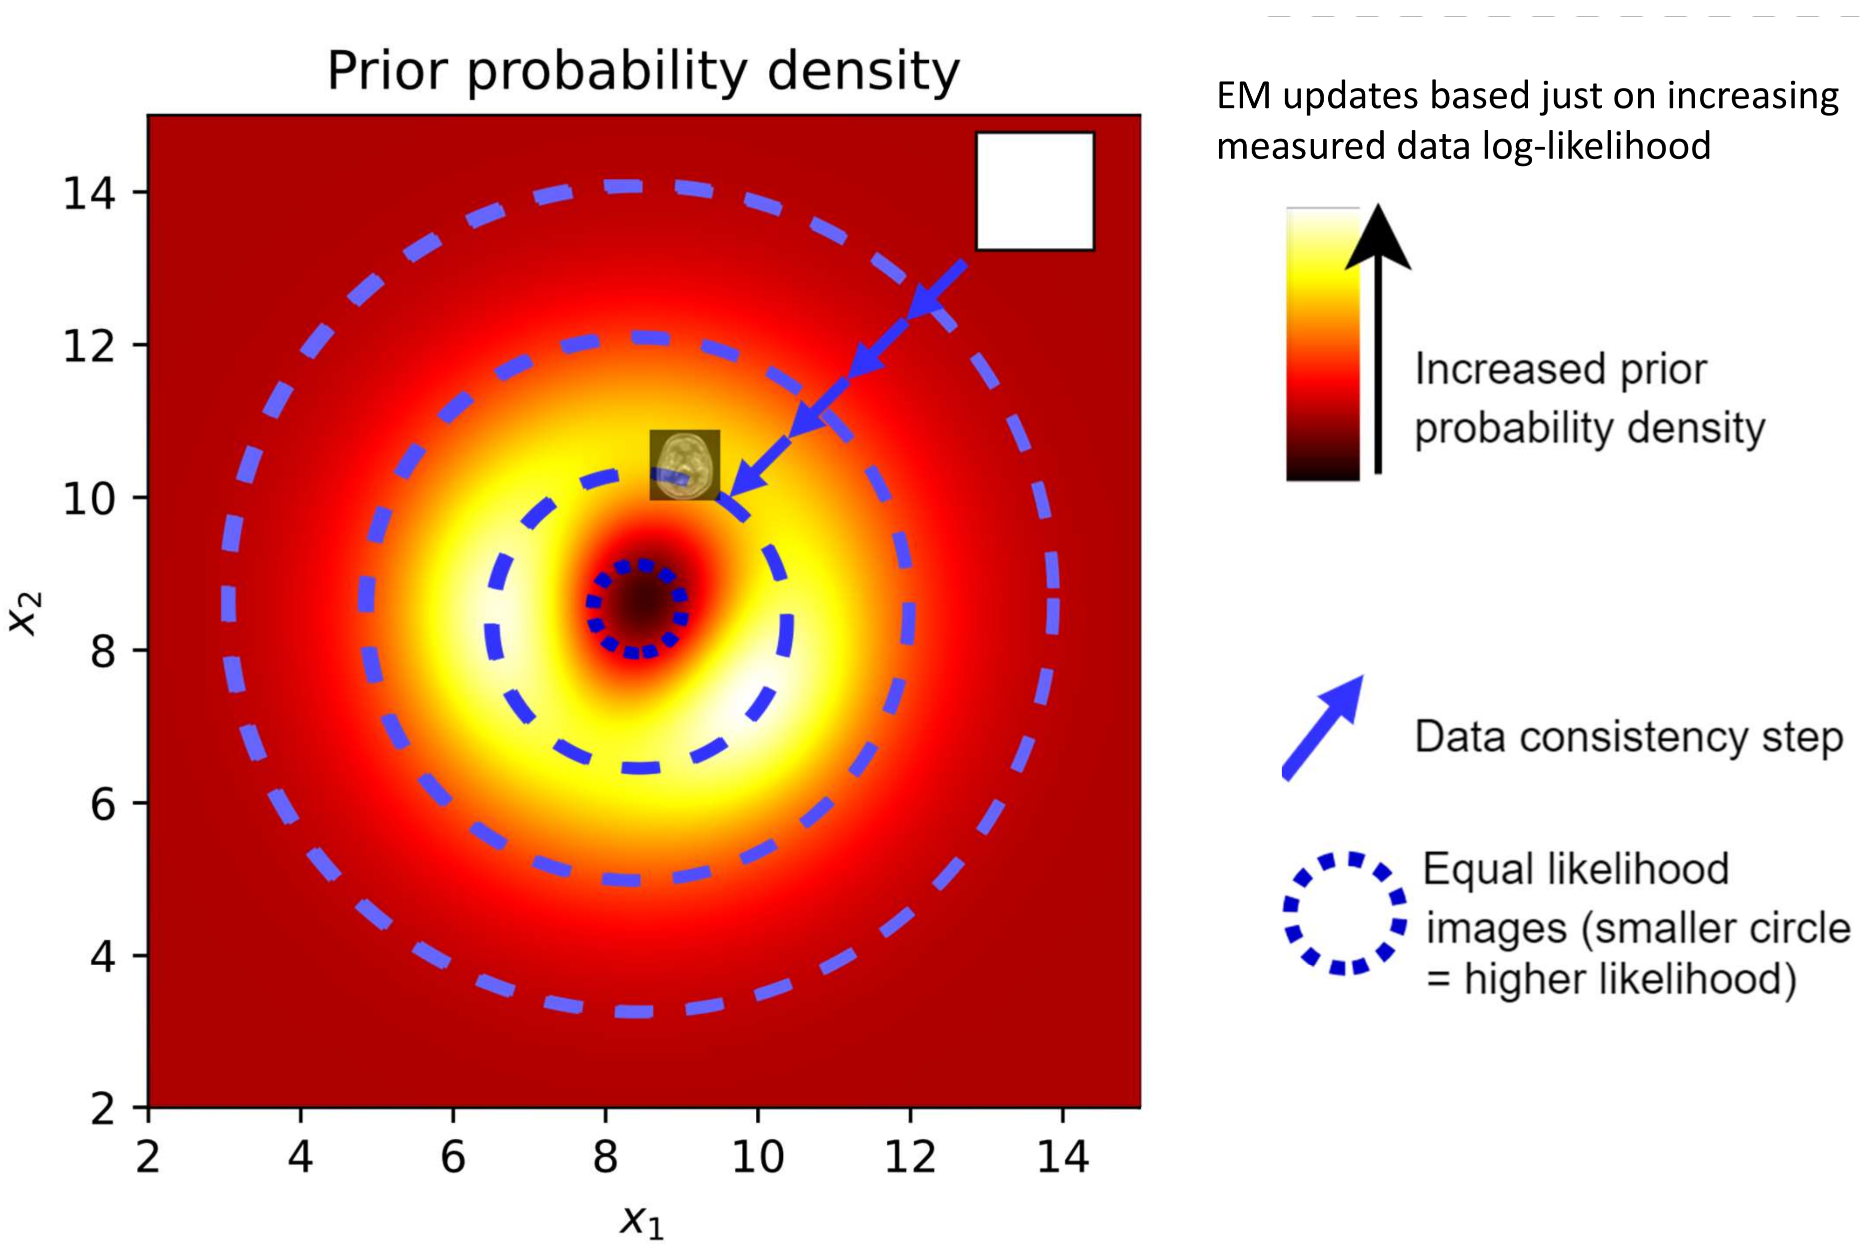
Prior probability density
14
12
10
8
6
4
2
2
4
6
8
10
12
14
x_1
x_2
EM updates based just on increasing measured data log-likelihood
Increased prior probability density
Data consistency step
Equal likelihood images (smaller circle = higher likelihood)

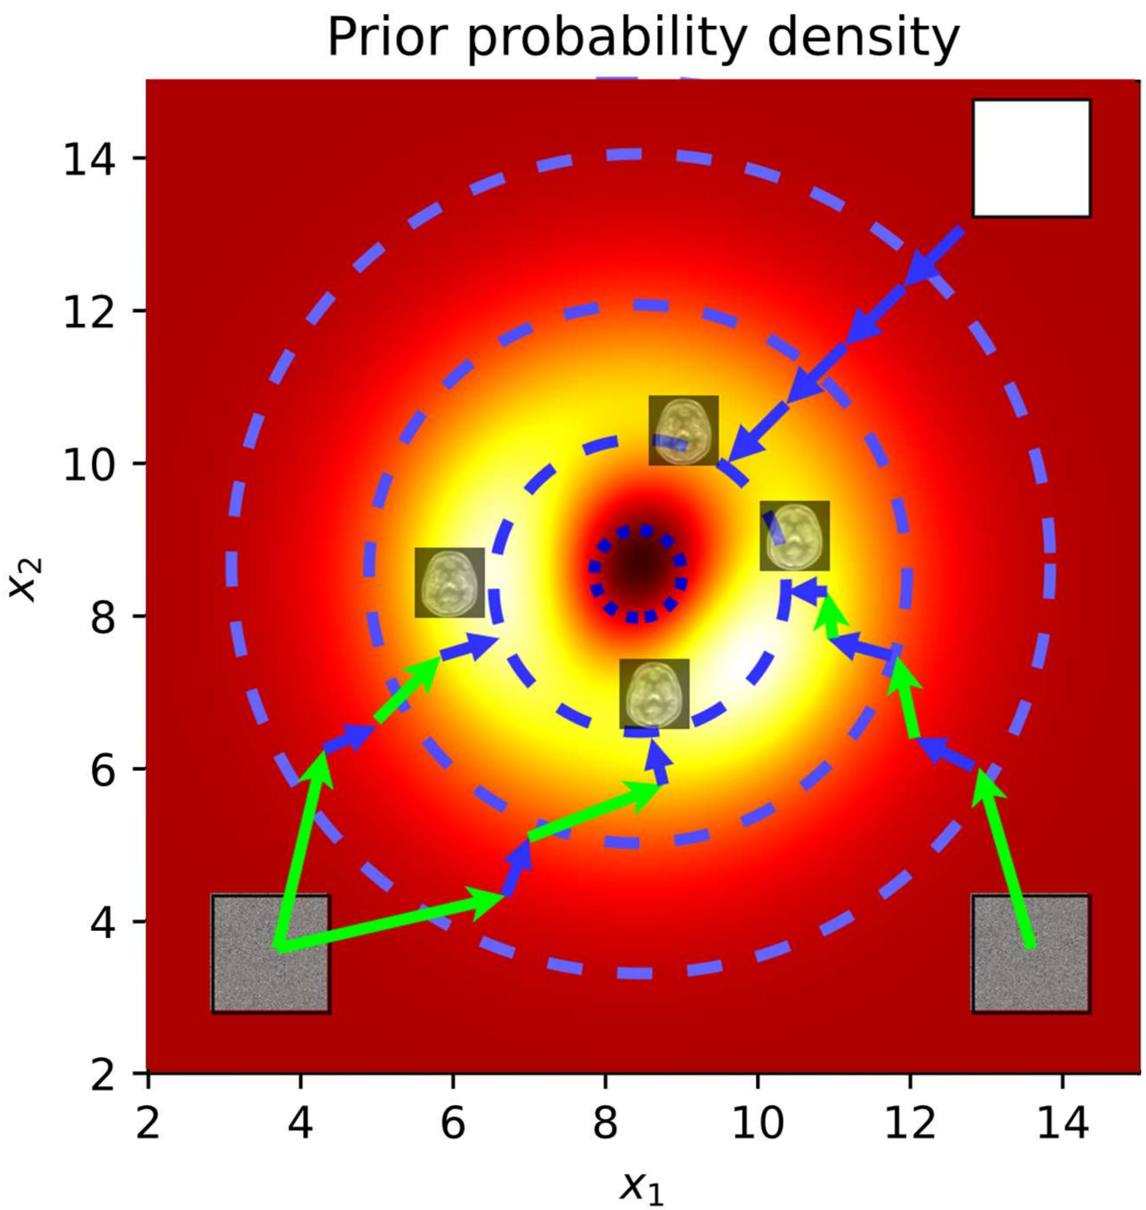

Image Reconstruction by Likelihood Scheduled Diffusion Sampling

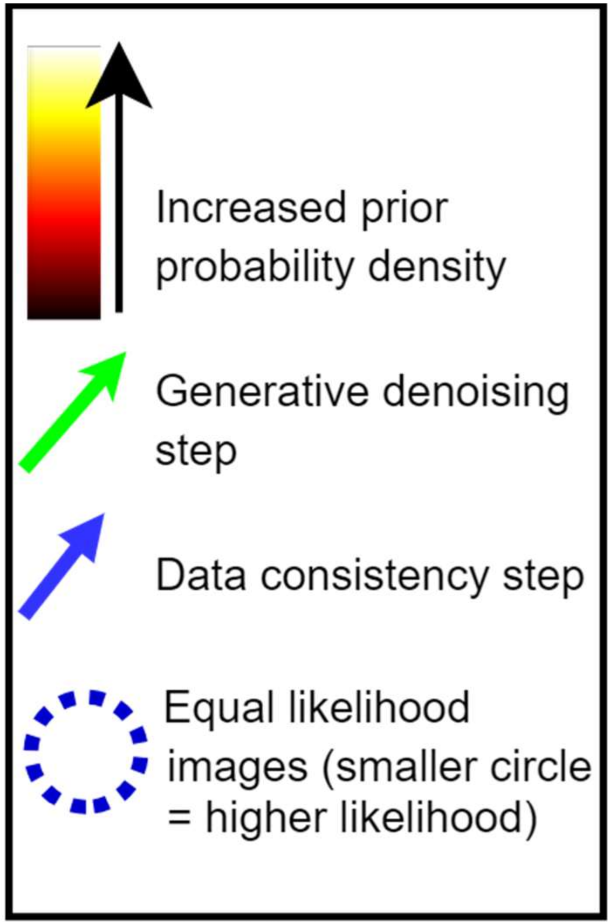

Many samples can be obtained
Balance between likelihood and prior determined by total iterations of the likelihood schedule

Fully 3D Reconstruction

PET-DDS: relative difference prior (RDP)

- Score update: s_θ^{xy}
- Likelihood update: proximal update
  - Gradient ascent step
  - Anchor step
  - Relative difference prior in z direction

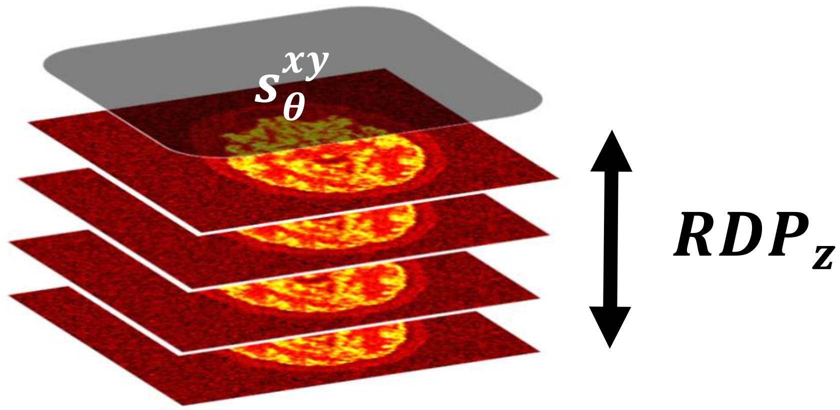

Our approach: 3 perpendicular score models

- Score update: $\frac{1}{3}\left[s_\theta^{xy} + s_\theta^{yz} + s_\theta^{xz}\right]$
- Likelihood update: gradient ascent steps (to a likelihood schedule)

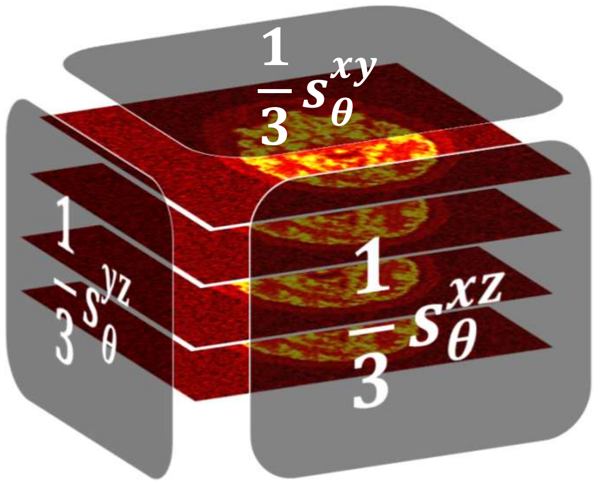

I. R. D. Singh *et al.* Score-based generative models for PET image reconstruction. *MELBA*. 2024;2(Generative Models):547-585.

Fully 3D Reconstruction

1. number of generative steps

PET-DDS: relative difference prior (RDP)

- Score update: s_θ^{xy}

2. iterations / generative step

- Likelihood update: proximal update
  - Gradient ascent step 3. step size
  - Anchor step 4. anchor weighting
  - Relative difference prior in z direction 5. RDP strength

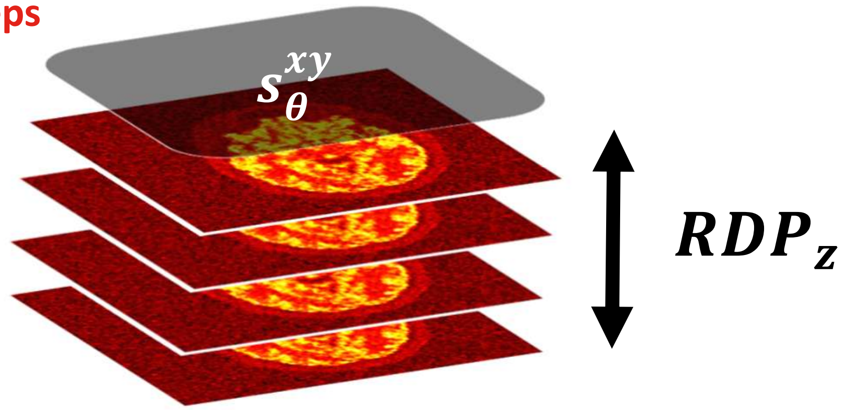

1. number of generative steps

Our approach: 3 perpendicular score models

- Score update: $\frac{1}{3}\left[s_\theta^{xy} + s_\theta^{yz} + s_\theta^{xz}\right]$
- Likelihood update: gradient ascent steps (to a likelihood schedule)

2. target likelihood
(3. step size of gradient ascent)

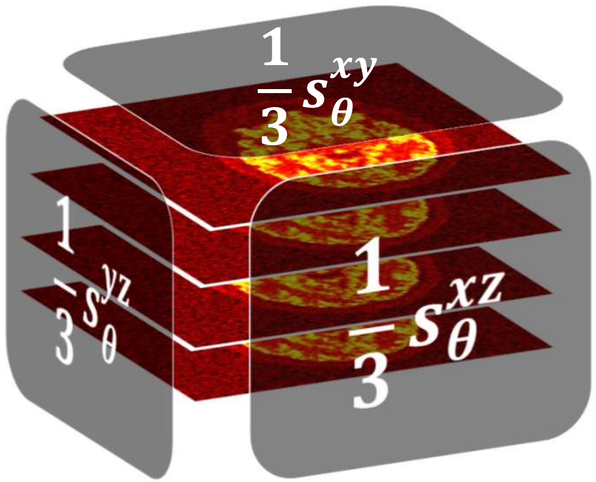

I. R. D. Singh *et al.* Score-based generative models for PET image reconstruction. *MELBA*. 2024;2(Generative Models):547-585.

