

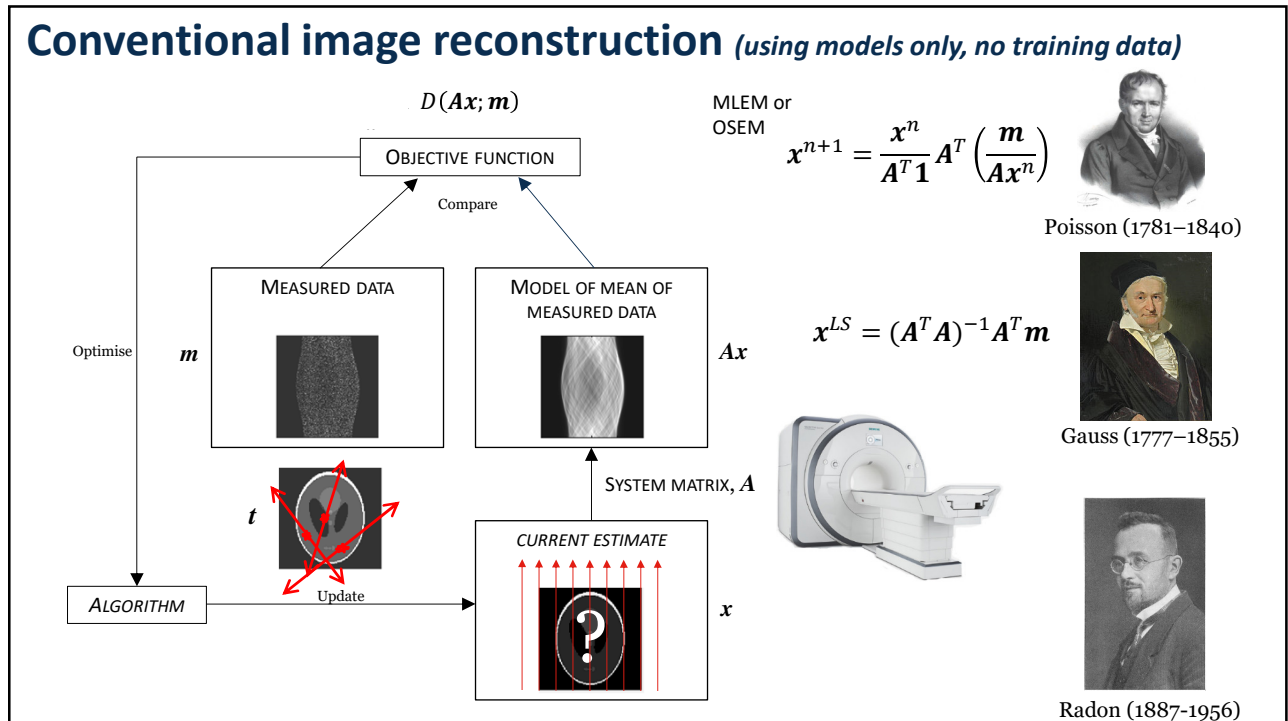
AI for PET Image Reconstruction

Andrew J. Reader

School of Biomedical Engineering and Imaging Sciences
King's College London
UK

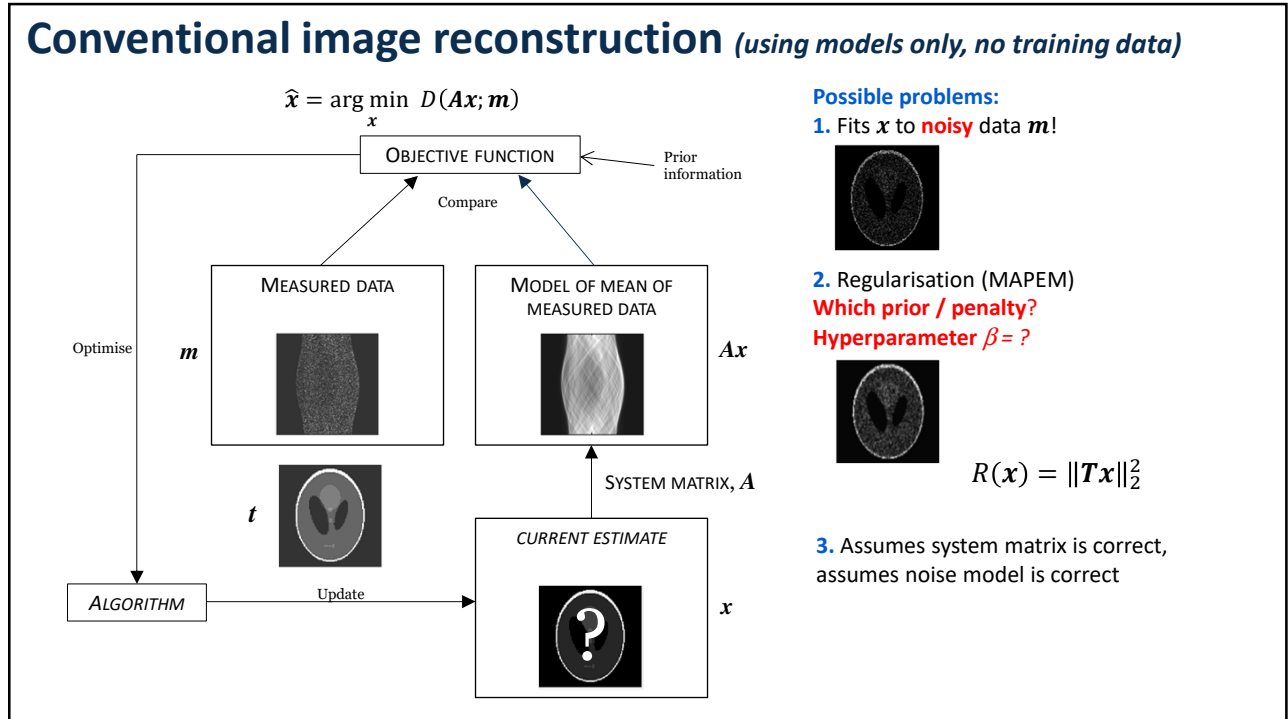
andrew.reader@kcl.ac.uk

1

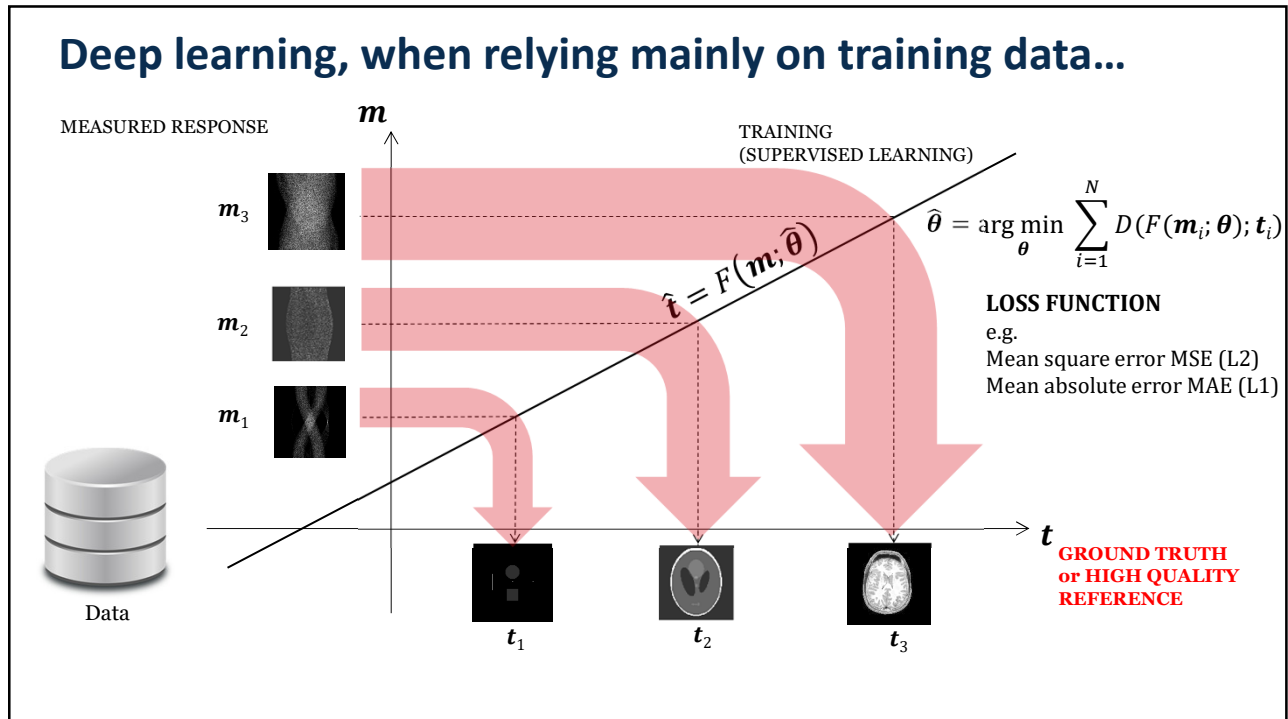


2

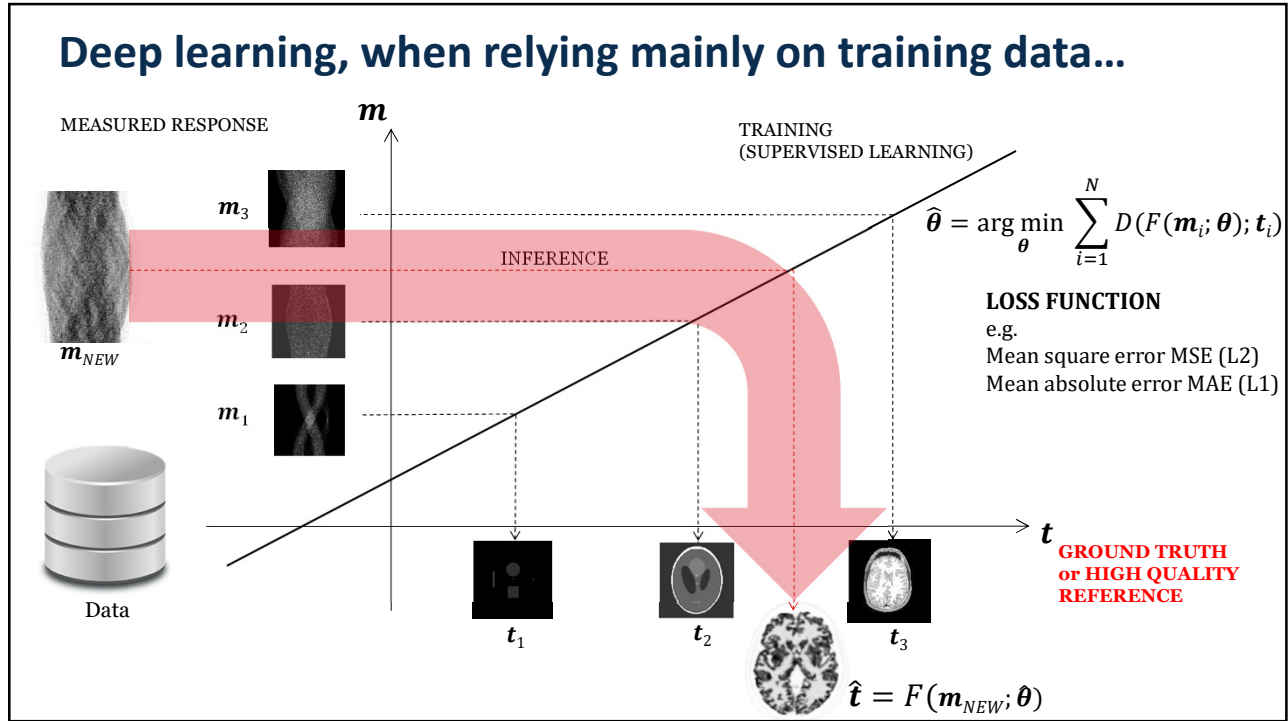
1



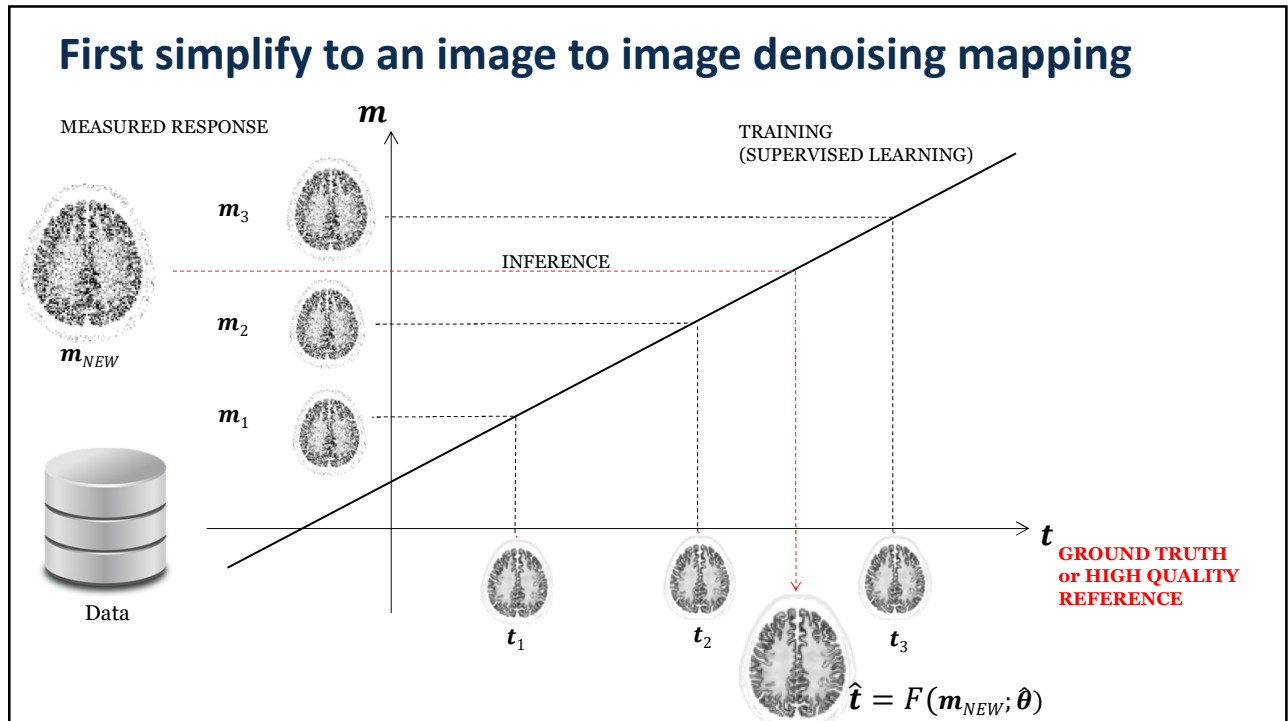
3



4




5



6


Image to image mapping with few parameters: convolution

Ground truth




t

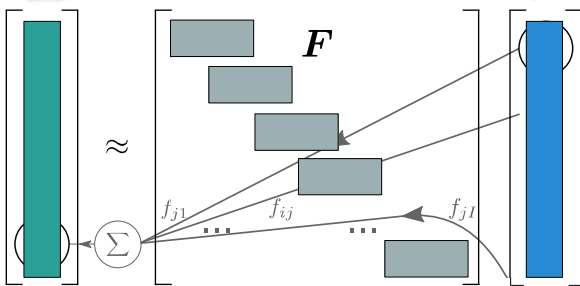
Noisy image



m

Learn this circulant matrix (convolution)






By Vincent Dumoulin,
Francesco Visin -