Advantages of our approach

- ✓ Fewer & independent hyperparameters
- ✓ Fewer likelihood evaluations
- ✓ Corresponds with clinical MLEM

Results

Dynamically varying numbers of likelihood steps

- Our method trivially adapts to different gradient ascent step sizes, enabling faster reconstruction by larger step sizes

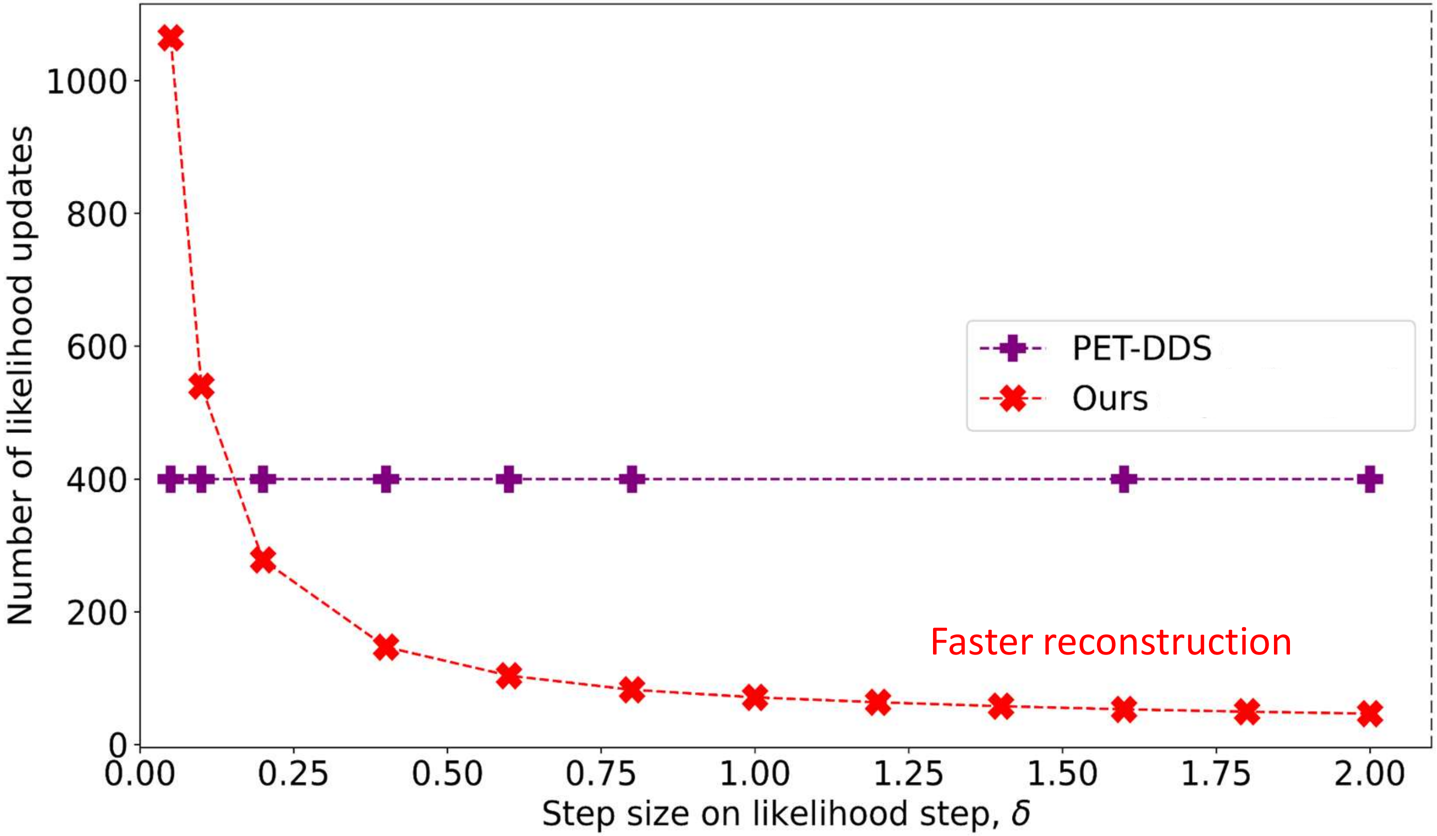

Likelihood updates recorded for 100 steps of reverse diffusion

Example reconstructions – 2D realistic simulation data (low count)

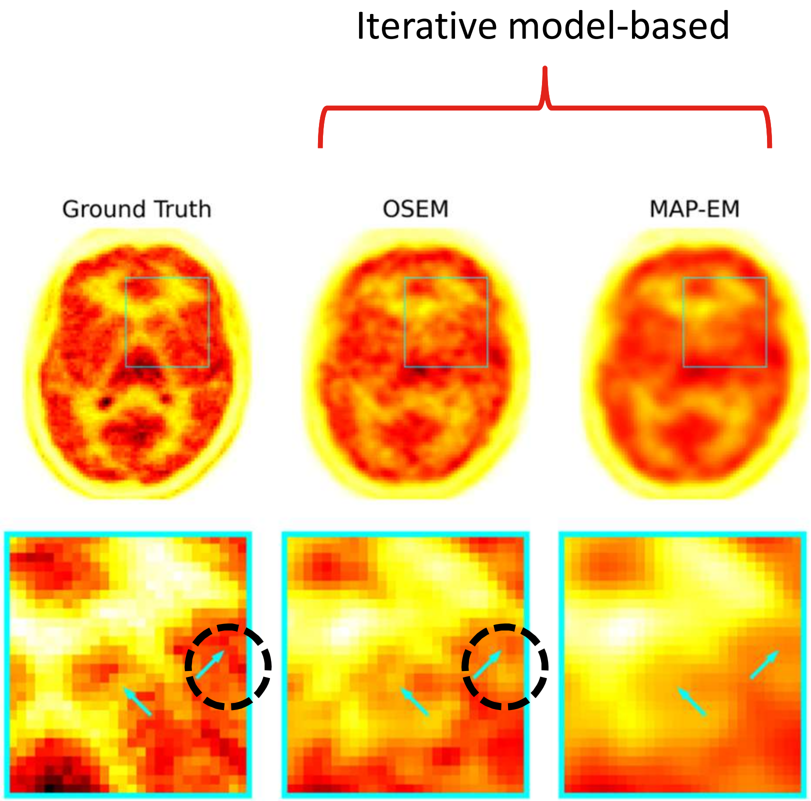

G. Wang, J. Qi. Penalized likelihood PET image reconstruction using patch-based edge-preserving regularization. IEEE Trans Med Imaging. 2012 Dec;31(12):2194-204.
A. Mehranian, A.J. Reader. Model-Based Deep Learning PET Image Reconstruction Using Forward-Backward Splitting Expectation-Maximization. *IEEE Trans Radiat Plasma Med Sci*. 2020 Jun 23;5(1):54-64.

Example reconstructions – 2D realistic simulation data (low count)

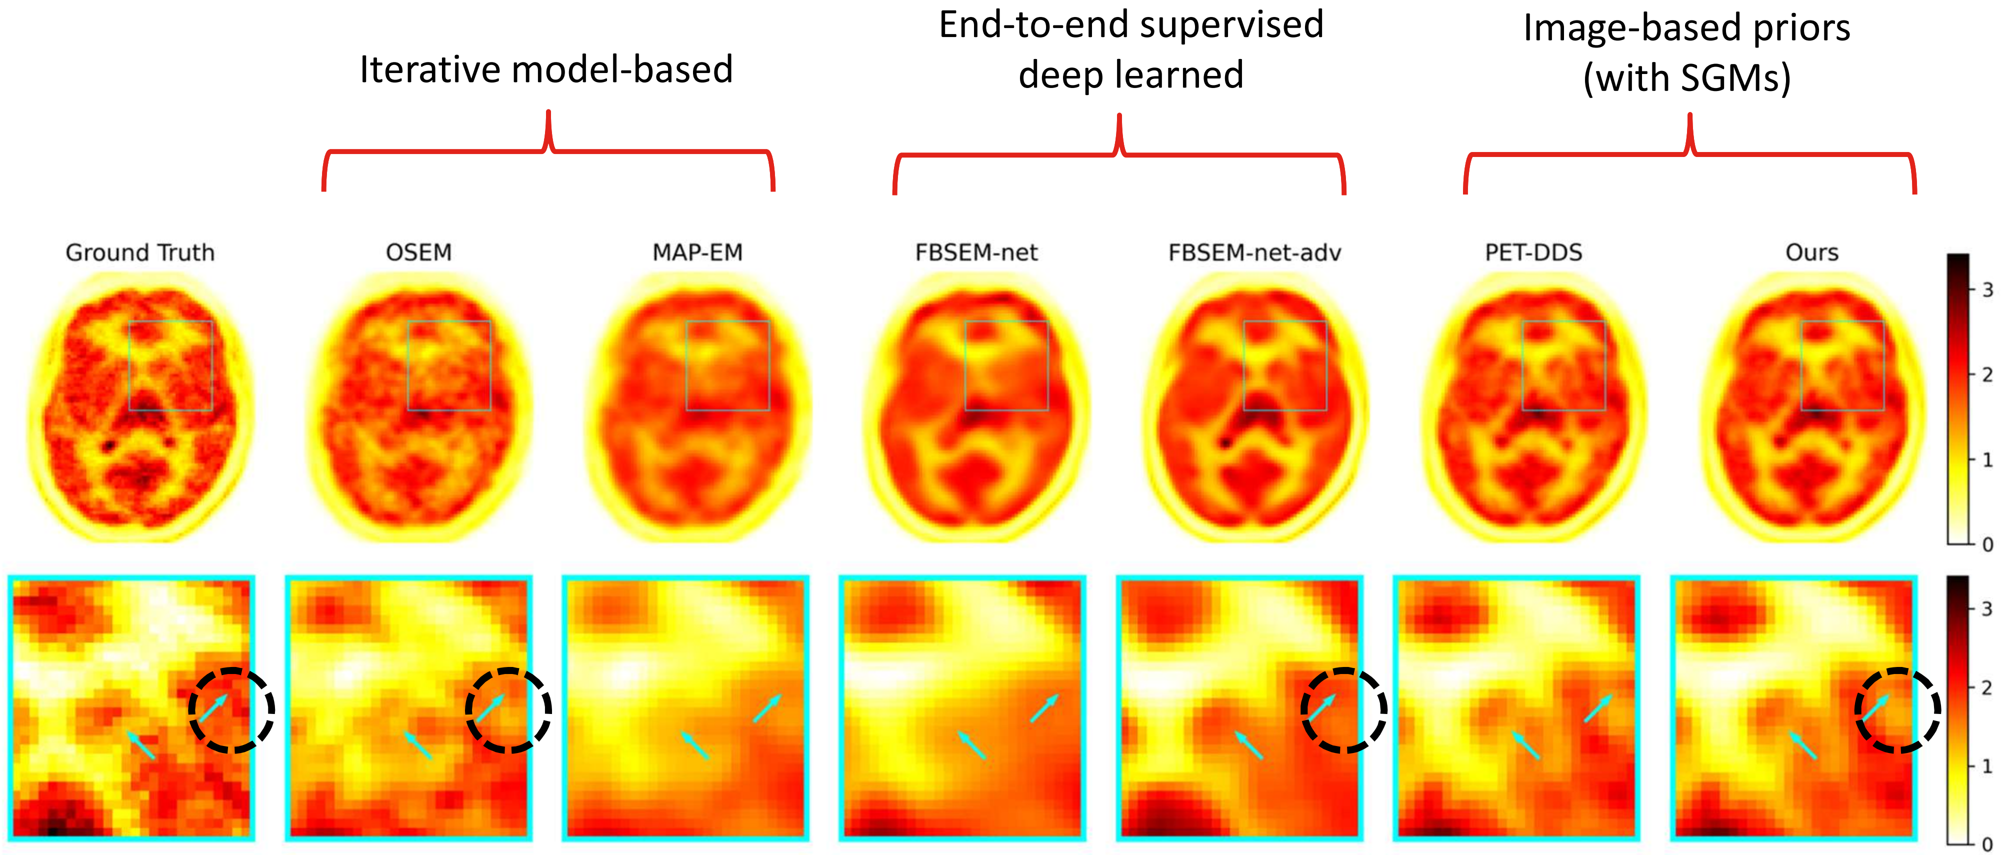

Reconstructions for SGM-based methods are the mean of 5 sample reconstructions.

G. Wang, J. Qi. Penalized likelihood PET image reconstruction using patch-based edge-preserving regularization. IEEE Trans Med Imaging. 2012 Dec;31(12):2194-204.
A. Mehranian, A.J. Reader. Model-Based Deep Learning PET Image Reconstruction Using Forward-Backward Splitting Expectation-Maximization. *IEEE Trans Radiat Plasma Med Sci*. 2020 Jun 23;5(1):54-64.

Example reconstructions – 2D realistic simulation data (low count)

Singh et al 2024

Webber et al 2025

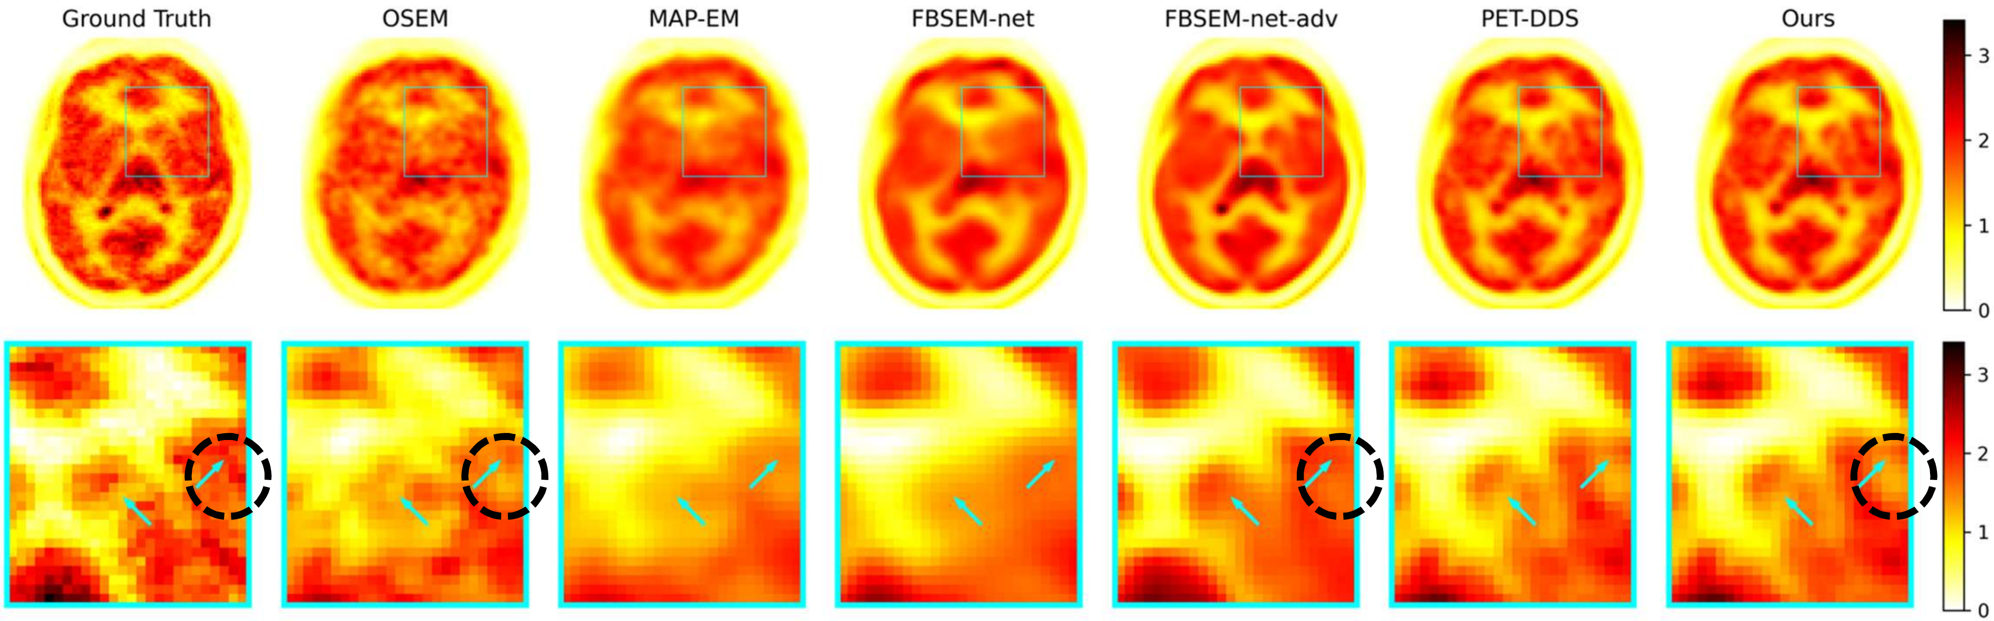

Reconstructions for SGM-based methods are the mean of 5 sample reconstructions.

G. Wang, J. Qi. Penalized likelihood PET image reconstruction using patch-based edge-preserving regularization. IEEE Trans Med Imaging. 2012 Dec;31(12):2194-204.

A. Mehranian, A.J. Reader. Model-Based Deep Learning PET Image Reconstruction Using Forward-Backward Splitting Expectation-Maximization. *IEEE Trans Radiat Plasma Med Sci*. 2020 Jun 23;5(1):54-64.

Bias-variance assessment – 2D realistic simulation data (low count)

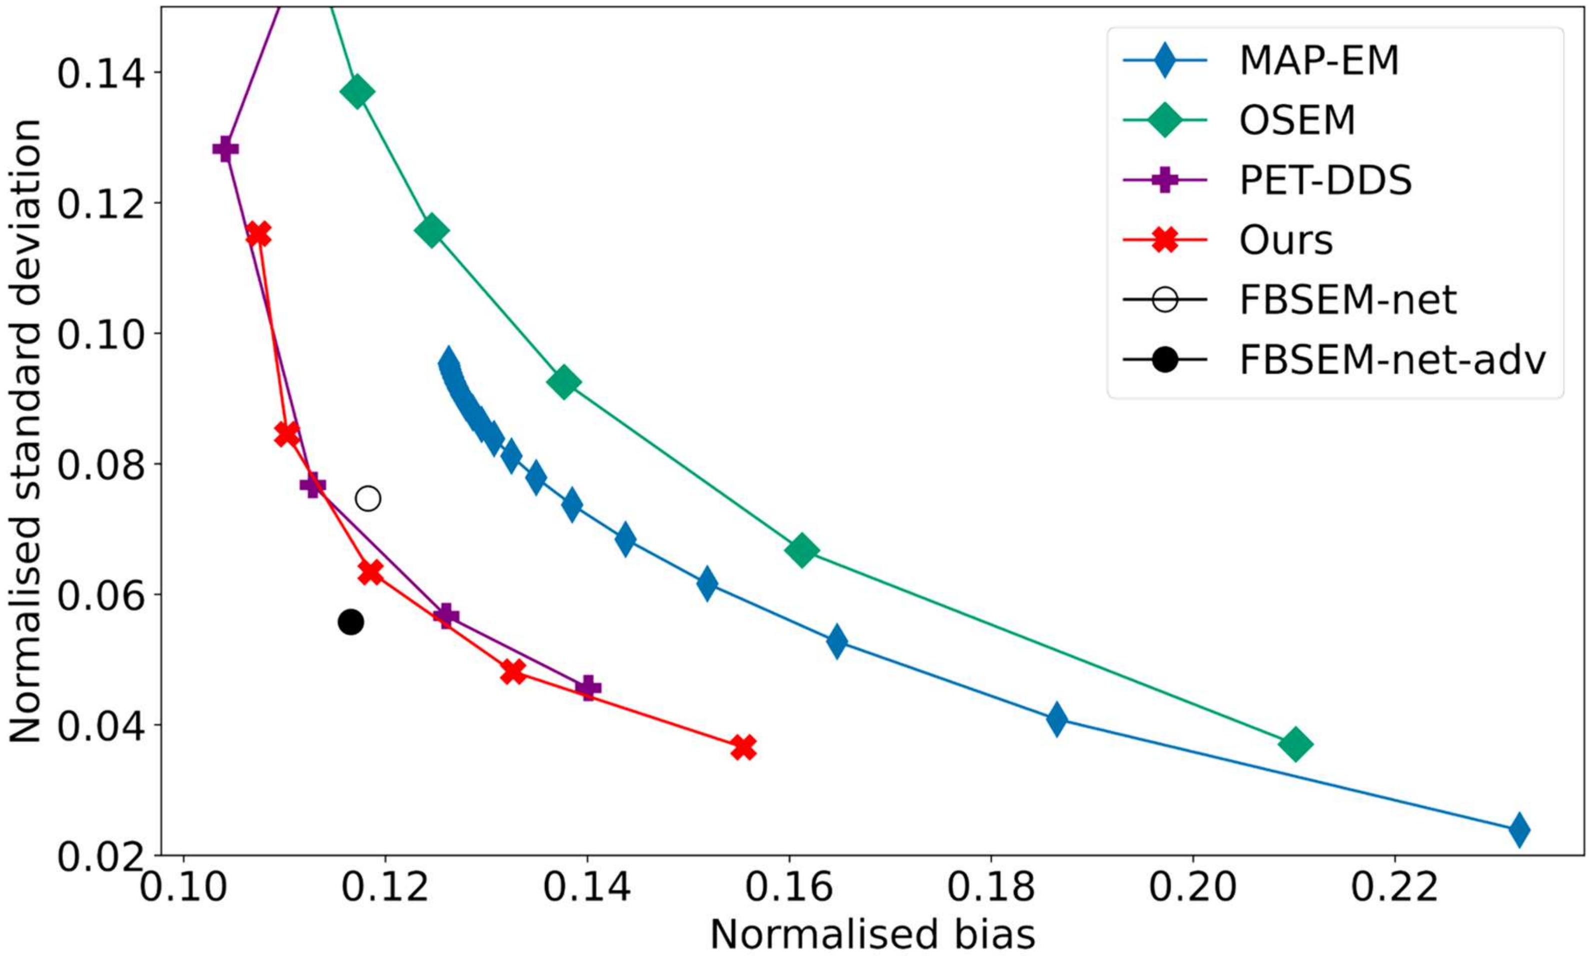

Uncertainty of reconstructed slices (2D realistic sim. data)