https://github.com/vdumoulin/conv_arithmetic, MIT,
<https://commons.wikimedia.org/w/index.php?curid=78003456>

**Convolution with a 3x3 kernel:
just 9 parameters for a whole image to image mapping!**

7

Learning 1 convolution kernel: denoising



Architecture:

One 3×3 convolution kernel
9 parameters to learn

Training data:

Input x : **noisy** image
Target t : ground truth

Loss function:

Mean square error (MSE)

$$\frac{1}{V} \sum_{v=1}^V (y_v - t_v)^2$$

Optimiser:

Gradient descent (GD)
update parameters by
subtracting gradient of loss function
Stochastic GD (SGD)

8

4

Learning 1 convolution kernel


<https://youtu.be/JvJgvjm1hco>

9

Learning 1 convolution kernel


<https://youtu.be/JvJgvjm1hco>

INPUT: brain (PET)




CNN →

OUTPUT



TARGET: brain (PET)



10

5

Learning 1 convolution kernel

<https://youtu.be/Jvjgvm1hco>

INPUT: brain (PET)

OUTPUT

TARGET: brain (PET)

11

Learning 1 convolution kernel : sharpening

INPUT
 x

5x5

OUTPUT
 y

TARGET
 t

Architecture:
One 5 × 5 convolution kernel
25 parameters to learn

Training data:
Input x : **blurred** image
Target t : ground truth

Loss function:
Mean square error (MSE)


$$\frac{1}{V} \sum_{v=1}^V (y_v - t_v)^2$$

Optimiser:
Gradient descent (GD)
Stochastic GD (SGD)

12

Convolution with bias & activation: **feature detection**

Architecture:
One 7 × 7 kernel




13

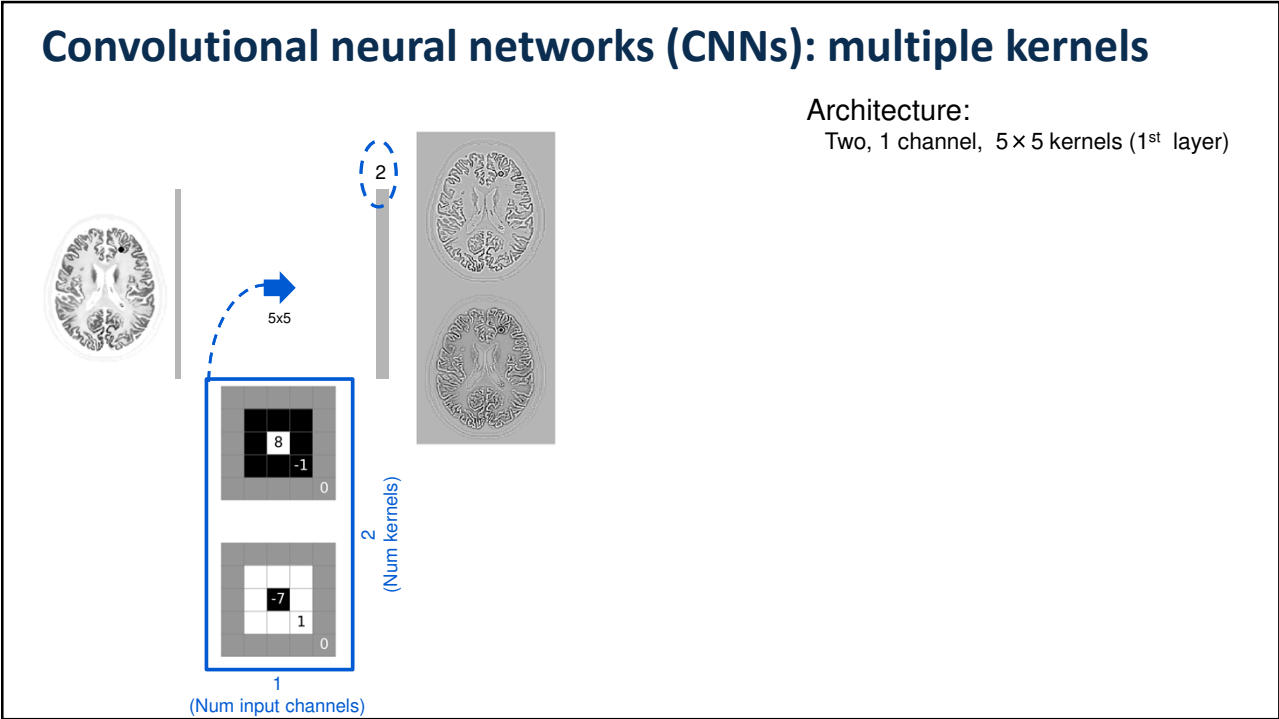
Convolution with bias & activation: **feature detection**

ReLU = rectified linear unit (just sets any negative values to zero!)

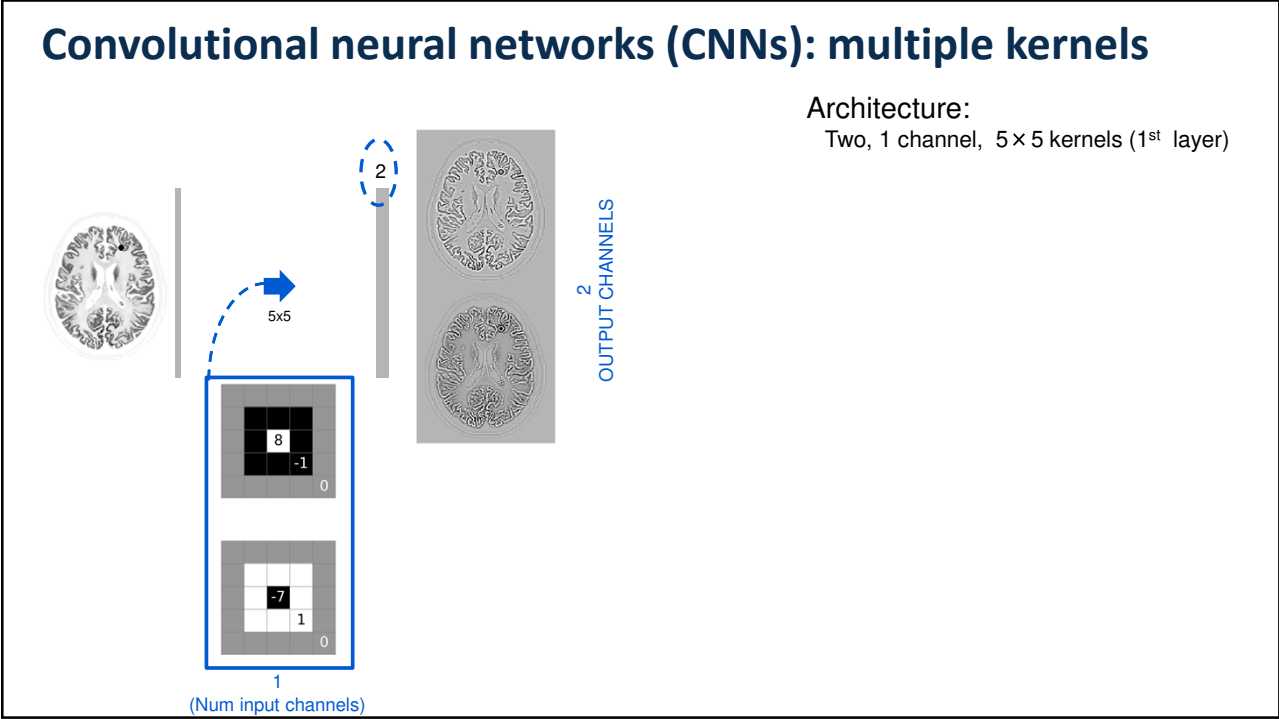
Architecture:
One 7 × 7 kernel + bias



14



15



16

Convolutional neural networks (CNNs): multiple kernels

Architecture:
 Two, 1 channel, 5 × 5 kernels (1st layer)
 One, 2 channel, 1 × 1 kernel (2nd layer)

**By fitting kernel parameters & biases...
 for many kernels and many layers...
 Task-specific processing can be learned
 ...to map inputs to desired outputs**

- >Pick out features
- >Denoise
- >Sharpen
- >Modify!