- Reduced variation between samples as count increases

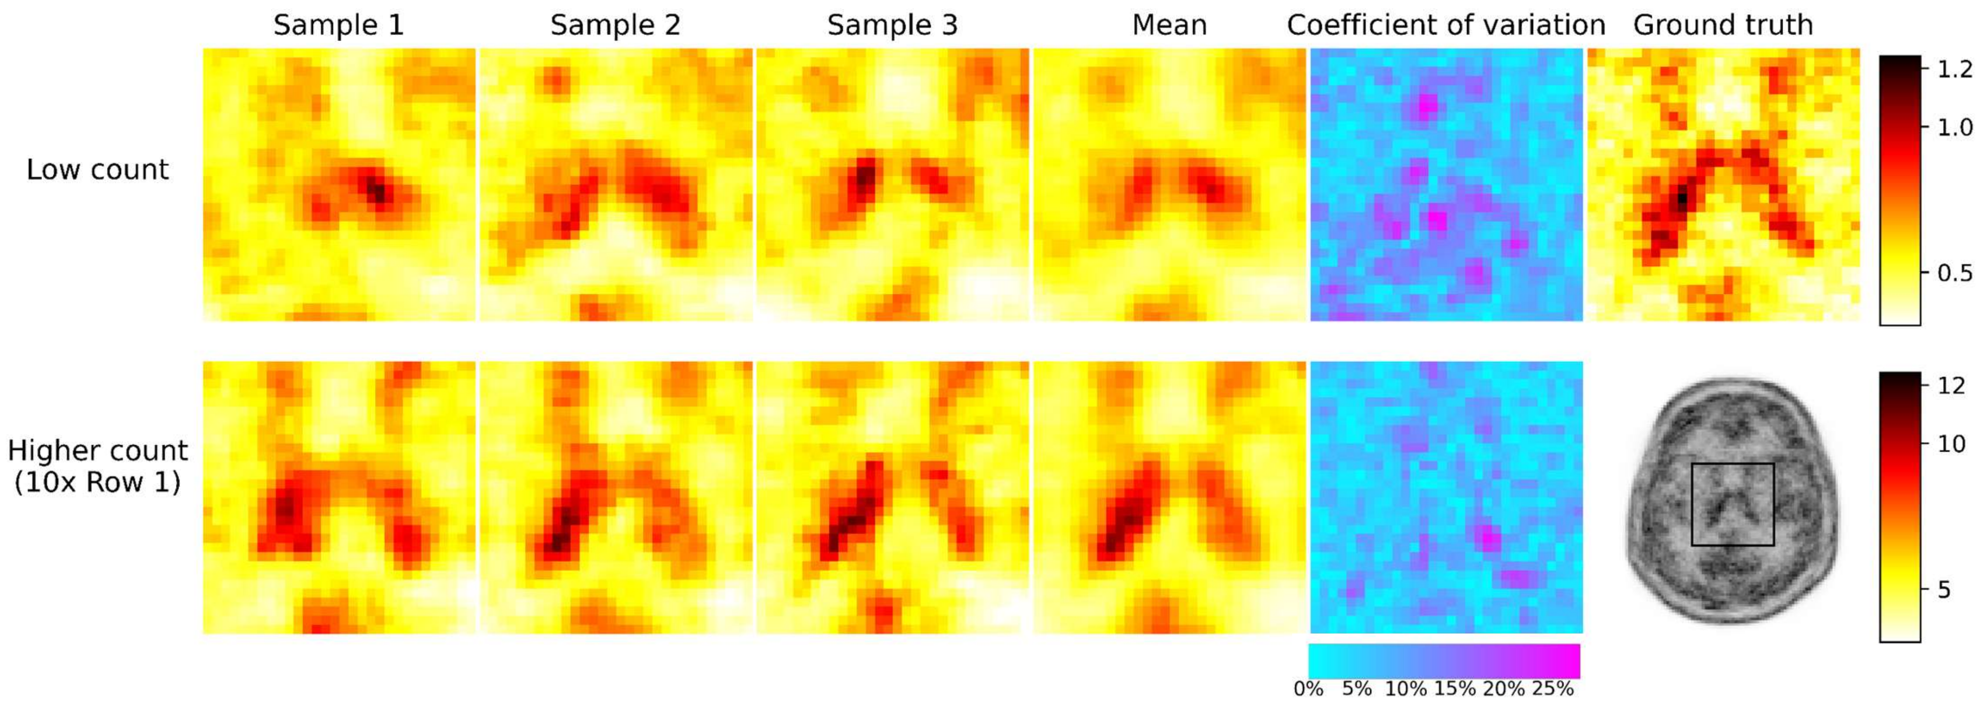

Fully 3D reconstruction (real [^{18}F]-DPA714 data)

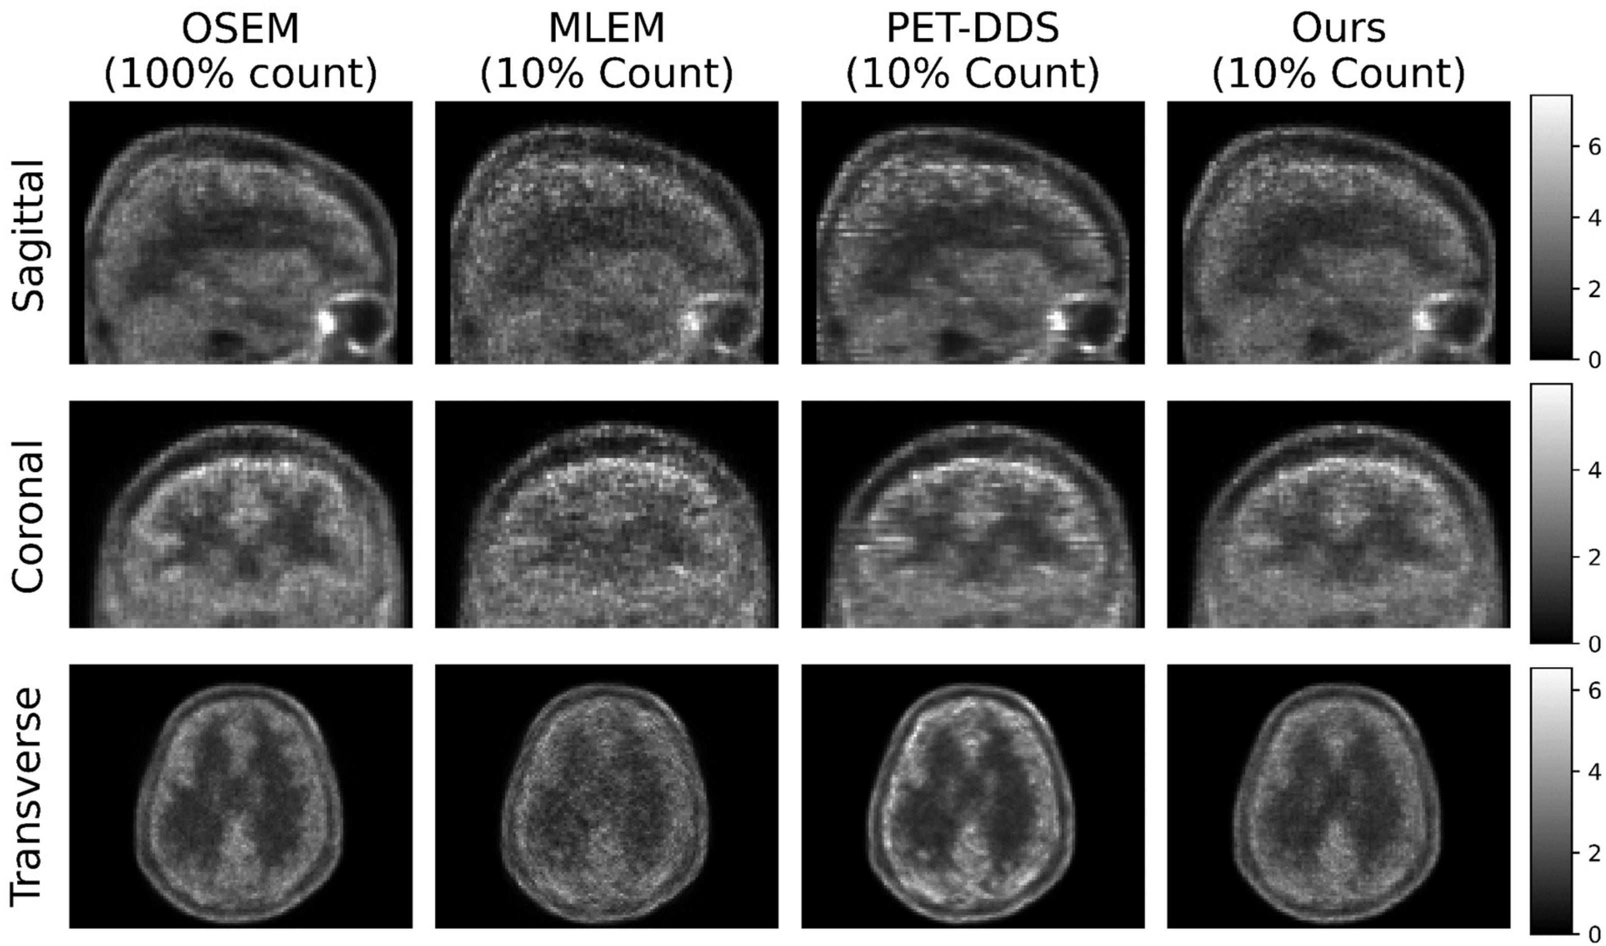

Quality of DM-based method: limited by training data (better training data -> better results!)

BUT FIRST… what if our actual data is out of distribution, if there is domain shift?

Steerable Diffusion Models for PET Image Reconstruction

G. Webber et al. IEEE Medical Imaging Conference 2025
arXiv:2510.13441

George Webber[1], Alexander Hammers[2], Andrew P King[1], Andrew J Reader[1]

1. School of Biomedical Engineering & Imaging Sciences, King's College London, UK
2. Guy's and St Thomas' PET Centre & King's College London, UK

EPSRC Centre for Doctoral Training
Smart Medical Imaging

Conditioning a pre-trained diffusion model prior for PET image recon

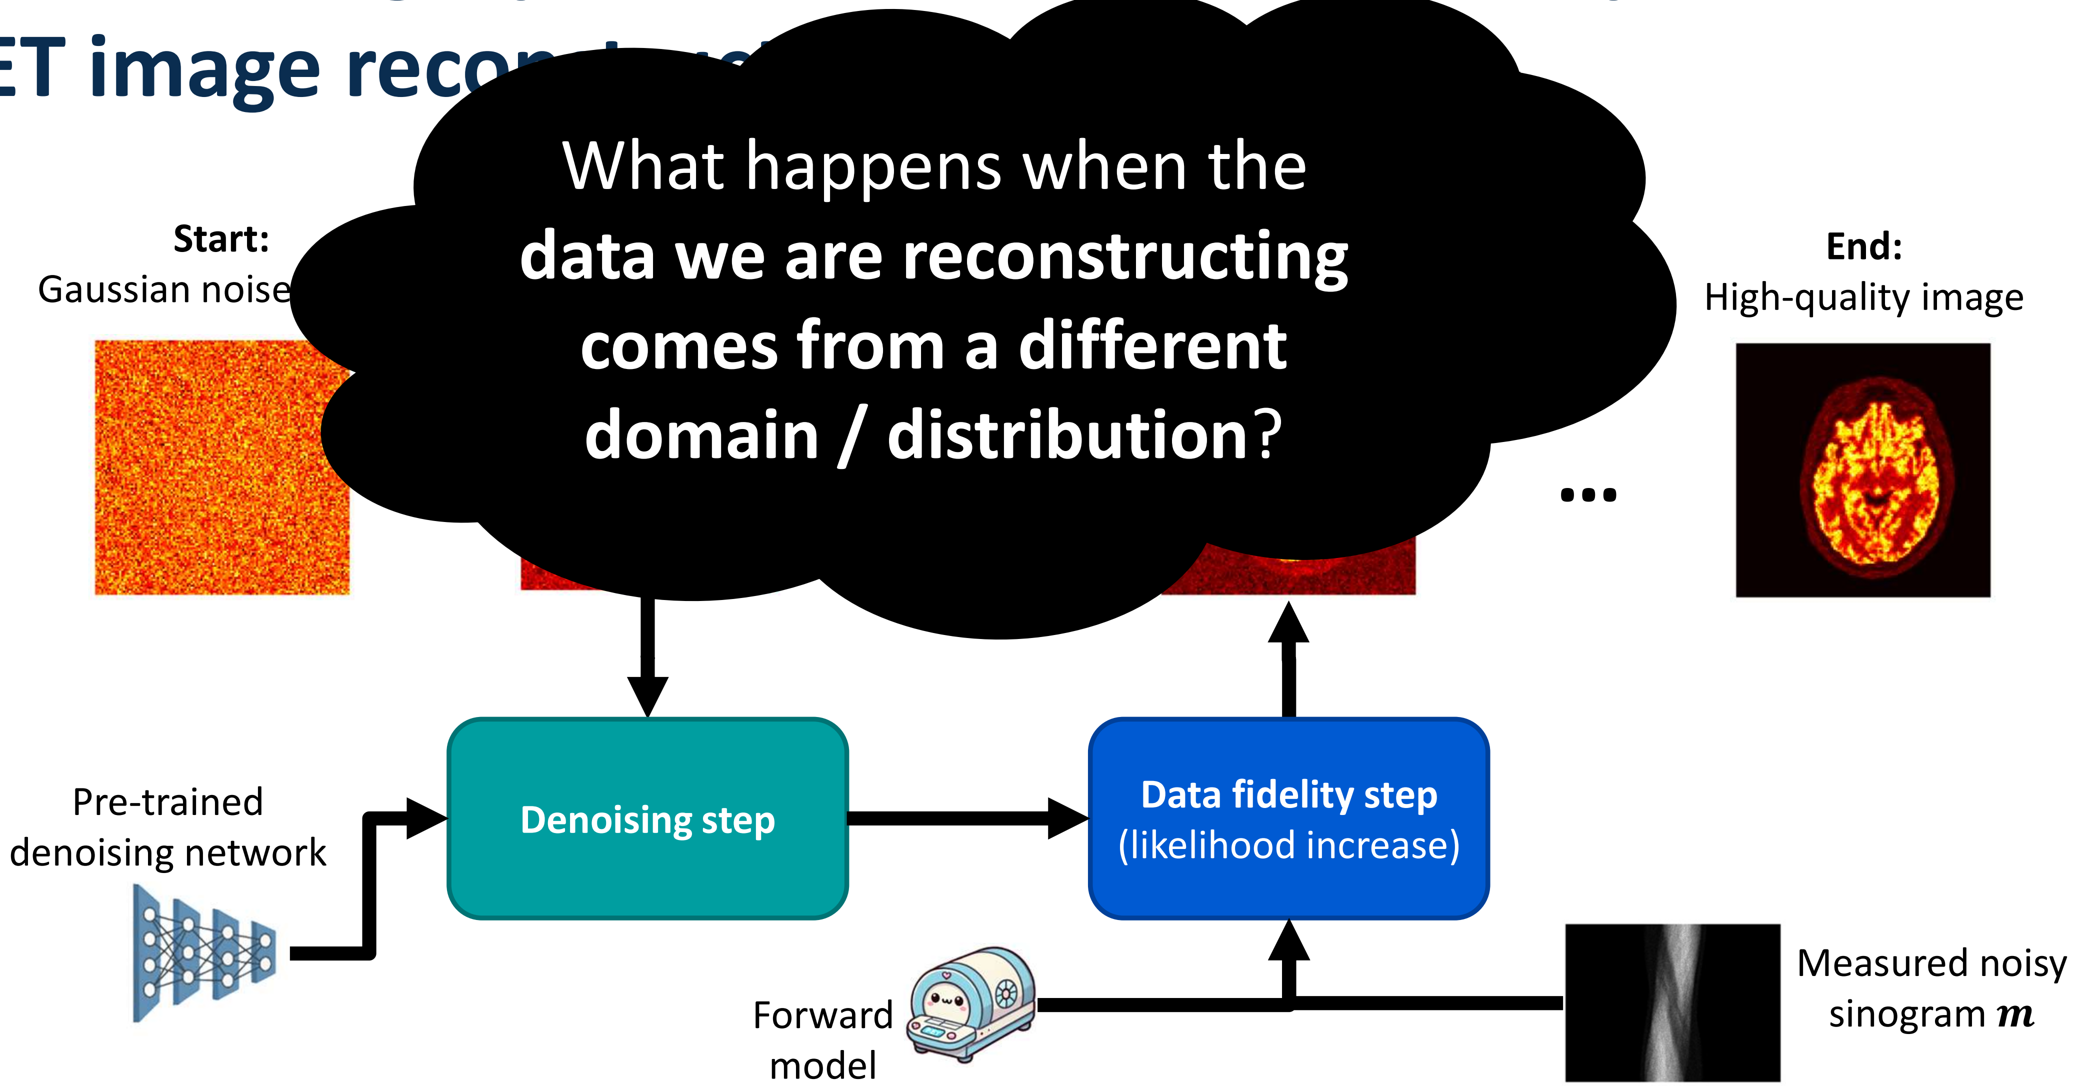

What happens when the **data we are reconstructing comes from a different domain / distribution?**