**With downsampling:
 Can learn compositions of features, feature
 hierarchies, for increasing abstraction**

17

Convolutional layer with 3 kernels

Kernel

0	0	0	0	0	0
0	0	0	0	0	0
0	0	-1	-1	-1	0
0	0	-1	8	-1	0
0	0	-1	-1	-1	0
0	0	0	0	0	0
0	0	0	0	0	0

Thresholded

*

=

18

Convolutional layer with 3 kernels

Kernel

0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	-1	-1	-1	0	0
0	0	8	-1	0	0	0
0	0	-1	-1	-1	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

Convolved

Bias

Activation

Edges

This is called a "Convolutional Layer"

3 kernels → 3 output channels

19

Convolutional layer with 3 kernels

Kernel

0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	-1	-1	-1	0	0
0	0	-1	8	-1	0	0
0	0	-1	-1	-1	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

Kernel

0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	-1	2	-1	0	0
0	0	-1	2	-1	0	0
0	0	-1	2	-1	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

Kernel

0	0	1	1	1	0	0
0	1	1	1	1	1	0
1	1	1	1	1	1	1
1	1	1	1	1	1	1
1	1	1	1	1	1	1
0	1	1	1	1	1	0
0	0	1	1	1	0	0

Convolved

Bias

Activation

Edges

Vertical Edges

Large Tumours

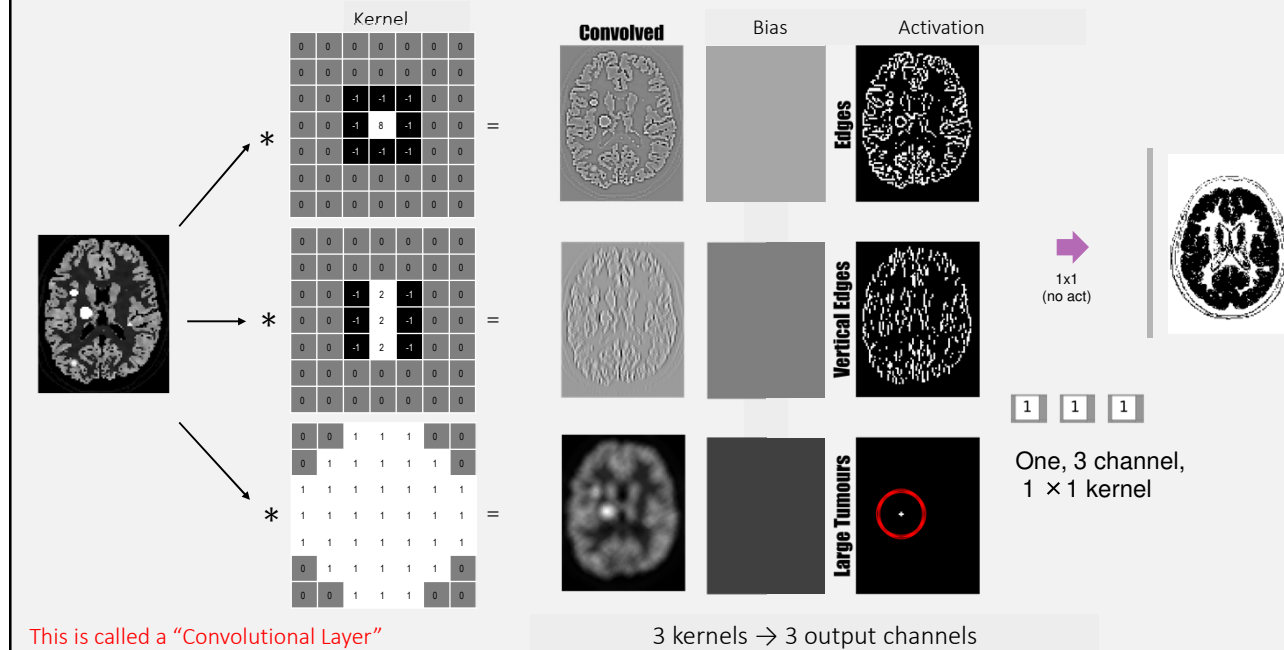
Activation function examples:
ReLU
PReLU
LReLU
...

This is called a "Convolutional Layer"

3 kernels → 3 output channels

20

Convolutional layer with 3 kernels



21

Deep learning components

1. Training data

From no training data...
...to tens of examples pairs... to thousands

2. Architecture / inductive prior for the mapping from input to output

Trainable parameters for a code structure
E.g. fully-connected (linear) layers, convolutional neural networks (CNNs), transformers