Adapt the prior to the data

Steerable Conditional Diffusion

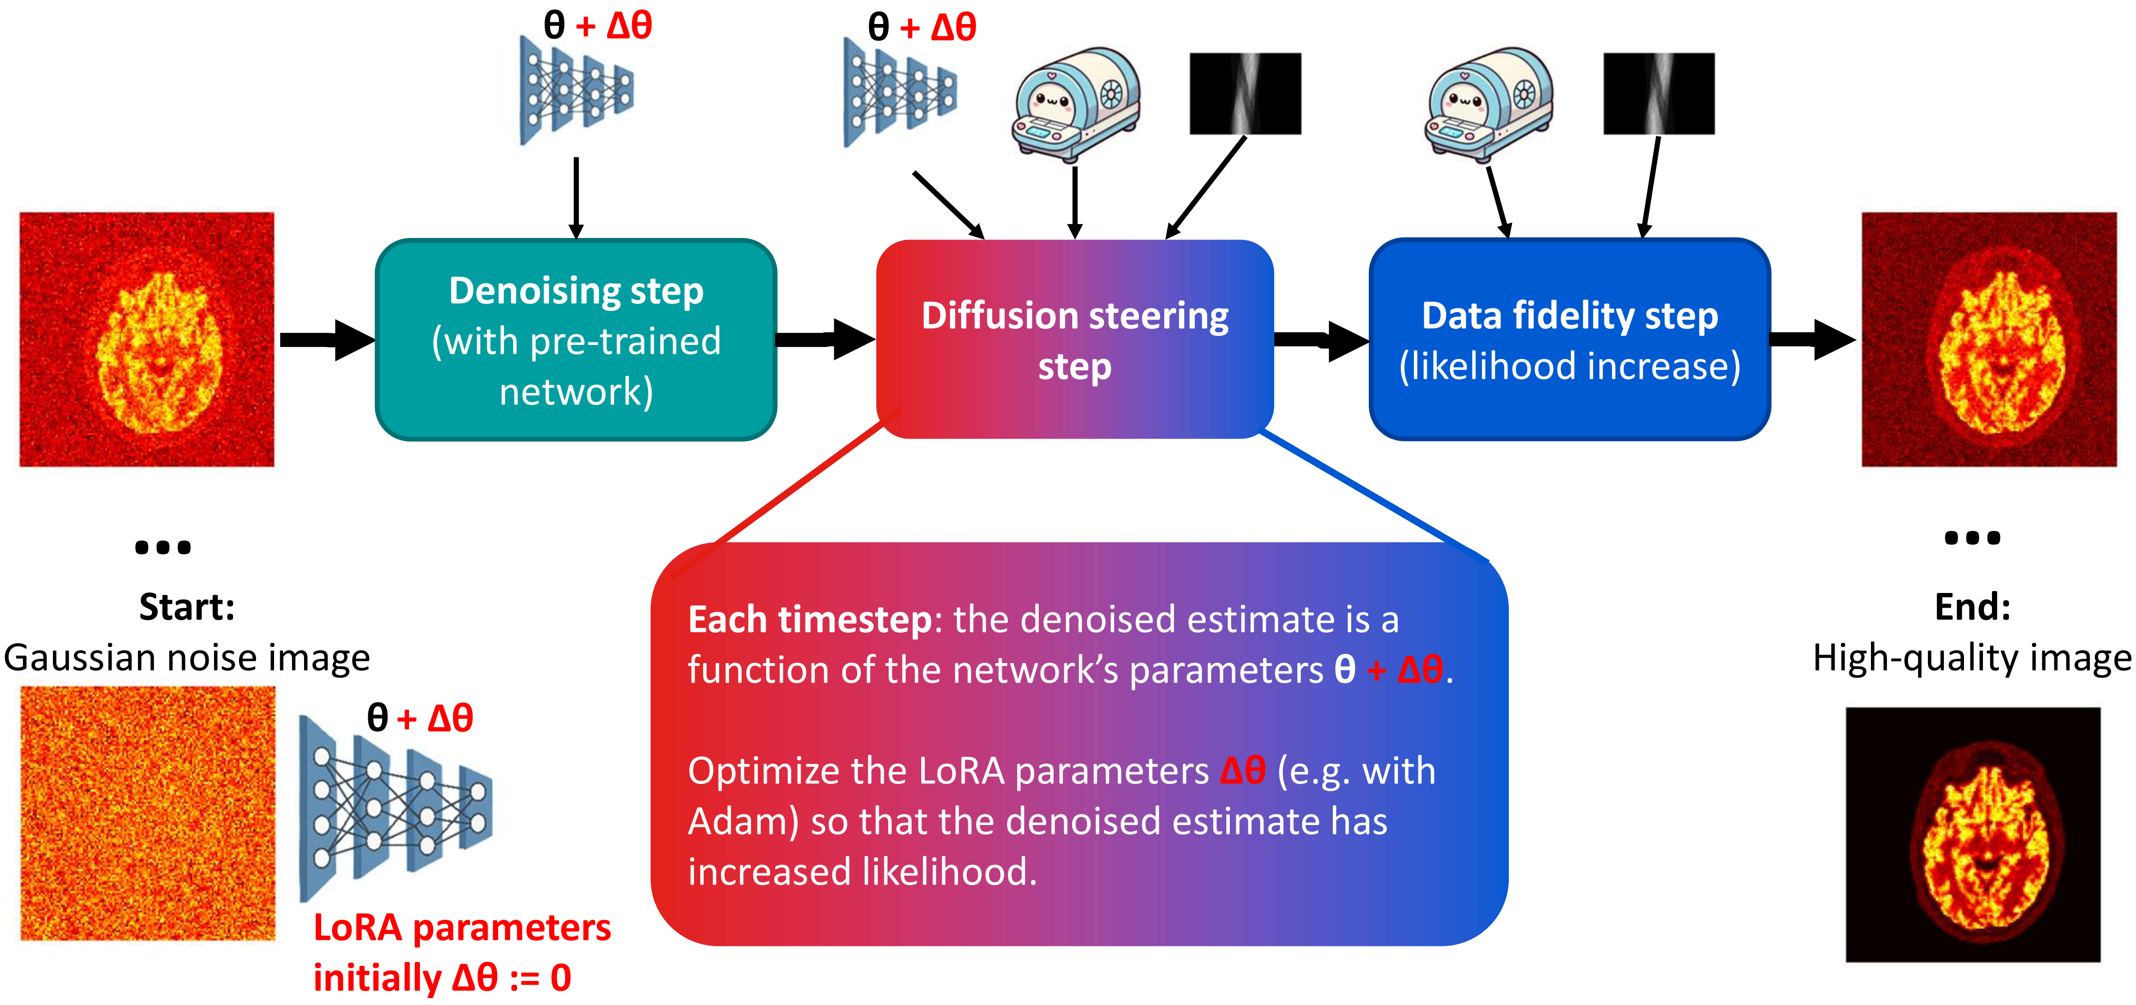

Results

Domain adaptation results

Steerable conditional diffusion does not compensate for noise spikes

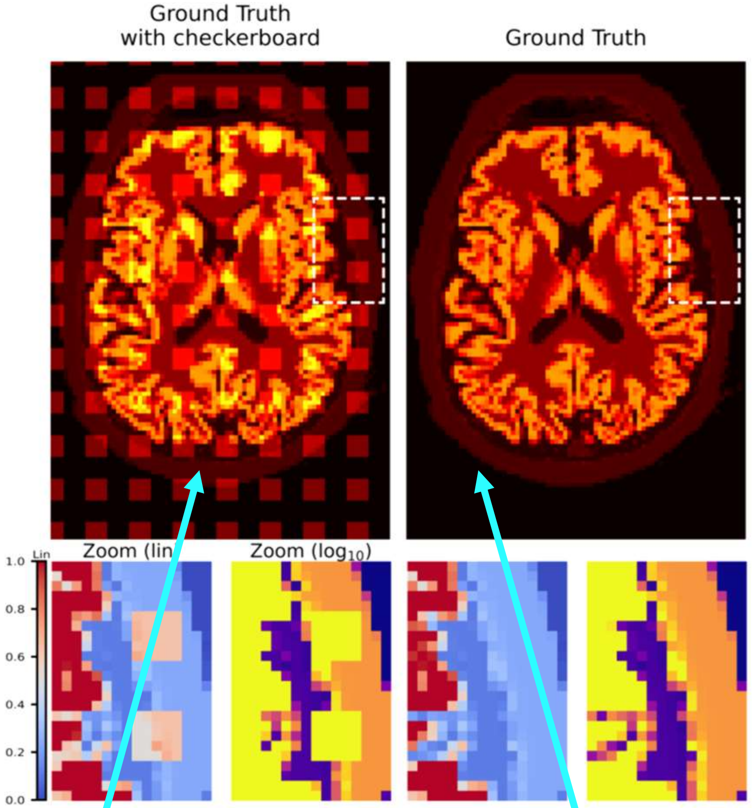

Train on images like this

Attempt to reconstruct this image

Steerable conditional diffusion adapts the prior to a new domain

Domain adaptation results

Steerable conditional diffusion does not compensate for noise spikes

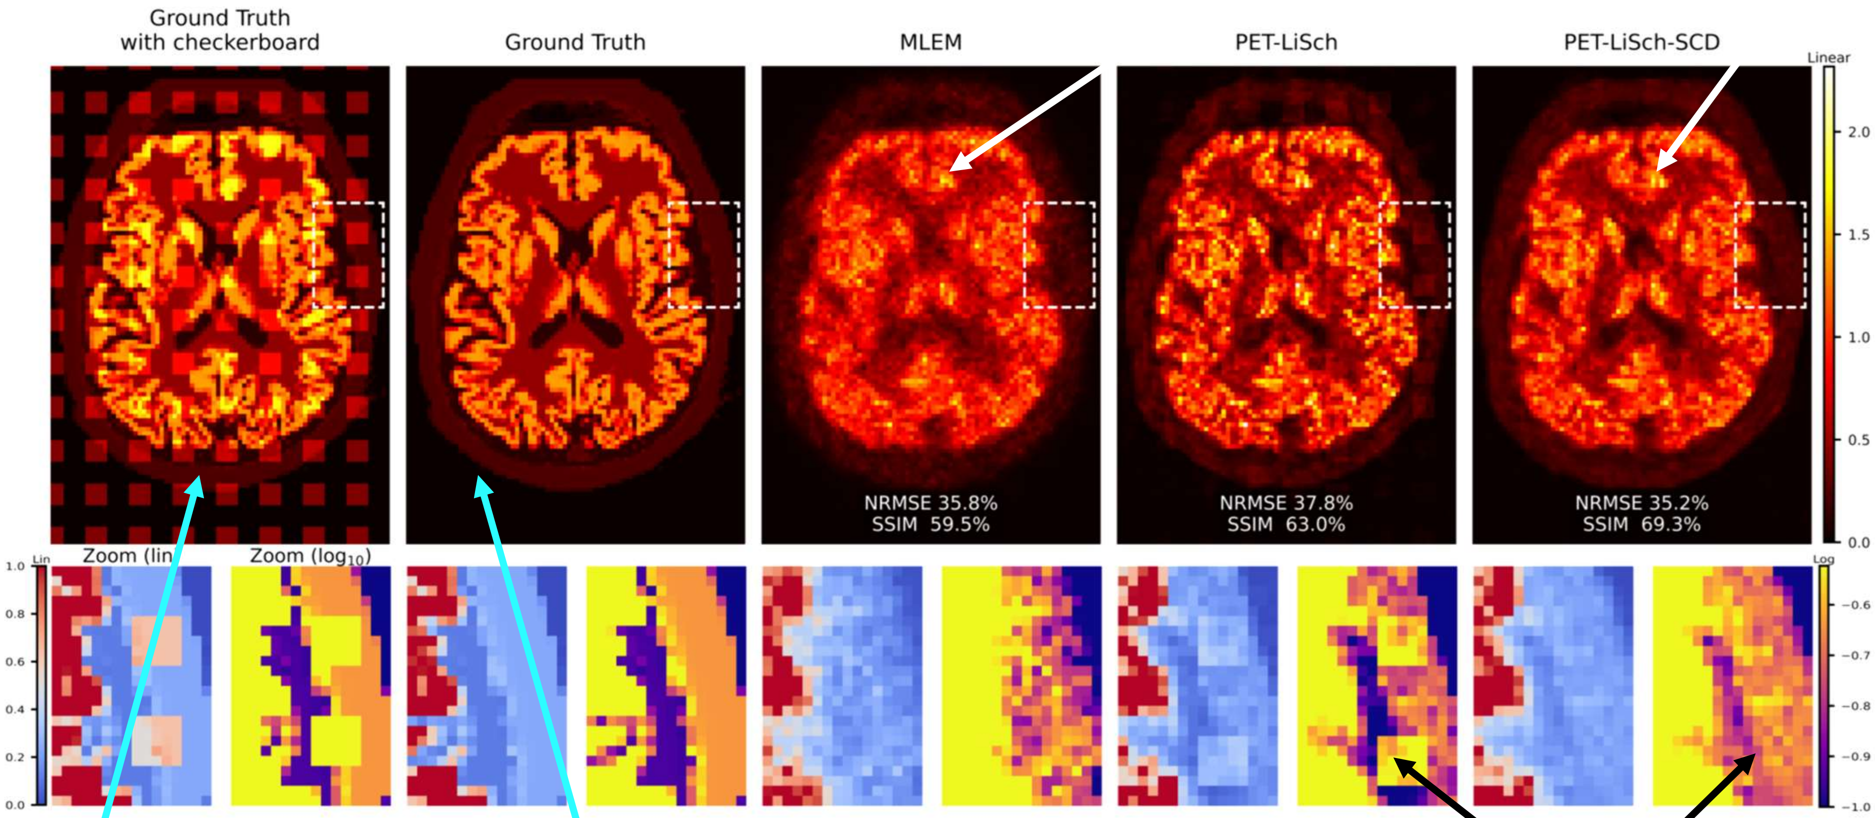

Train on images like this

Attempt to reconstruct this image

Steerable conditional diffusion adapts the prior to a new domain

Steerable Conditional Diffusion: Summary

- **Enables** some domain adaptation, where the training data has predictable structure (perhaps from a different scanner)

- More work needed to characterise the limits of steerability

Can we upgrade the quality of our training data?

Personalised Diffusion Models for PET Image Reconstruction

G. Webber et al. IEEE TRPMS 2025

George Webber[1], Alexander Hammers[2], Andrew P King[1], Andrew J Reader[1]

1. School of Biomedical Engineering & Imaging Sciences, King's College London, UK
2. Guy's and St Thomas' PET Centre & King's College London, UK

EPSRC Centre for Doctoral Training

Smart Medical Imaging

UKRI

GSK

KING'S College LONDON

ATLASES

Atlas Construction for Dynamic (4D) PET Using Diffeomorphic Transformations

Marie Bieth[1], Hervé Lombaert[1], Andrew J. Reader[2], and Kaleem Siddiqi[1]

[1] School of Computer Science and Centre for Intelligent Machines, McGill University, Canada
[2] Montreal Neurological Institute, McGill University, Canada

Abstract. A novel dynamic (4D) PET to PET image registration procedure is proposed and applied to multiple PET scans acquired with the high resolution research tomograph (HRRT), the highest resolution human brain PET scanner available in the world. By extending the recent diffeomorphic log-demons (DLD) method and applying it to multiple dynamic [^{11}C]raclopride scans from the HRRT, an important step towards construction of a PET atlas of unprecedented quality for [^{11}C]raclopride imaging of the human brain has been achieved. Accounting for the temporal dimension in PET data improves registration accuracy when compared to registration of 3D to 3D time-averaged PET images. The DLD approach was chosen for its ease in providing both an intensity and shape template, through iterative sequential pair-wise registrations with fast convergence. The proposed method is applicable to any PET radiotracer, providing 4D atlases with useful applications in high accuracy PET data simulations and automated PET image analysis.

1 Introduction

Medical image registration methods are necessary in a variety of clinical and research studies, whether it be aligning data between different subjects acquired with the same imaging modality (multi-subject single modality), or aligning data obtained from different modalities for the same subject (multi-modality single subject). The latter case is the most frequently used when dual-modality imaging is not available (e.g. to identify anatomical regions in PET images), and the challenge is to relate functional and structural images, e.g., PET to CT or PET to MR, as described in [6]. Such multi-modal registration is often limited to different images of the same subject. In the brain therefore, a rigid transformation can often be used, particularly when the acquisitions are close in time. However, the case of inter-subject single-modality registration requires more complex non-rigid methods to be deployed. Among these, Collins *et al* [3] develop a method for MR brain data that is now part of the MINC software suite. The Hammer algorithm [11] and the LDDMM framework [2] are also very popular for 3D to 3D non-rigid registration.

The problem of inter-subject PET to PET registration, however, is relatively unexplored in the medical imaging community. As an early example, the approach by Alpert *et al.* [1] and Eberl *et al.* uses only 6 parameters and cannot

K. Mori et al. (Eds.): MICCAI 2013, Part II, LNCS 8150, pp. 35–42, 2013.

MICCAI 2013

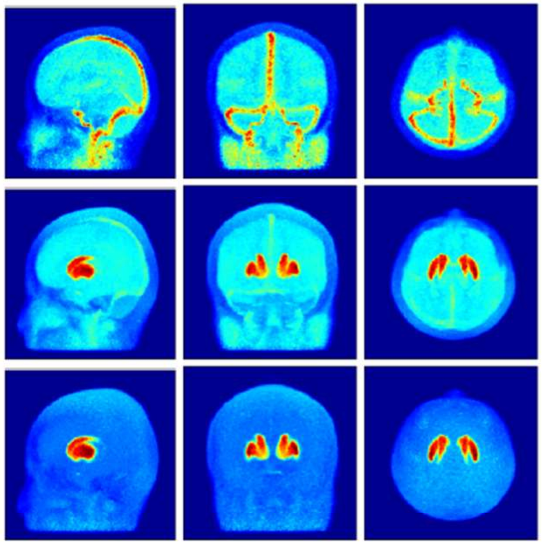

Fig. 2. Coronal, sagital and transverse maximum intensity projections of the 4D [^{11}C]Raclopride atlas. *First Row:* frame 3 in the temporal sequence. *Second Row:* frame 12 in the temporal sequence. *Third Row:* frame 22 in the temporal sequence.

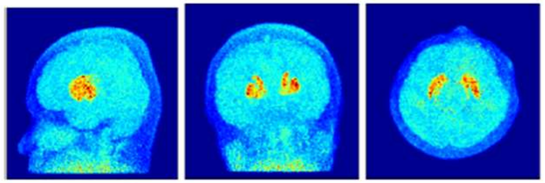

Fig. 3. Coronal, sagital and tranverse maximum intensity projections of a single [^{11}C]Raclopride subject, frame 12 in the temporal sequence. This is included for visual comparison with the template in Fig. 2.

4 Discussion

We have developed a new method for inter-subject dynamic (4D) PET image registration, based on an extension of the recent DLD method. Our method

outperforms two 3D registration methods we have compared it against in terms of intensity difference. It also appears to be more resistant to local minima. By applying it initially to 15 dynamic [^{11}C]raclopride scans from the HRRT, which is the highest resolution human brain PET scanner available in the world, we have taken an important step towards constructing a PET atlas of unprecedented quality for [^{11}C]raclopride imaging of the human brain. The DLD approach was chosen for its ease in providing both an intensity and shape-based template. The proposed method is in principle applicable to any PET radiotracer, providing 4D atlases which will find useful application in high accuracy PET data Monte Carlo simulations as well as for automated PET image analysis. Furthermore, when used with appropriate care, such atlases could provide spatiotemporal priors for 3D and fully 4D PET image reconstruction.

Acknowledgments. The authors would like to thank Alain Dagher for helpful discussions and data and NSERC and FQRNT for research funding.

References

1. Alpert, N., Berdichevsky, D., Levin, Z., Morris, E., Fischman, A.J.: Improved methods for image registration. NeuroImage 3 (1996)
2. Beg, M., Miller, M., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. IJCV 61 (2005)
3. Collins, D., Neelin, P., Peters, T., Evans, A.: Automatic 3D intersubject registration of MR volumetric data in standardized talairach space. Journal of Computer

"…when used with appropriate care, such atlases could provide spatiotemporal priors for 3D and fully 4D PET image reconstruction."

Learning prior information with a diffusion model

Forward diffusion process (random)

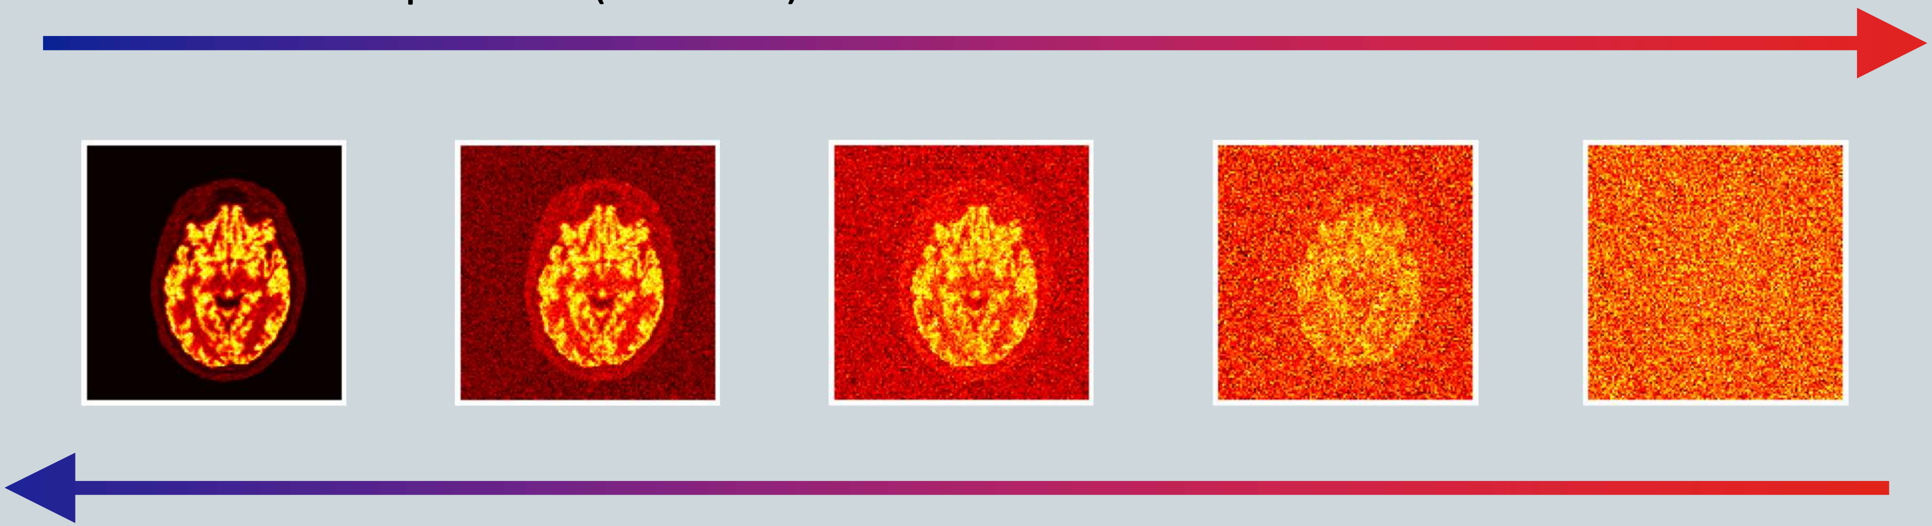

Reverse generative process (learnt)

- ✓ Scanner-agnostic
- ✓ Easy to train
- ✓ Generalises well from few datasets

Conditioning a trained diffusion model


JOURNAL ARTICLE

Diffusion models for medical image reconstruction

George Webber, MMathCompSci, Andrew J Reader, PhD

BJR|Artificial Intelligence, Volume 1, Issue 1, January 2024, ubae013, https://doi.org/10.1093/bjrai/ubae013
Published: 29 August 2024 **Article history ▾**

Forward diffusion process (random)

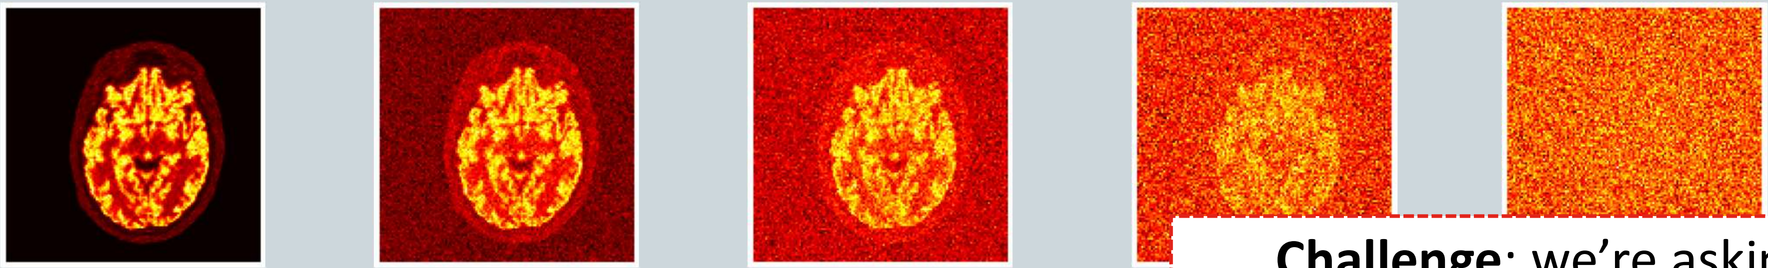

PET physics

\+

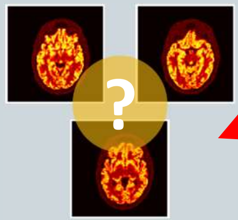

Pre-trained diffusion model

Challenge: we're asking our diffusion model to generalise to all unknown brain anatomies from <50 datasets.
Possible, but needlessly tricky!

Reverse generative process (learnt)

Personalised diffusion model

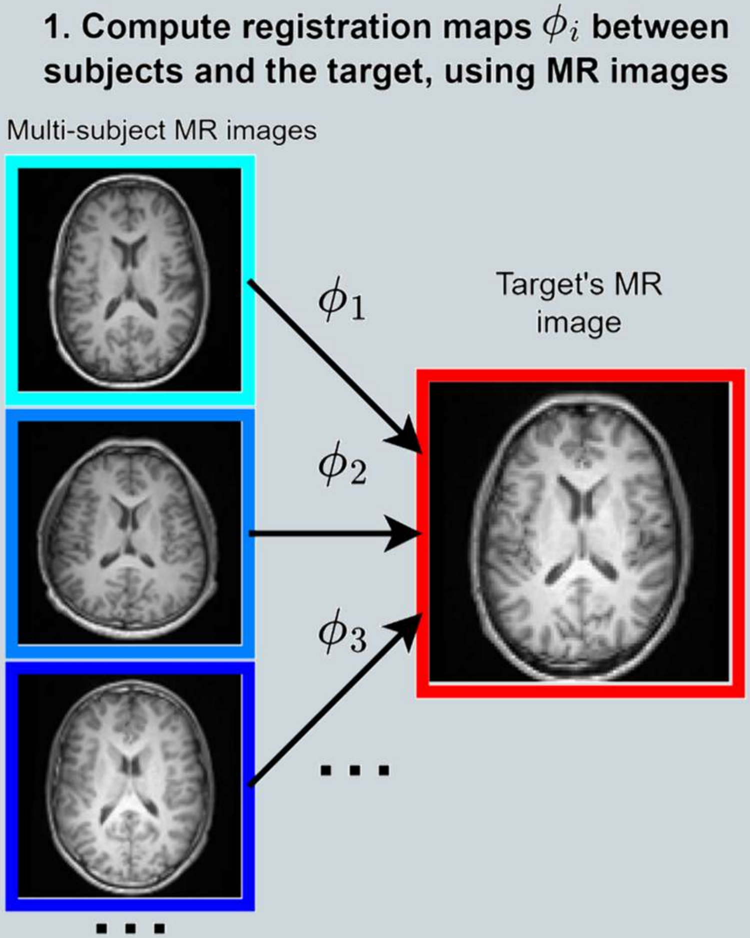

Affine +
VoxelMorph

Personalised diffusion model

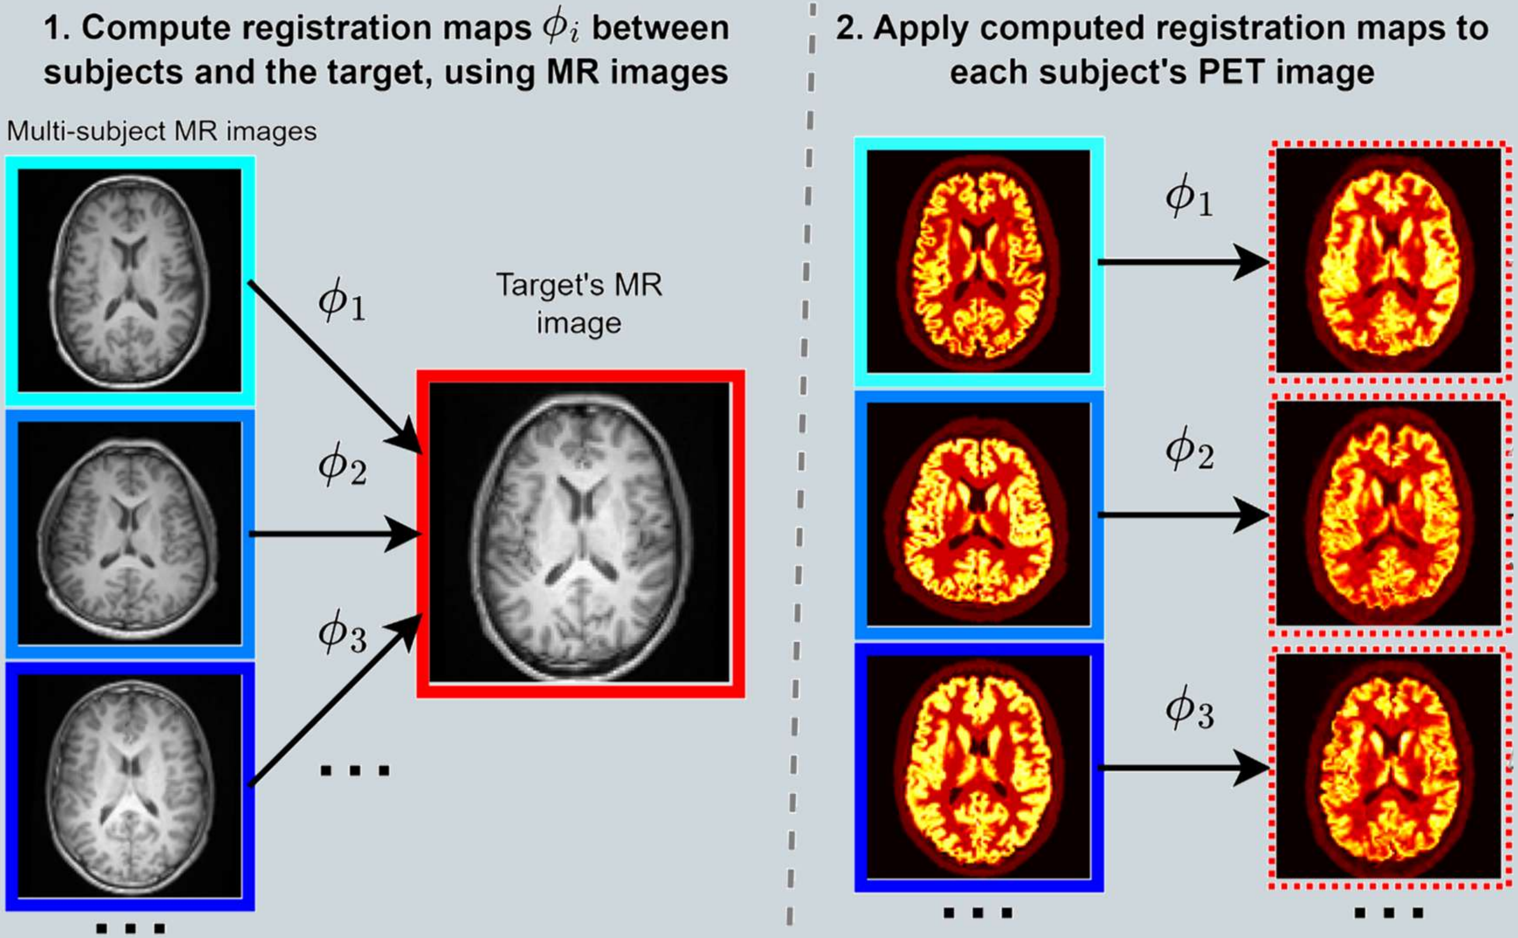

Can now train a patient-anatomy unique diffusion model

Personalised diffusion model

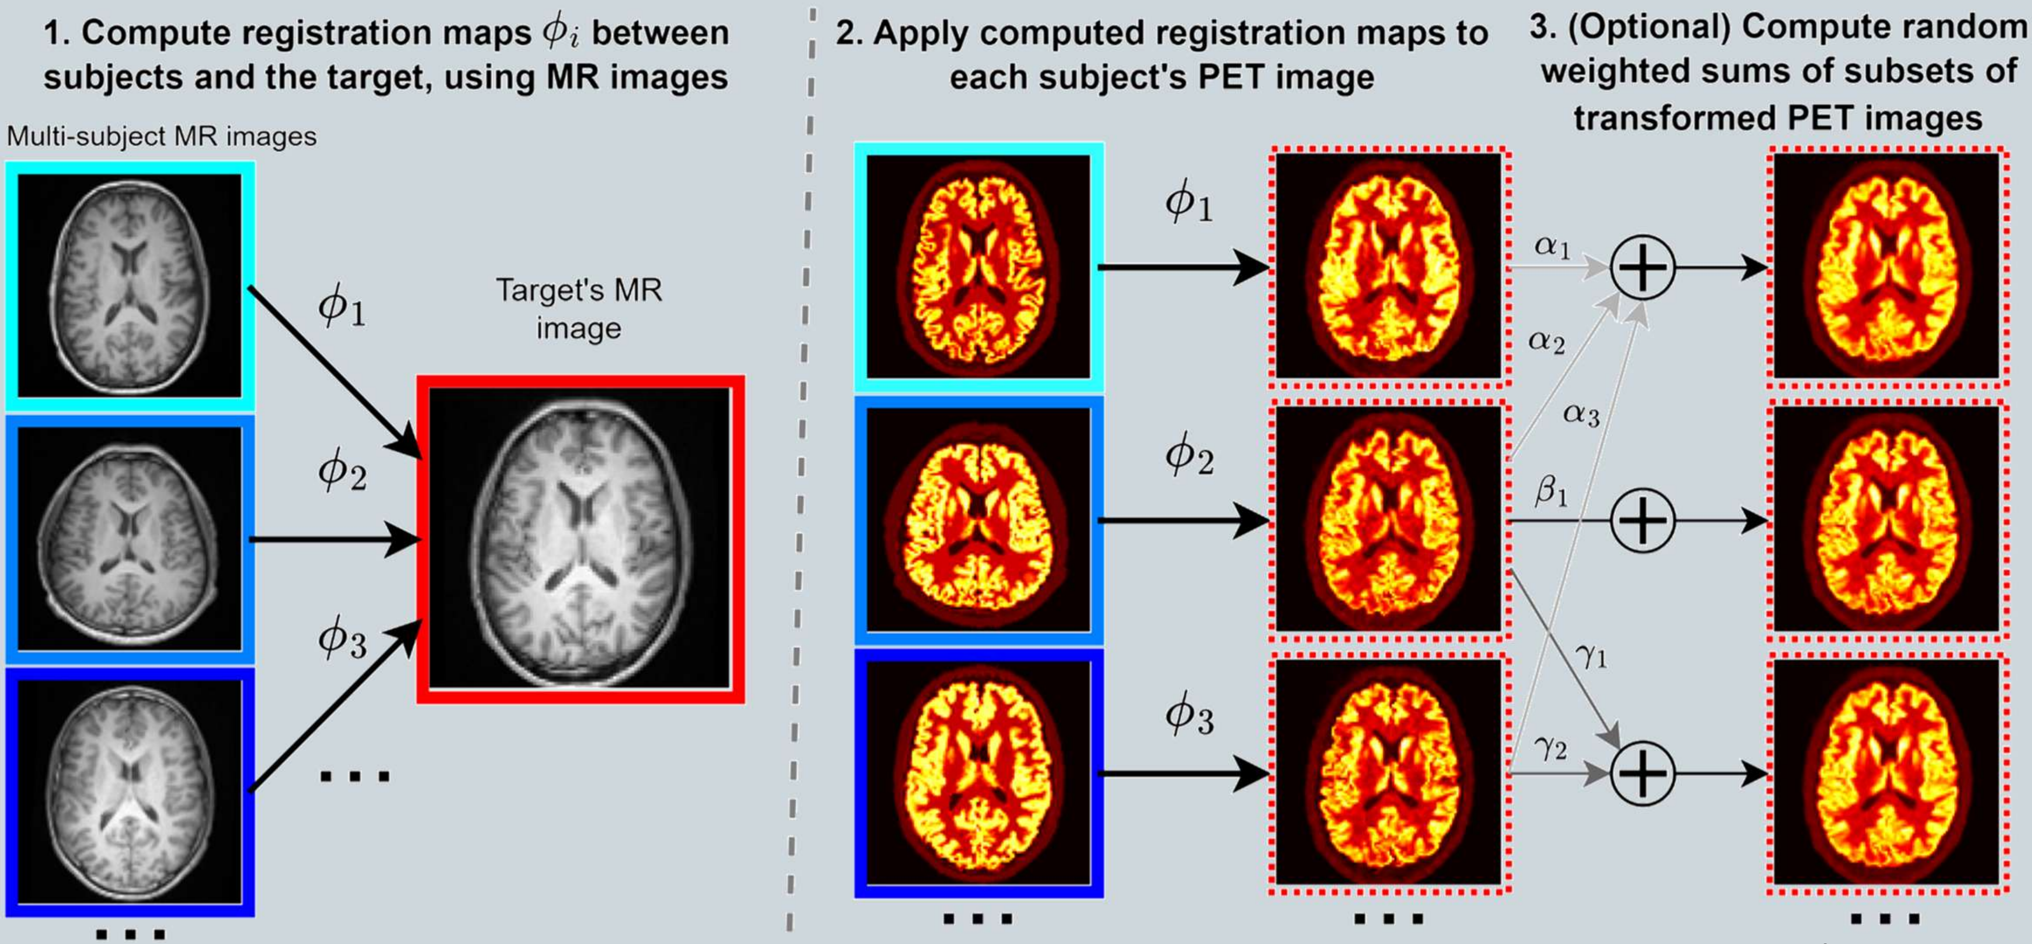

Richer SNR to train a diffusion model

Personalised diffusion model

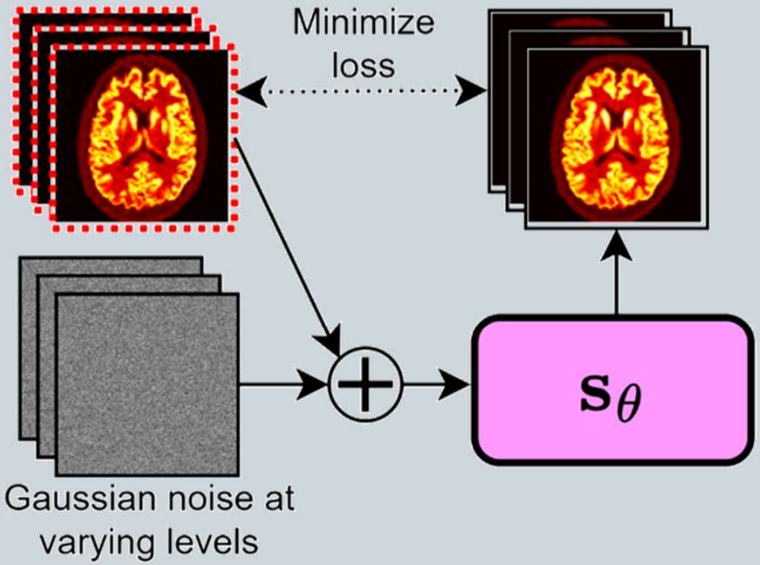

4. Train diffusion model on target-specific transformed PET images ("pseudo-PET")

Personalised diffusion model

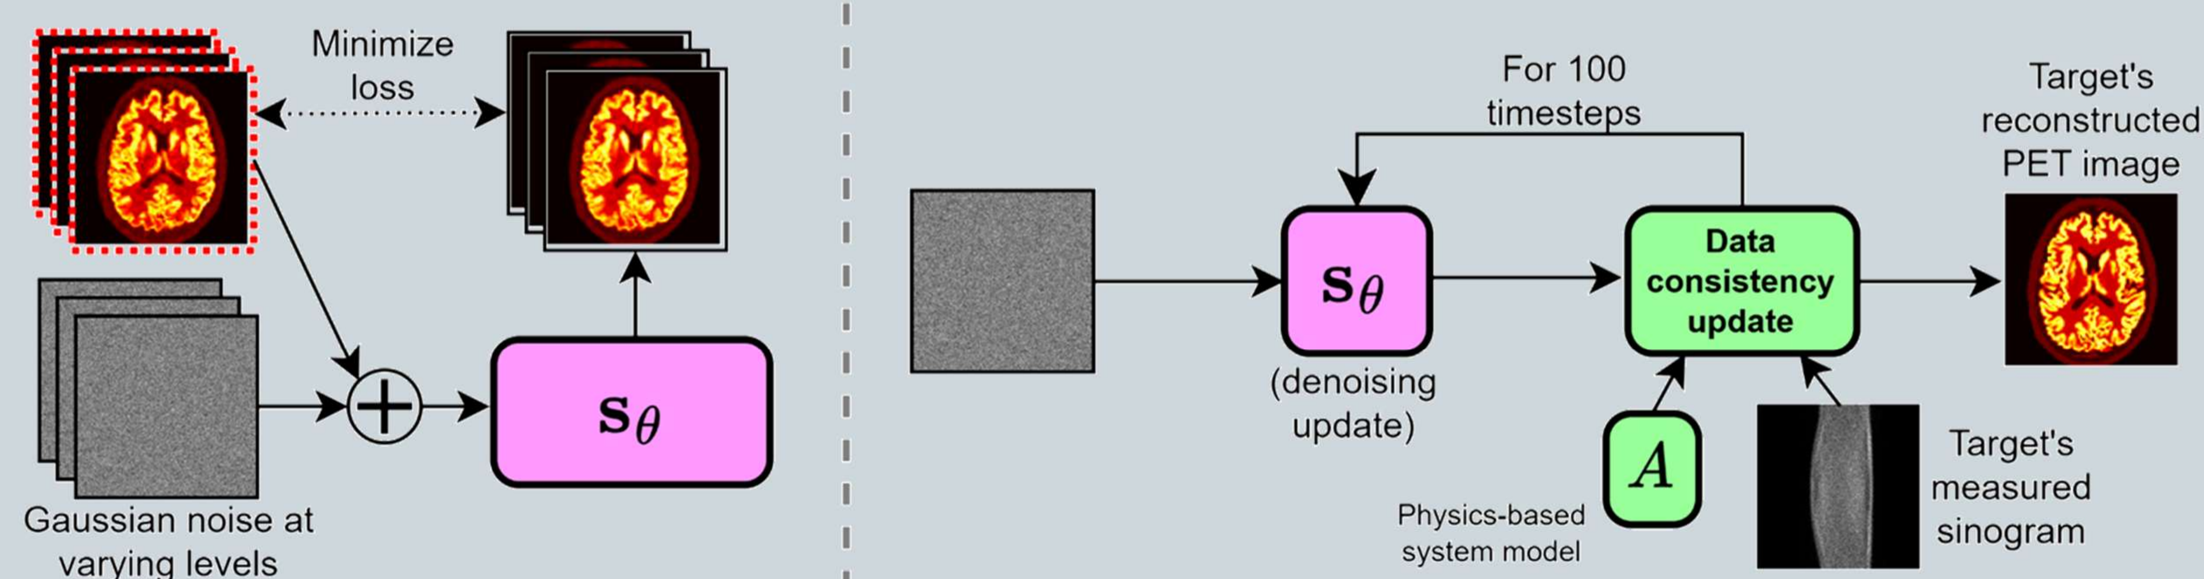

4. Train diffusion model on target-specific transformed PET images ("pseudo-PET")

5. Use PET-DDS algorithm with trained diffusion model to reconstruct subject's PET image from sinogram data

Results

Results – quantitative images (simulated 2D FDG case)