3. Loss functions to decide how well a mapping is doing its job

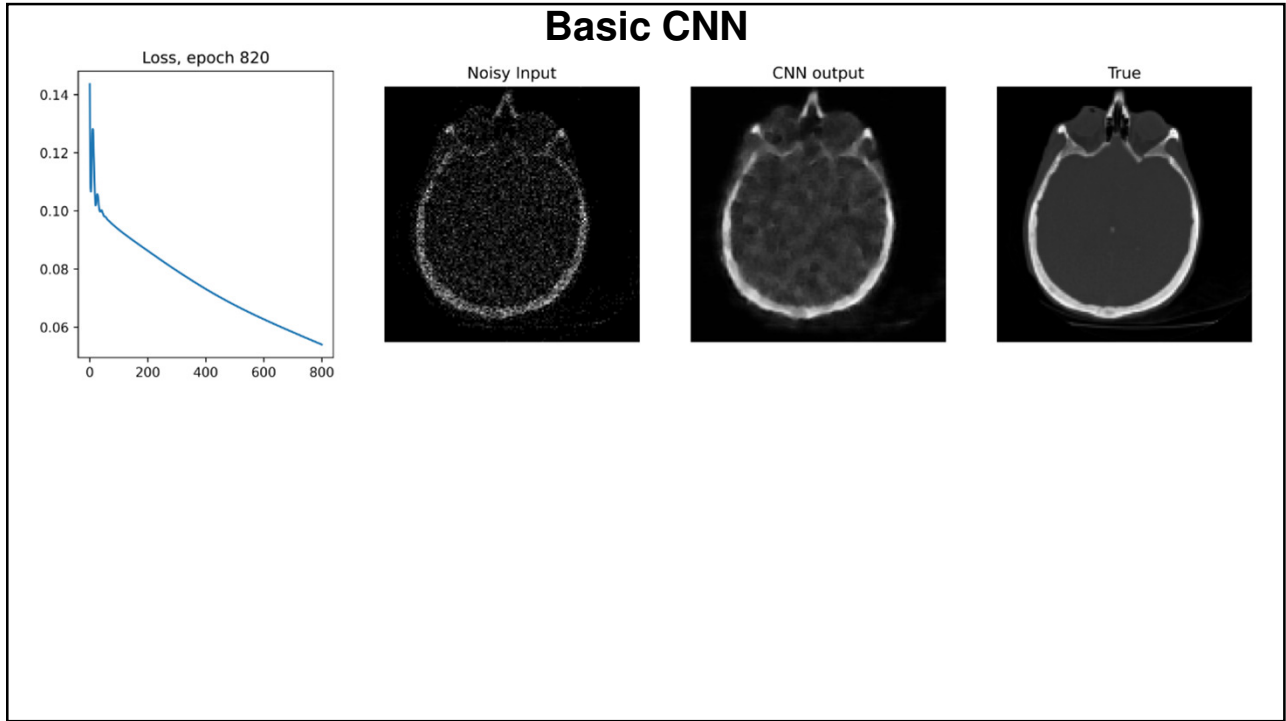
Mean squared error (MSE) or L2 norm
Mean absolute error (MAE) or L1 norm
Perceptual loss
Adversarial loss

4. Optimisers

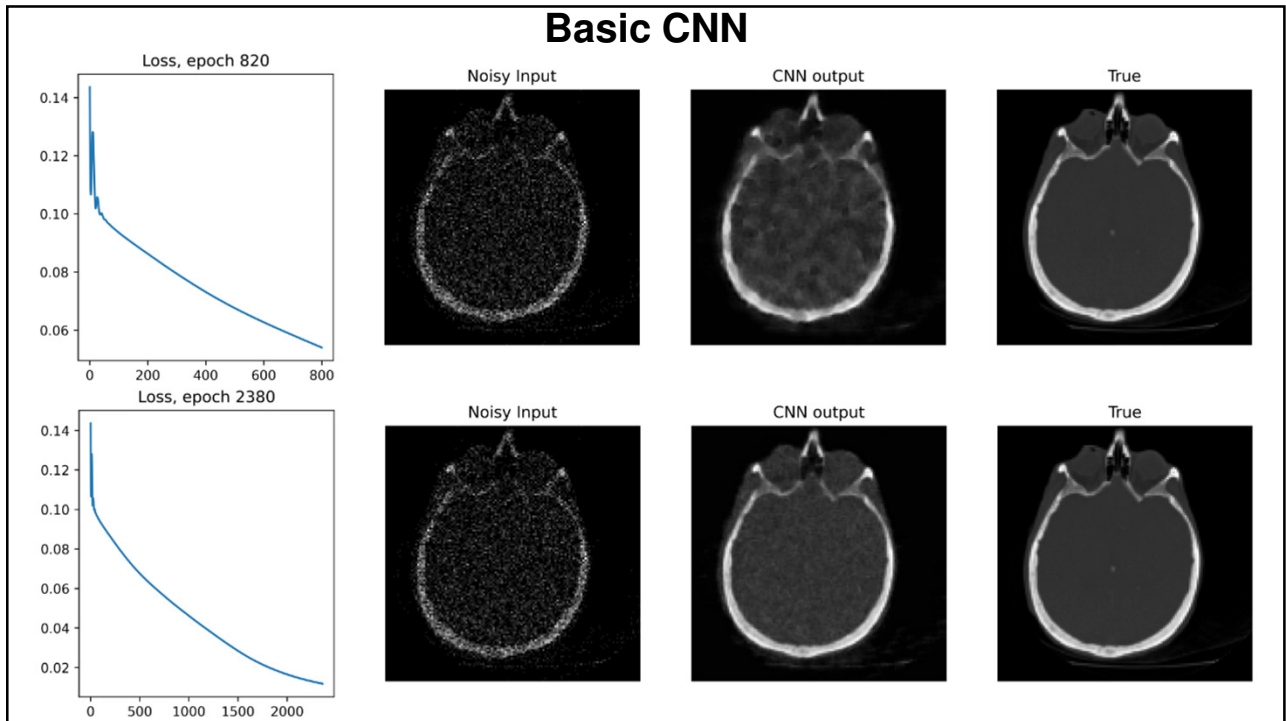
Stochastic gradient descent (SGD)
Adam
...and many more

22

11




23




24

Mapping to anatomically-guided reconstructions by a CNN


OSEM PET input



T1 MR input



anatomically-guided PET - output




Schramm *et al*, Neuroimage 2021 (<https://github.com/gschramm/pyapetnet>)