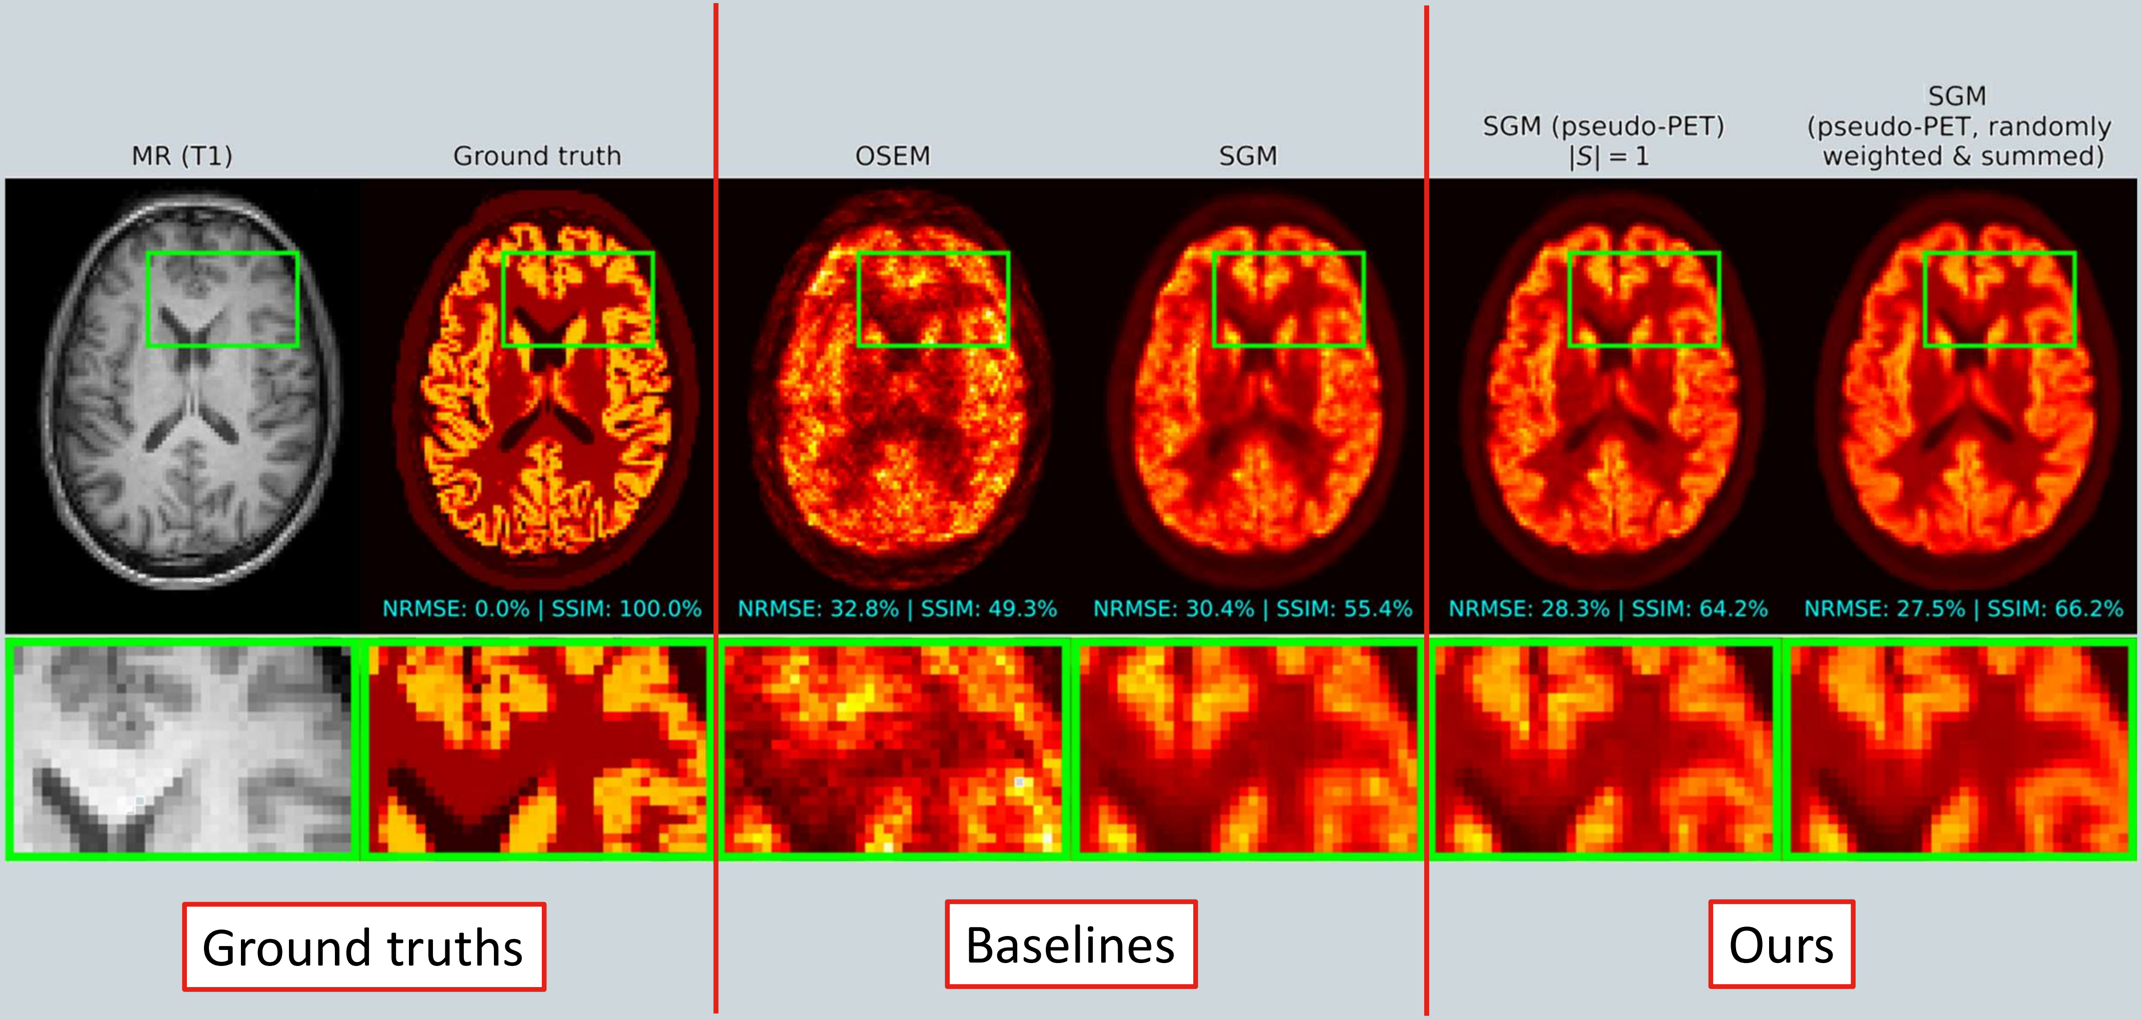

Results – reconstruction accuracy (simulated 2D FDG case)

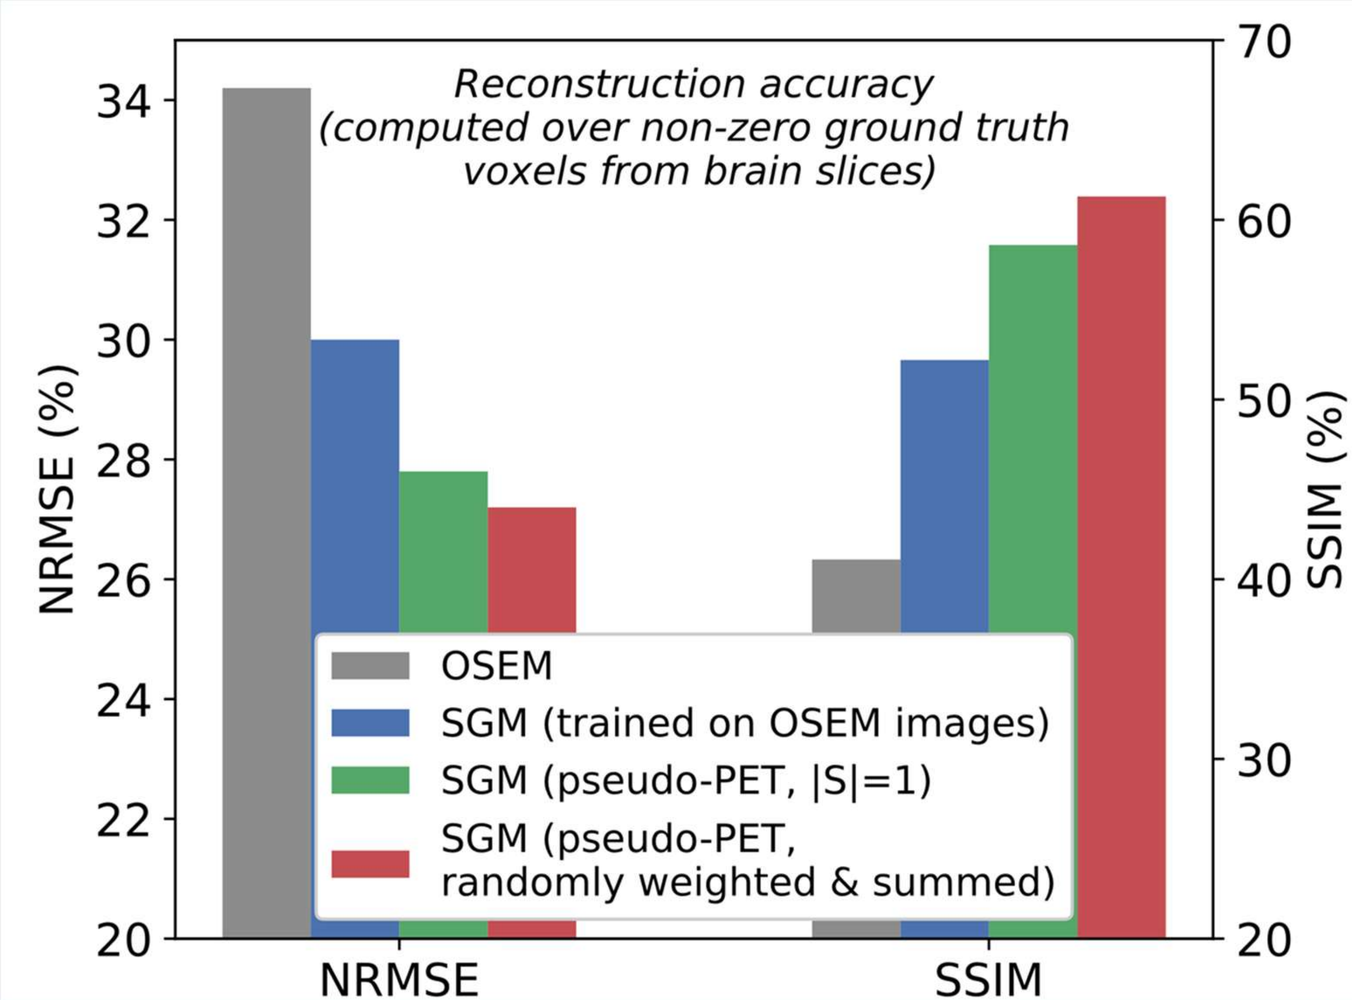

Results – real [^{18}F]DPA-714 images, 2D real-data reconstructions

MR (T1) | OSEM 3D | OSEM 2D | SGM 2D | SGM (pseudo-PET) 2D $|S| = 1$ | $|S| = 4$

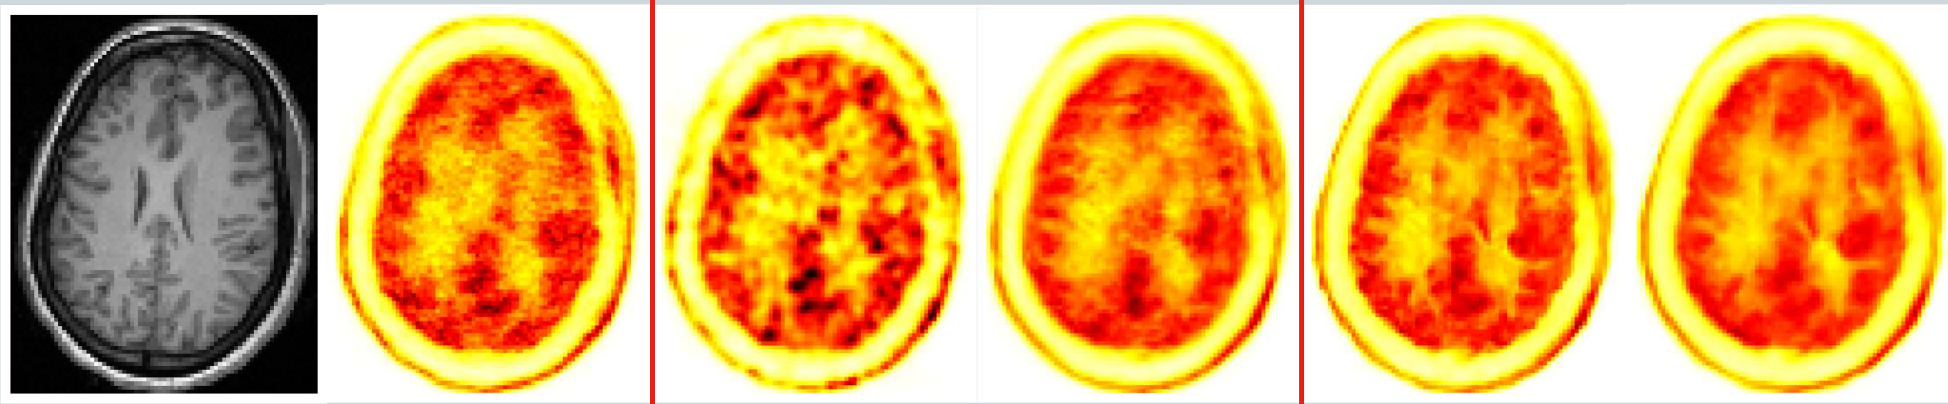

Ground truths

Baselines

Ours

Real data 3D FDG case: reconstructing low count data

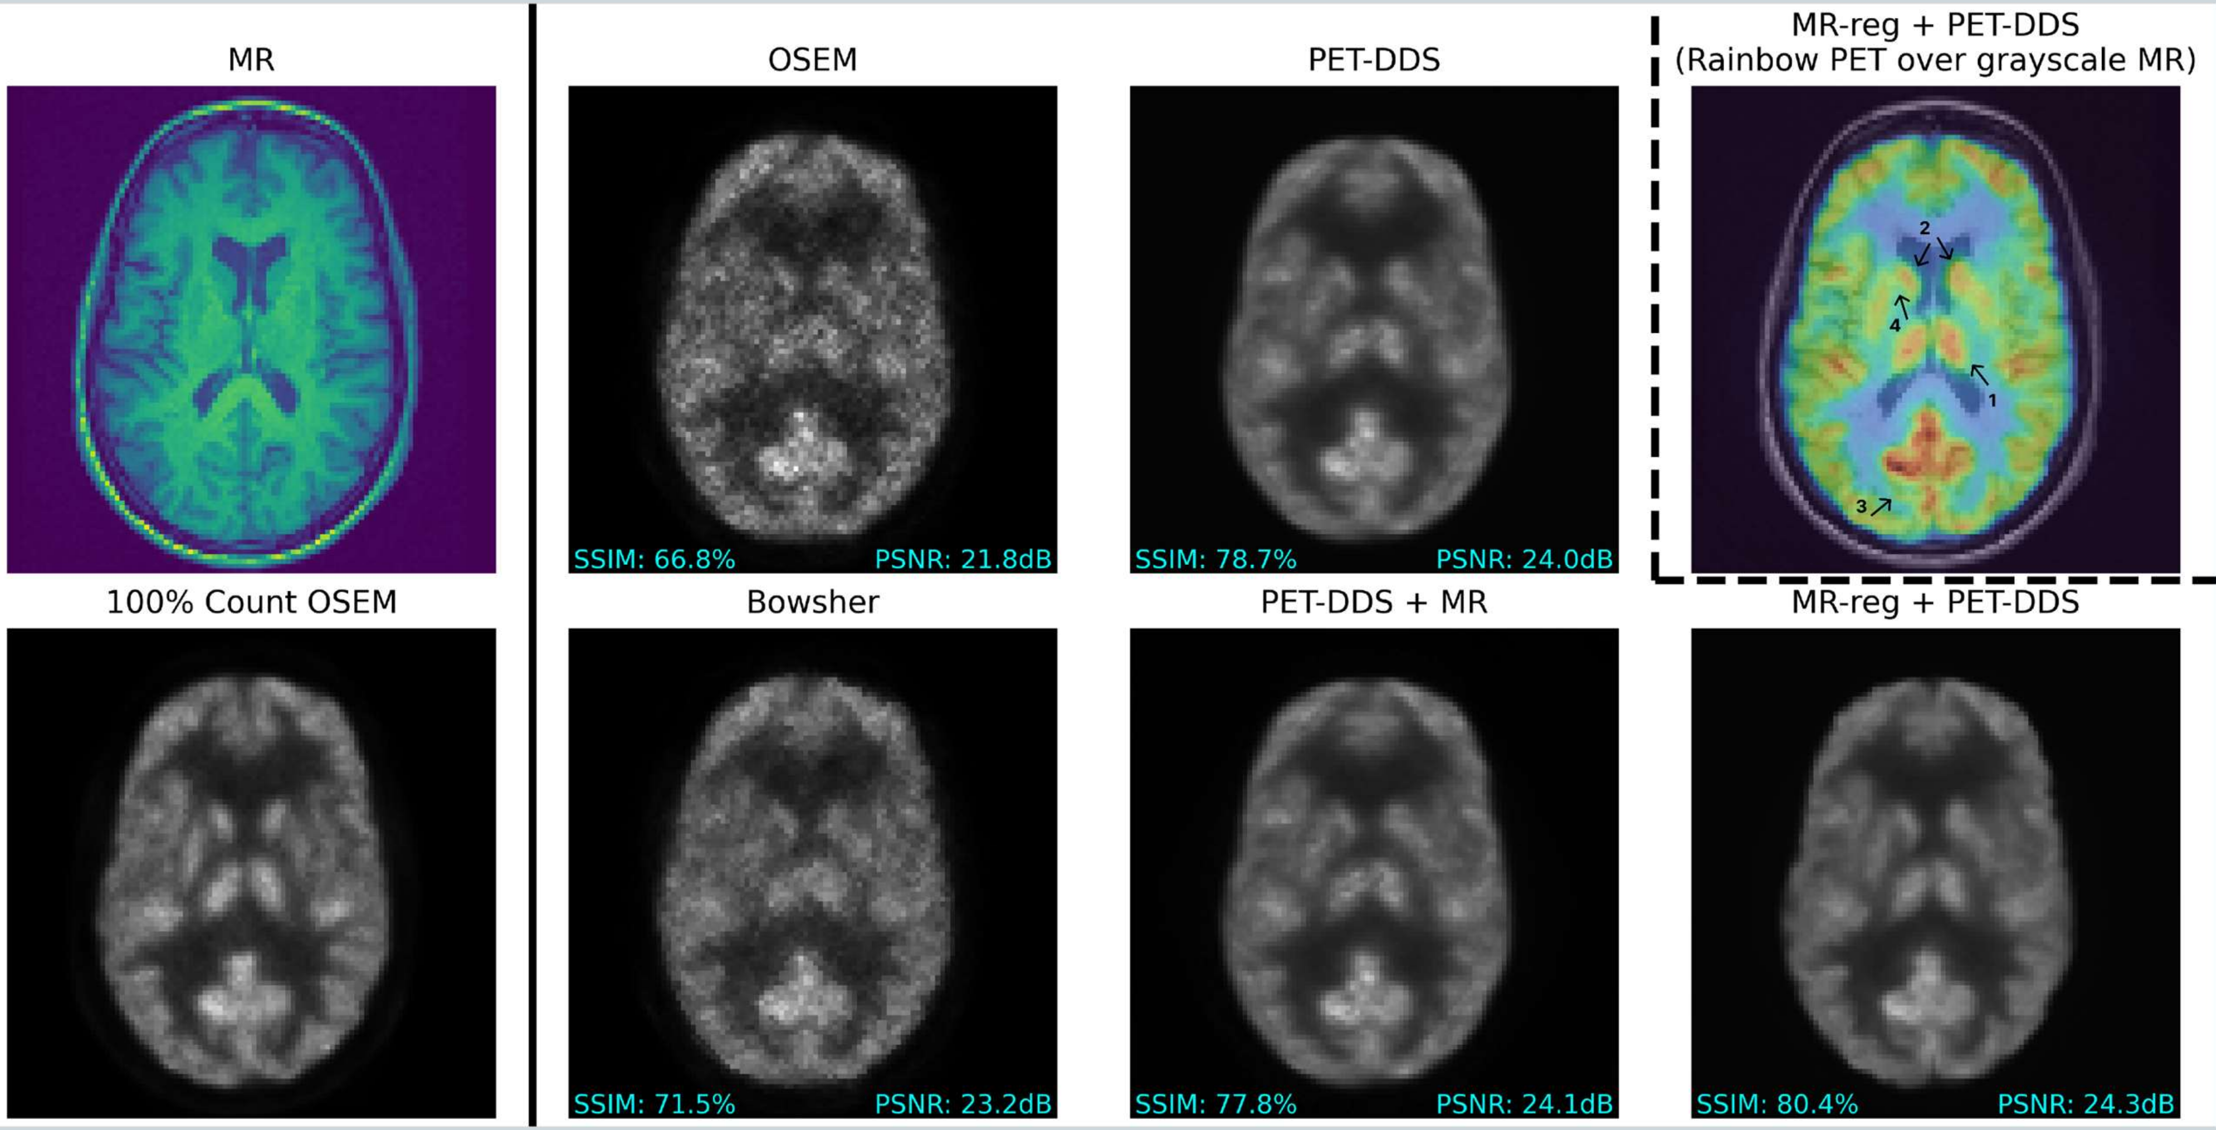

Pairing high-quality images with noisy measured data

Supervised Diffusion Models for PET Image Reconstruction

G. Webber et al. MICCAI 2025

George Webber[1], Alexander Hammers[2], Andrew P King[1], Andrew J Reader[1]

1. School of Biomedical Engineering & Imaging Sciences, King's College London, UK
2. Guy's and St Thomas' PET Centre & King's College London, UK

EPSRC Centre for Doctoral Training
Smart Medical Imaging
UKRI
GSK

KING'S College LONDON

Supervised diffusion models for PET reconstruction

- **Unsupervised diffusion model (DM) methods**
  - ✓ Can improve image quality relative to conventional MAP methods
  - DM does not learn how Poisson noise, system model, and iterative data-consistency steps affect the image
  - Limits image quality and quantitative accuracy

- **Conventional supervised methods**

  Trained end-to-end to map noisy measurements to higher-quality reference images
  - ✓ Excellent image quality and quantitative accuracy for specific acquisition protocols (e.g. low dose imaging)
  - But no model of probability distribution over images
  - Point estimates only, no principled uncertainty quantification
  - No generative mechanism to explore posterior variability

- **Supervised DM**
  - Train the DM with high-quality images paired with noisy PET measurement data, using the system model
  - ✓ High accuracy
  - ✓ Posterior sampling (uncertainty quantification)

Diffusion Model Training

Add Gaussian noise with varying σ to training images
Train a network to remove the noise

Pick
a training image

$$\mathbf{x} \ \& \ t \sim U[0,1]$$

with associated
noise level

$$\sigma = \beta(t).$$

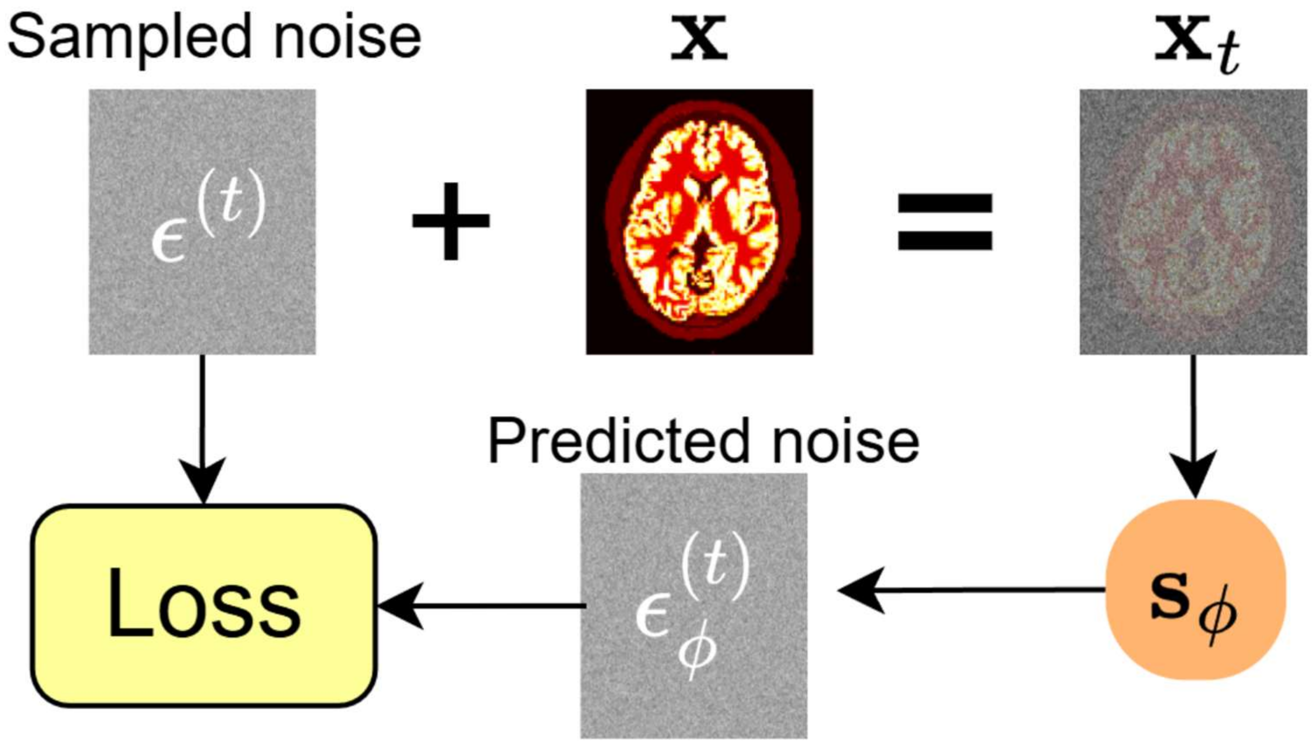

Generating Images with Diffusion Models

Begin with random noise. At each time step, predict and remove a small amount of noise, until a high-quality image remains

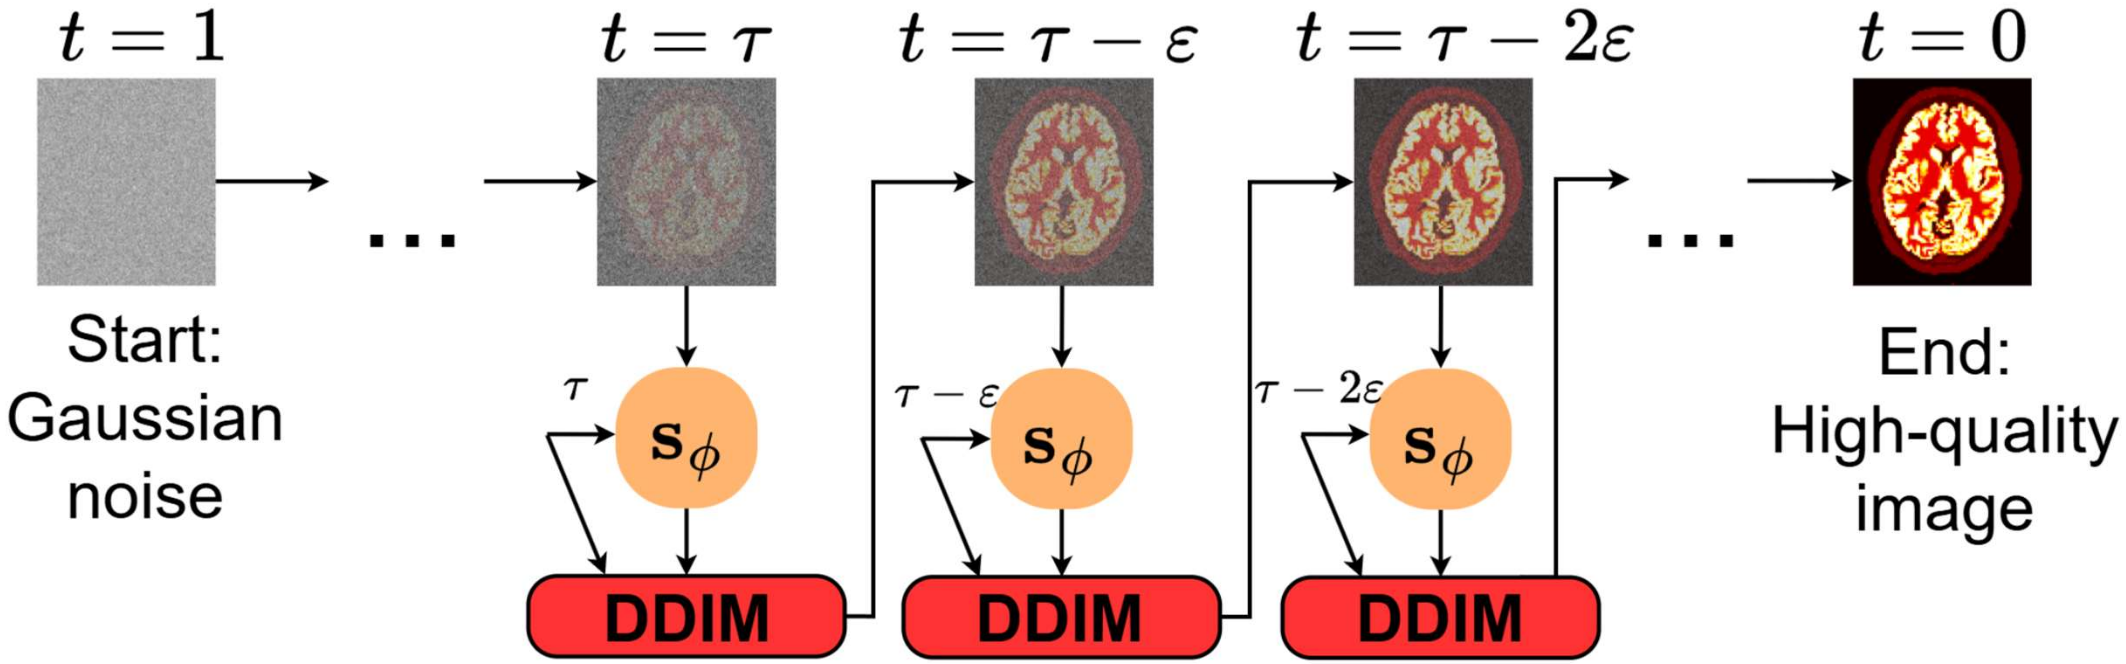

PET-DEFT: conditioning a DM on raw PET data

> Get current unconditioned high-quality estimate via a pre-trained unconditional DM

> Use that estimate to find the gradient of the Poisson log-likelihood (PLL) wrt noisy measured data

> Input i) the current iterate, ii) unconditional estimate and iii) gradient of the PLL to the conditional DM denoiser

> Apply a non-negativity operation, ensures the conditional score estimate results in a non-negative final image

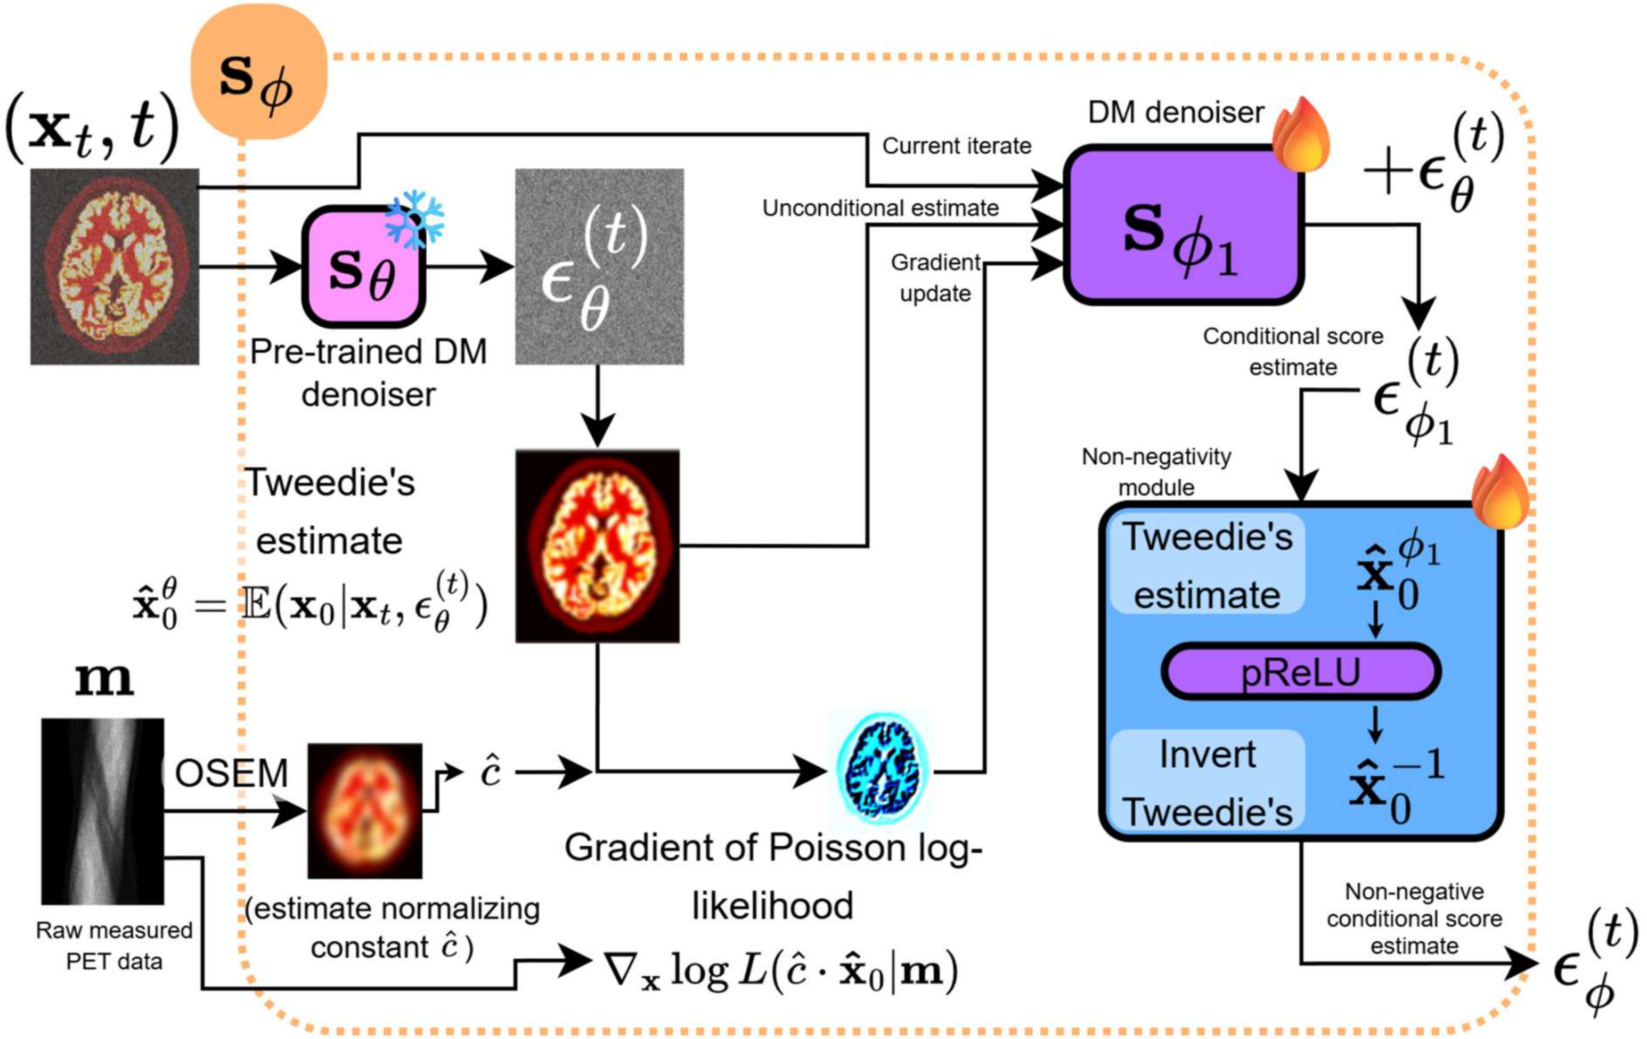

Denker et al. DEFT Advances in Neural Information Processing Systems 2024

Results

Results: Qualitative Comparison

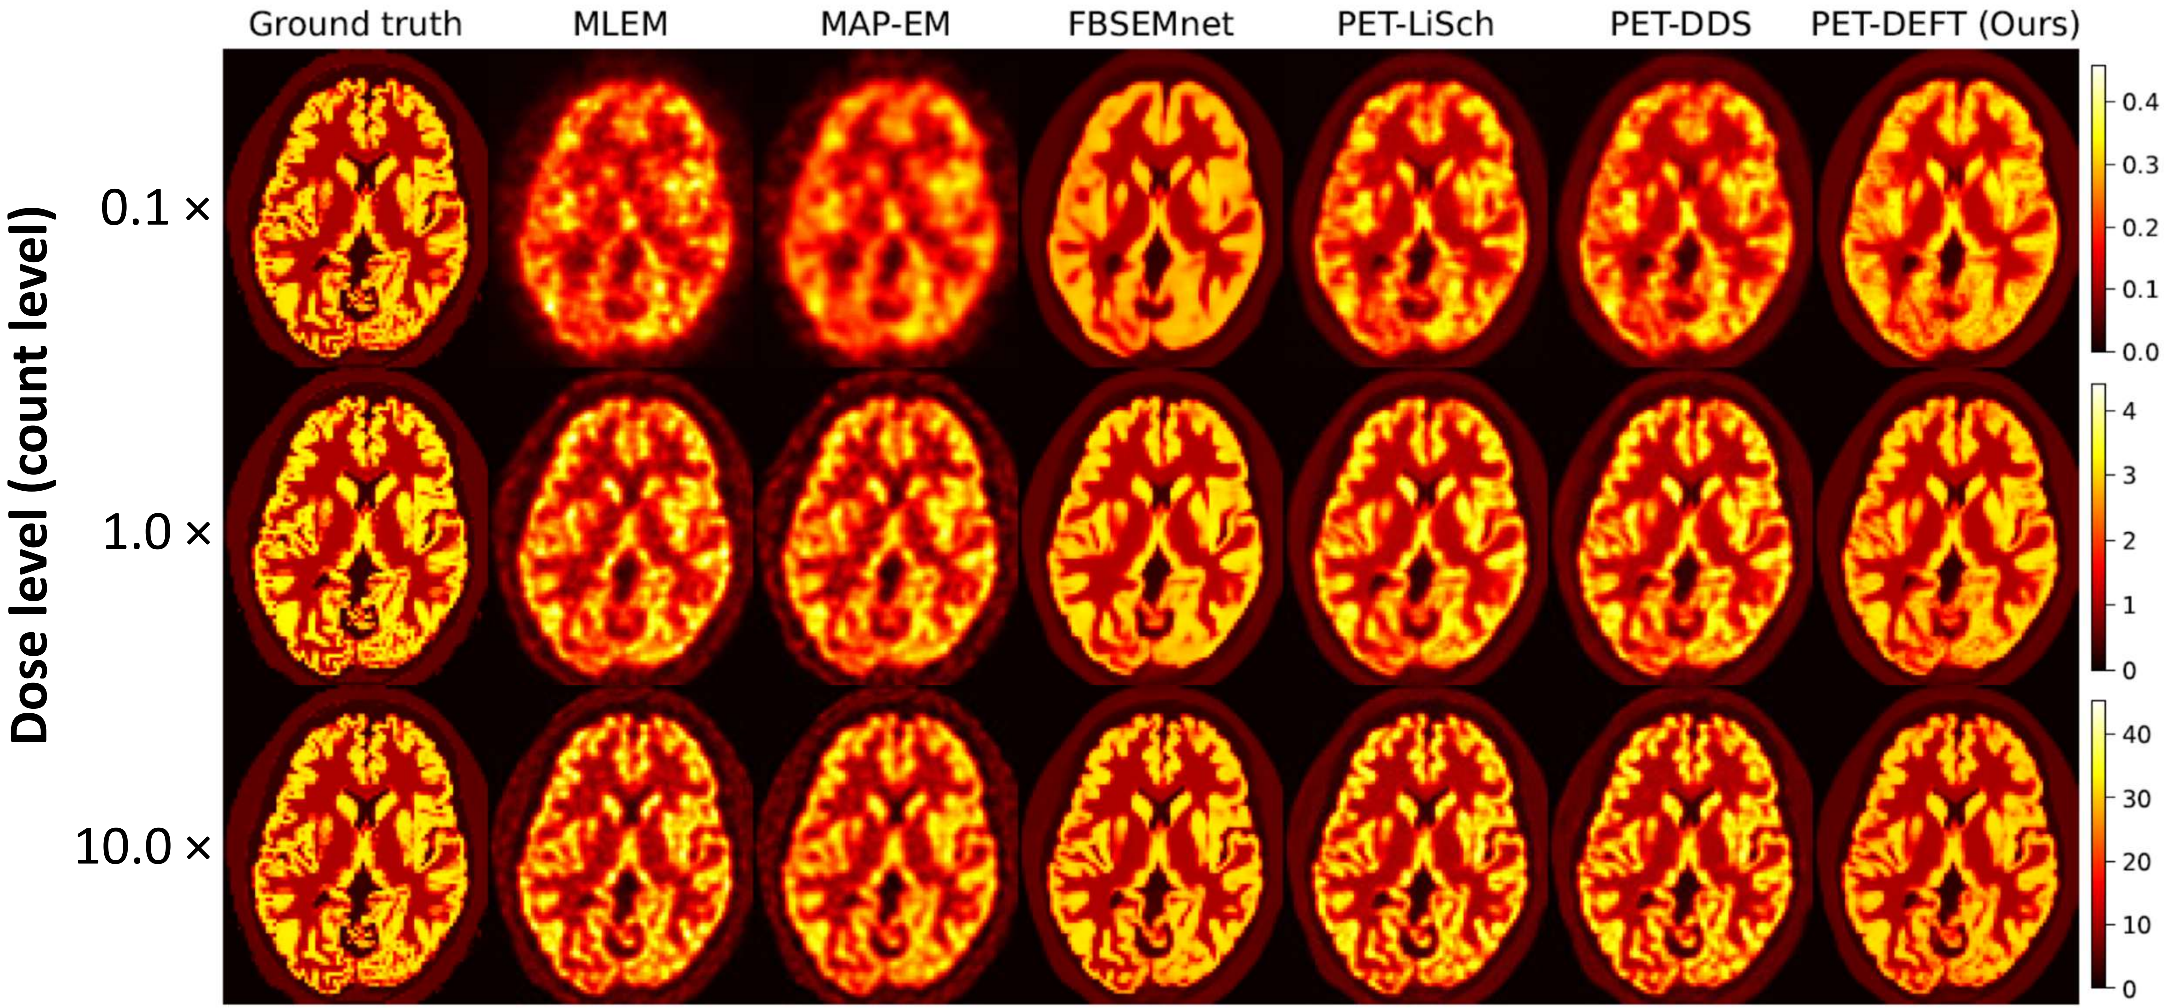

As dose level increases, all methods converge towards the ground truth at different rates

Results: NRMSE

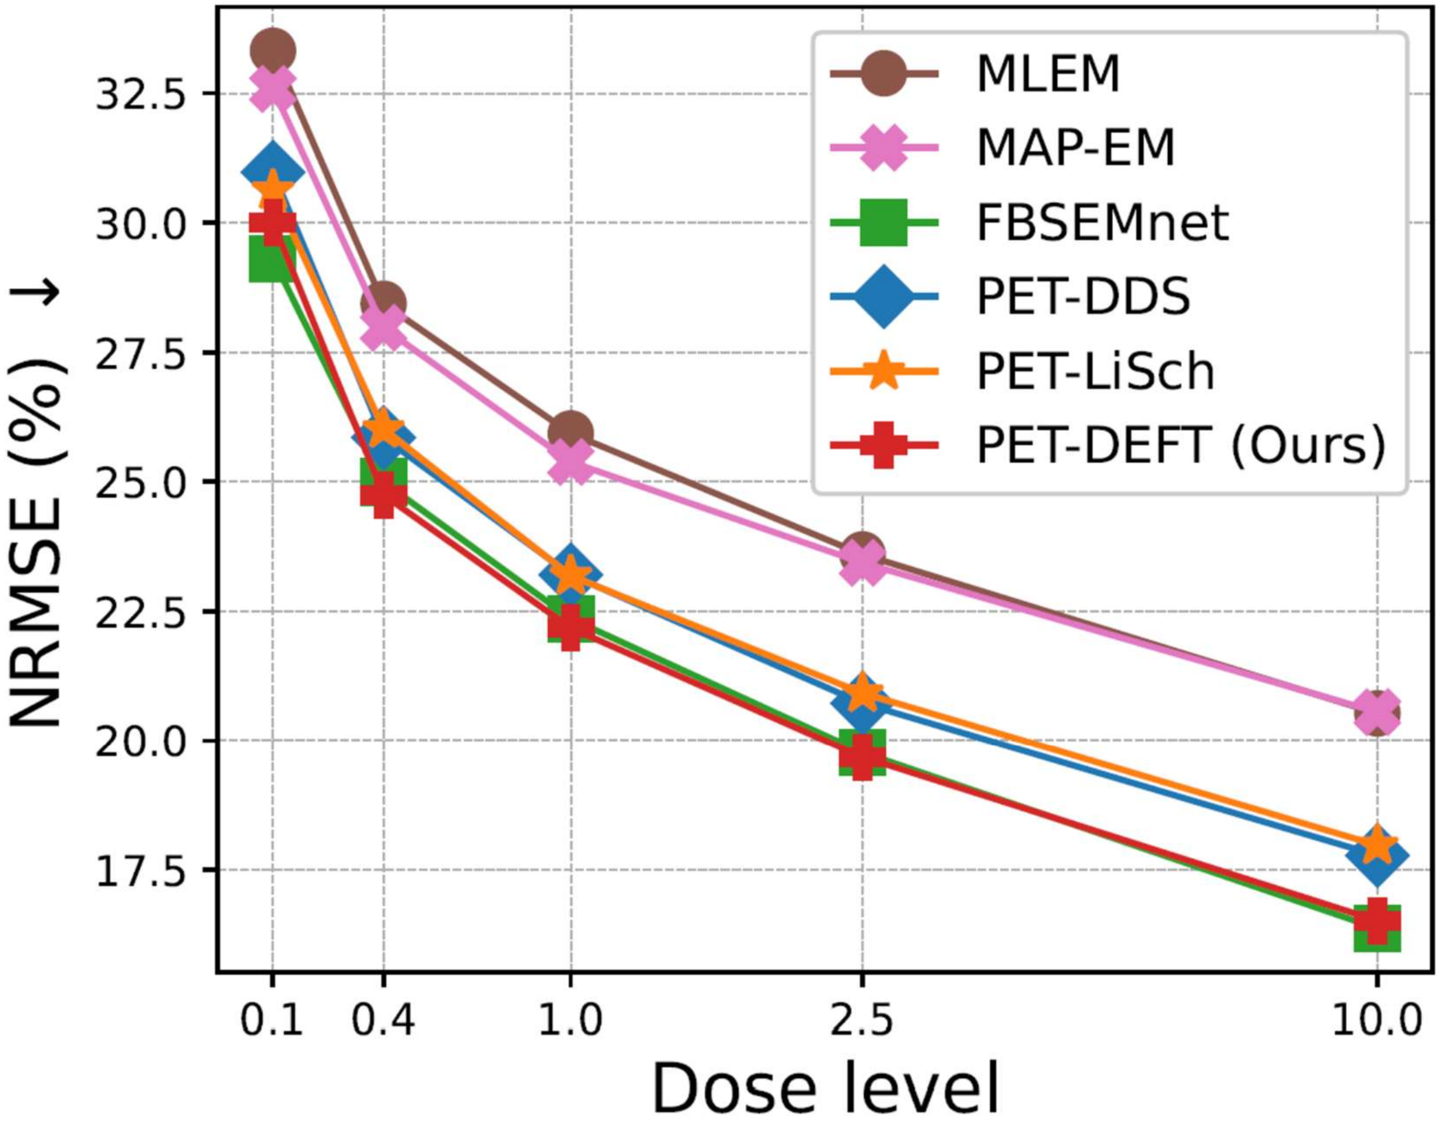

Supervised approaches PET-DEFT and FBSEMnet (i.e. those with access to raw data for training) yield the most accurate reconstructed Images

Unsupervised DM-based approaches PET-DDS and PET-LiSch perform worse

Results: SSIM

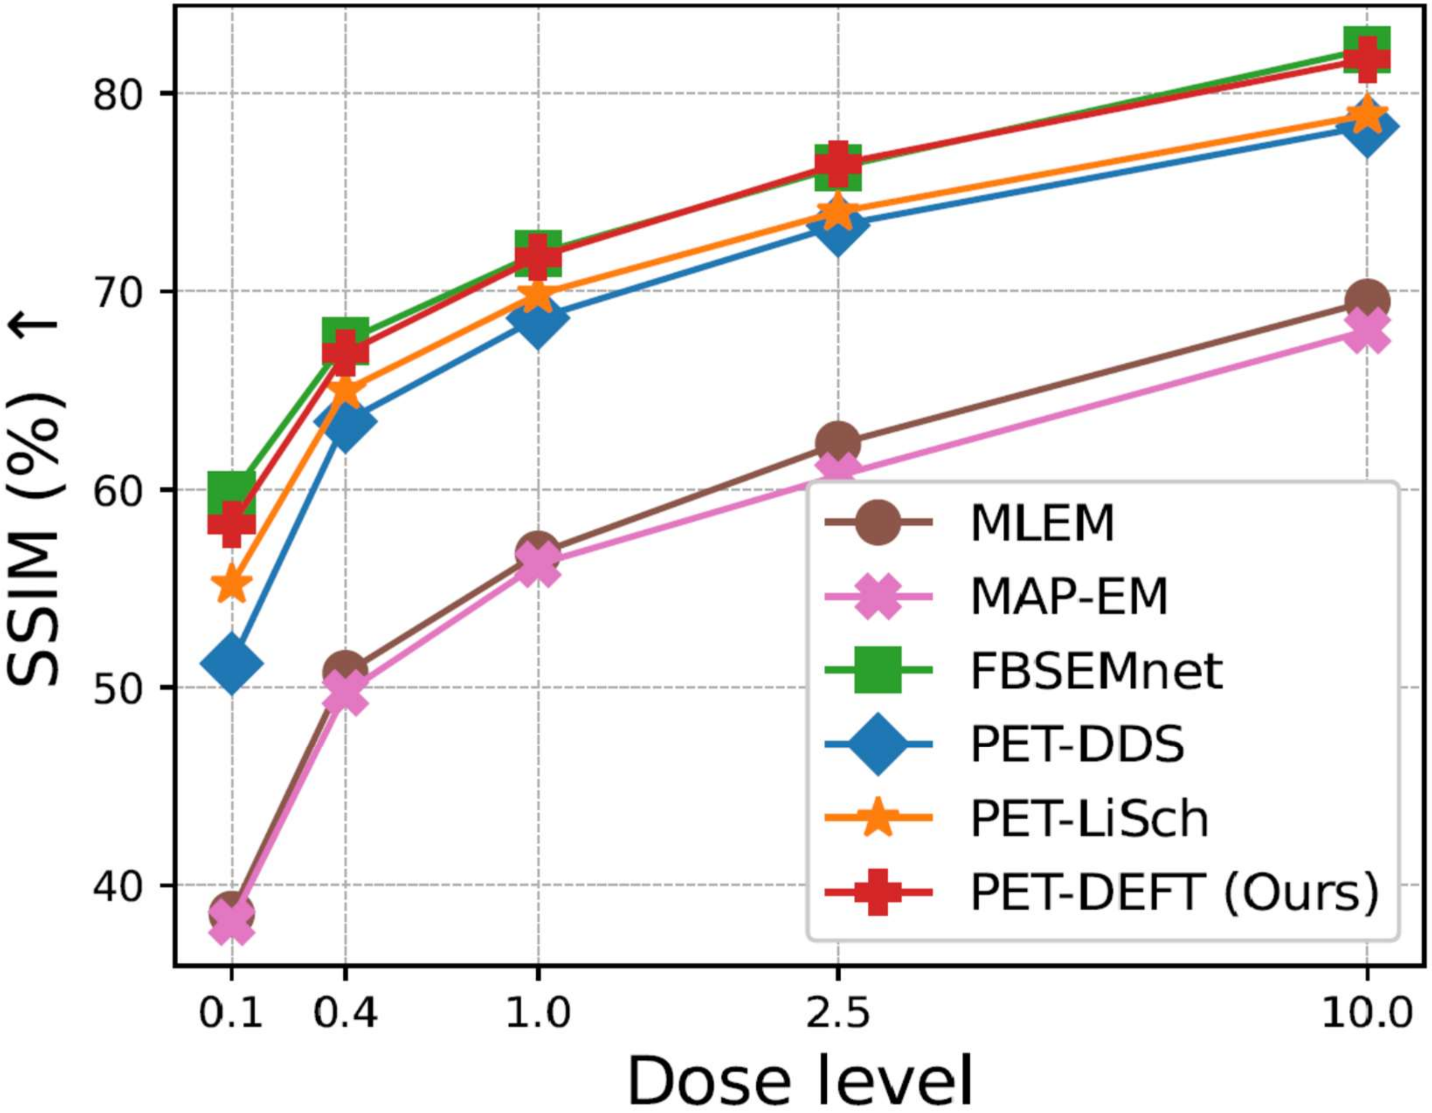

Supervised approaches PET-DEFT and FBSEMnet (i.e. those with access to raw data for training) yield the most accurate reconstructed Images

Unsupervised DM-based approaches PET-DDS and PET-LiSch perform worse

Results: Samples from the Posterior

Samples from PET-DEFT (supervised) visually resemble the ground truth manifold much more than samples from PET-LiSch (unsupervised)

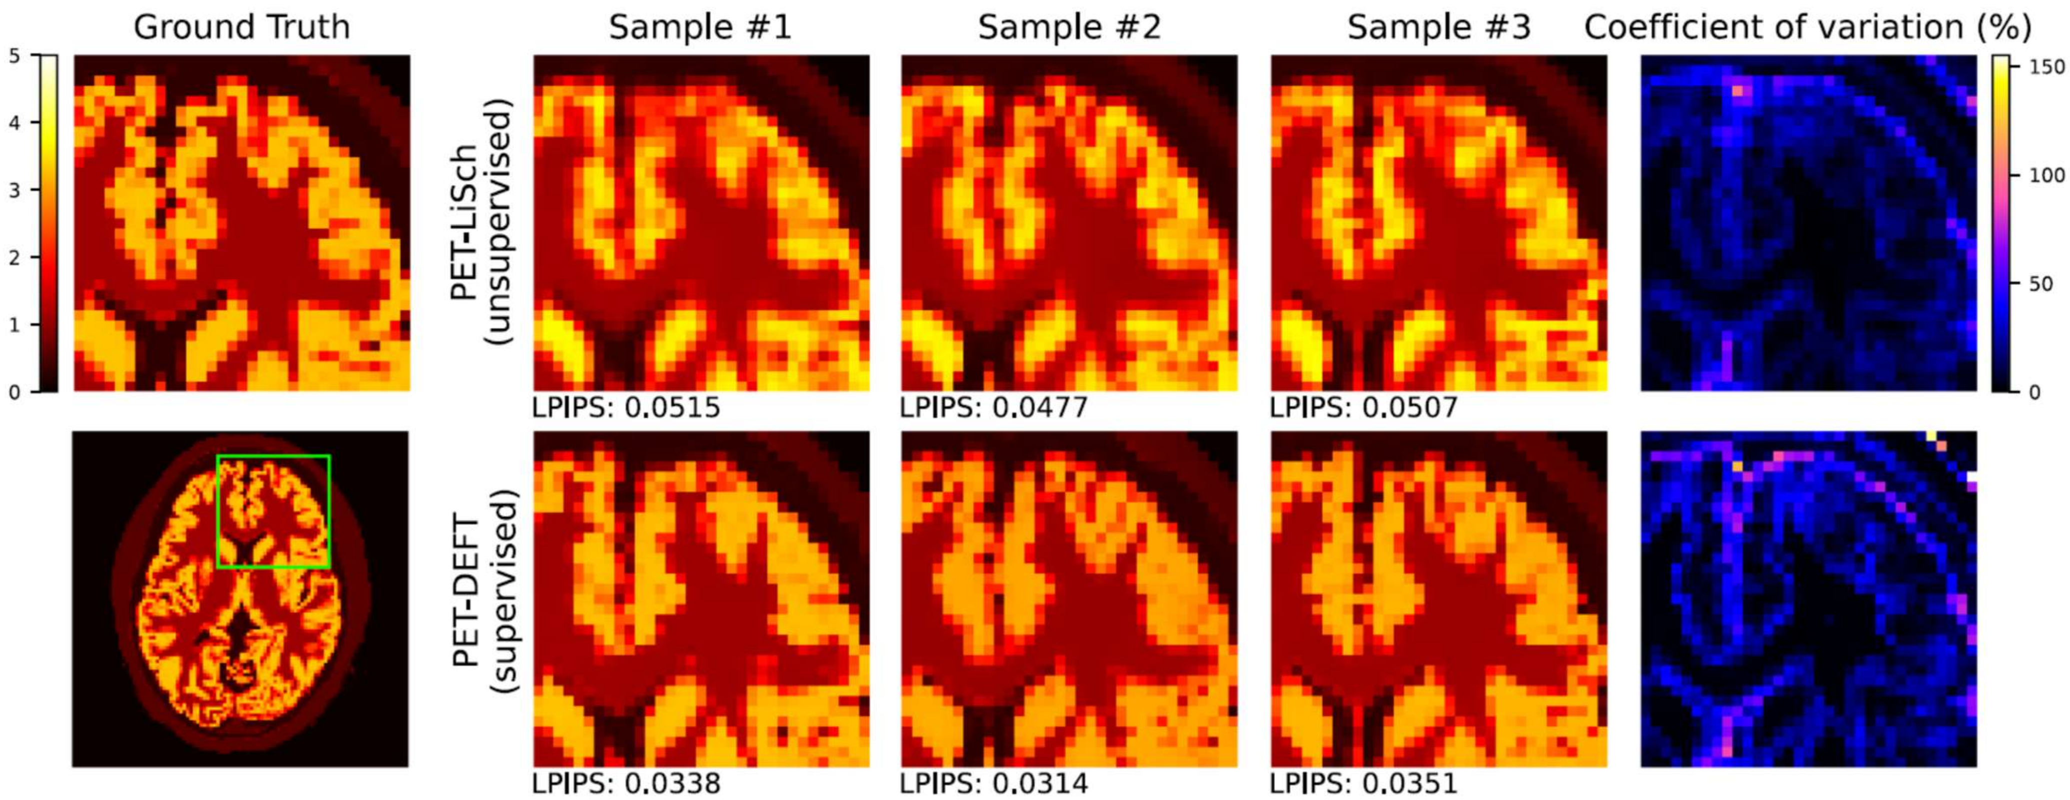

Results: Reconstruction of real 3D data

Samples show meaningfully different cortical folding patterns

The mean image delivers greater separation between the caudate and putamen

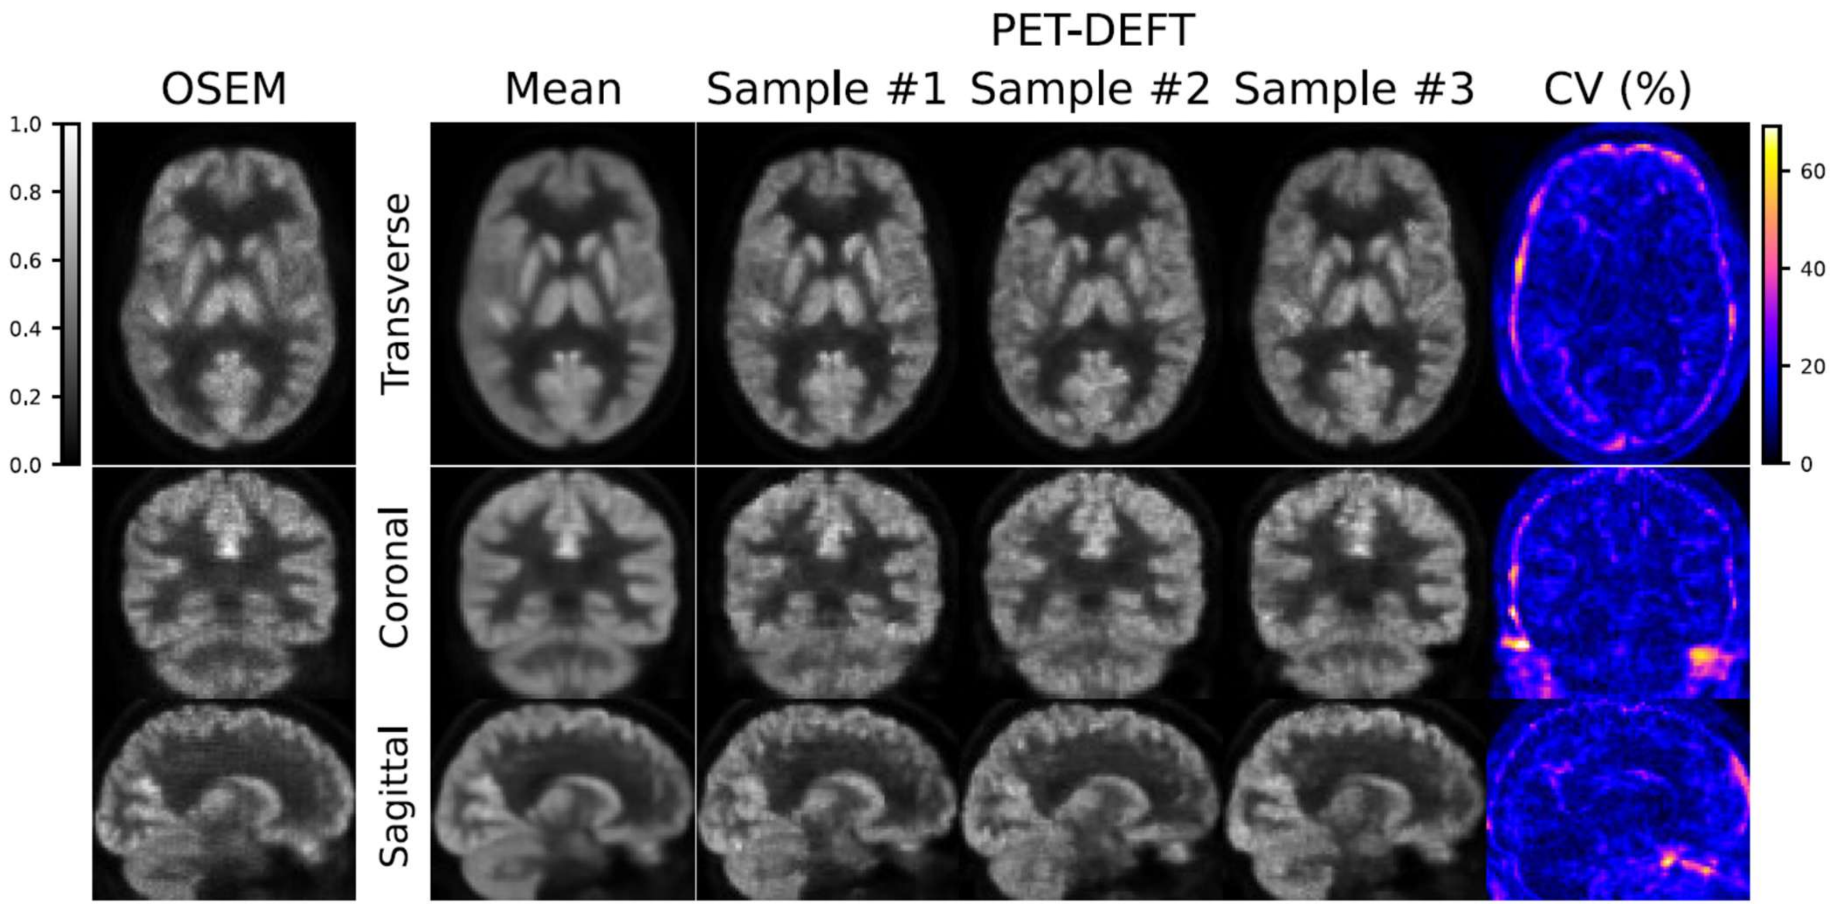

Supervised **diffusion models: Discussion**

- Incorporating measurement information into the DM's sampling process **improves reconstruction accuracy** of mean reconstructed images
- Samples generated by our method PET-DEFT **more closely match the ground truth**
- Introducing measurement information during training **may limit generalisability** to different doses and acquisition setups
- Explicitly encouraging image non-negativity enables reduction of the **reverse diffusion steps as low as 5** (may also apply to unsupervised methods)
- Computationally feasible in 3D with real data

Acknowledgements

George Webber

School of Biomedical Engineering & Imaging Sciences, King's College London

Yuya Mizuno

Institute of Psychiatry, Psychology & Neuroscience (IoPPN) at King's College London

Oliver D. Howes

Institute of Psychiatry, Psychology & Neuroscience (IoPPN) at King's College London

Alexander Hammers

Guy's and St Thomas' PET Centre & King's College London, UK

Andy P. King

School of Biomedical Engineering & Imaging Sciences, King's College London

Thank you

JOURNAL ARTICLE

Diffusion models for medical image reconstruction

George Webber, MMathCompSci ✉, Andrew J Reader, PhD

BJR|Artificial Intelligence, Volume 1, Issue 1, January 2024, ubae013, https://doi.org/10.1093/bjrai/ubae013

Published: 29 August 2024 **Article history ▾**

KING'S College LONDON