25

Mapping to anatomically-guided reconstructions by a CNN



OSEM PET input



T1 MR input



anatomically-guided PET - output

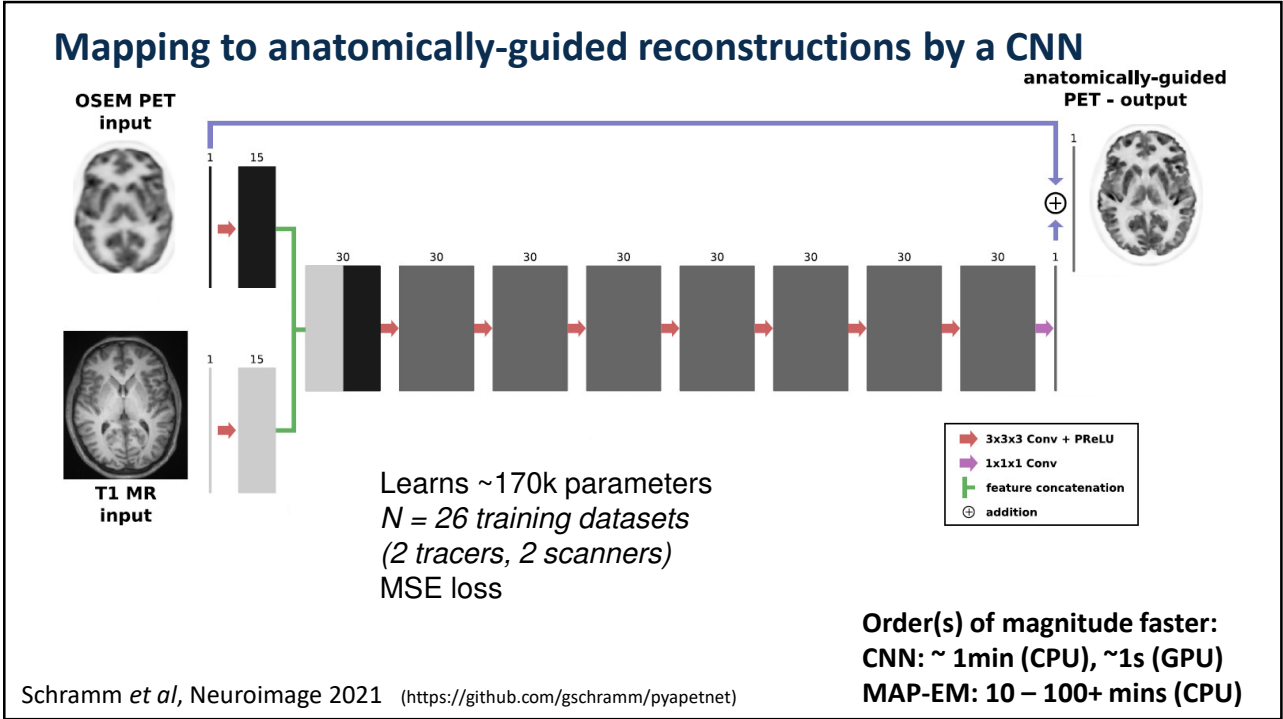


Learns ~170k parameters
N = 26 training datasets
 (2 tracers, 2 scanners)
 MSE loss

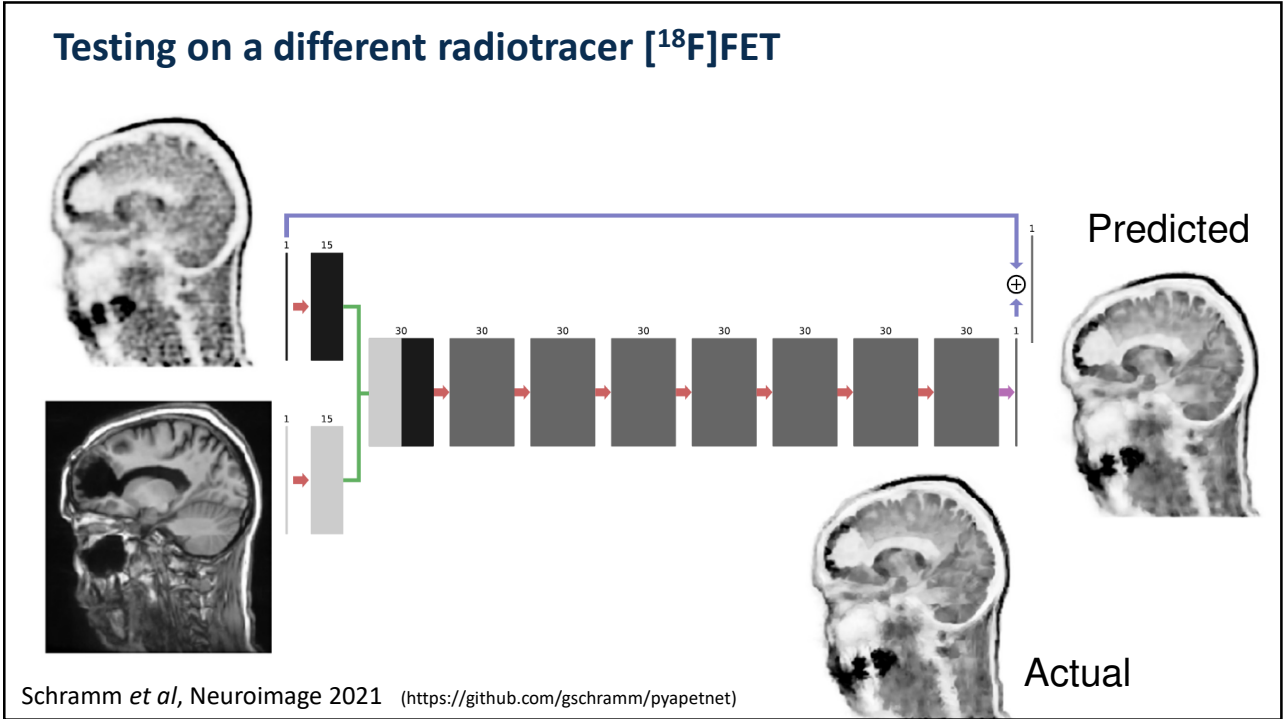
- ➔ 3x3x3 Conv + PReLU
- ➔ 1x1x1 Conv
- └ feature concatenation

Schramm *et al*, Neuroimage 2021 (<https://github.com/gschramm/pyapetnet>)

26



27



28

So far:

No downsampling / upsampling: shift-equivariant mappings

Suitable for image to image mappings

29

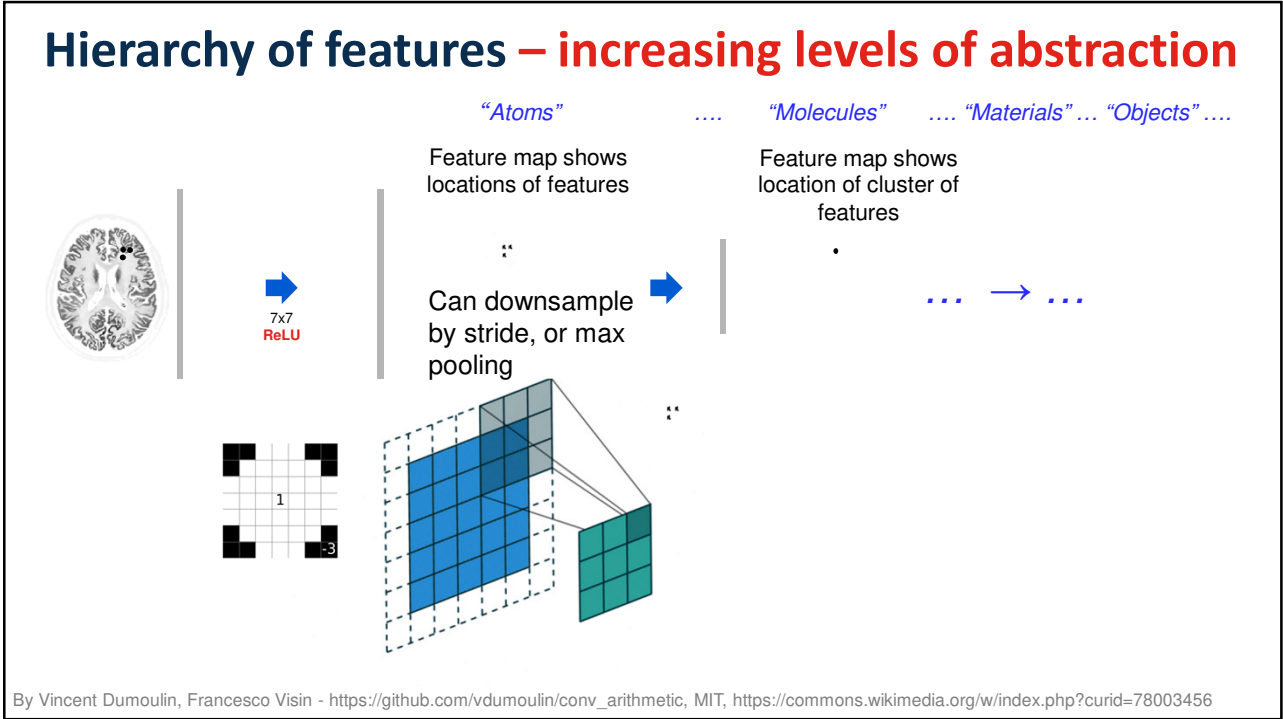
CNNs can do more

Using downsampling (or max pooling) / upsampling

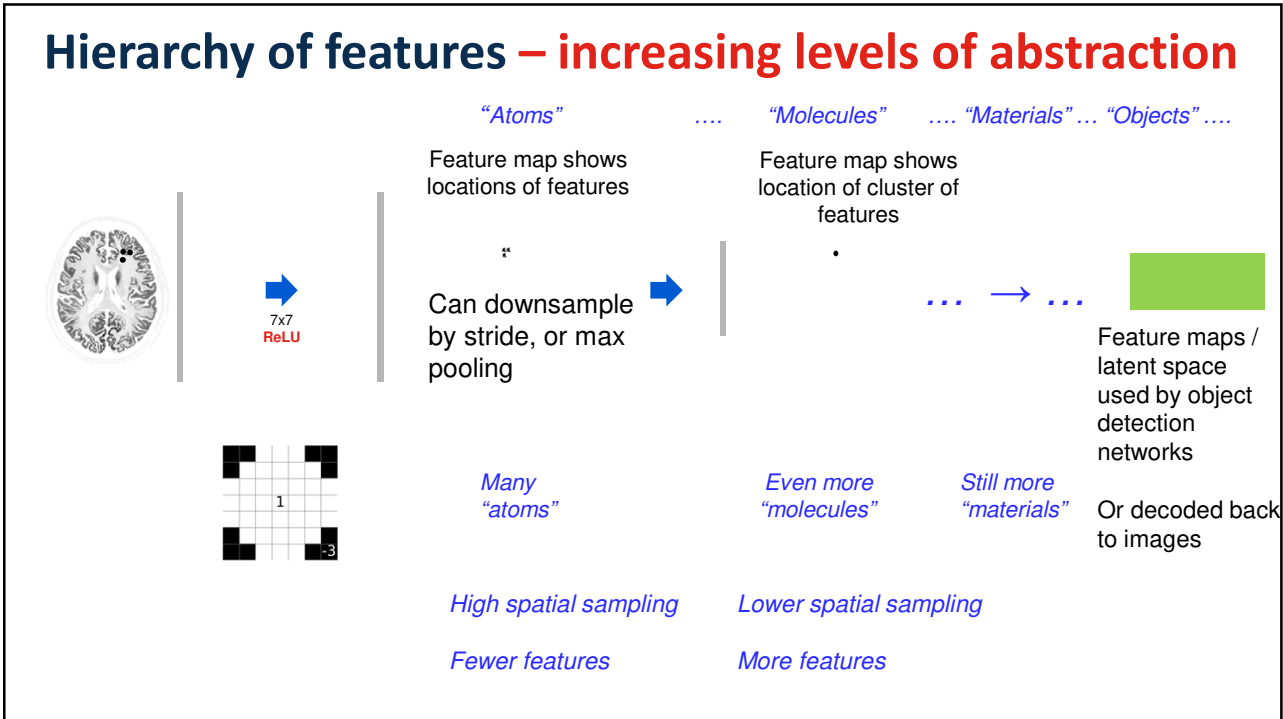
Non-linear, shift-variant mappings

Suitable for sinogram to image mappings

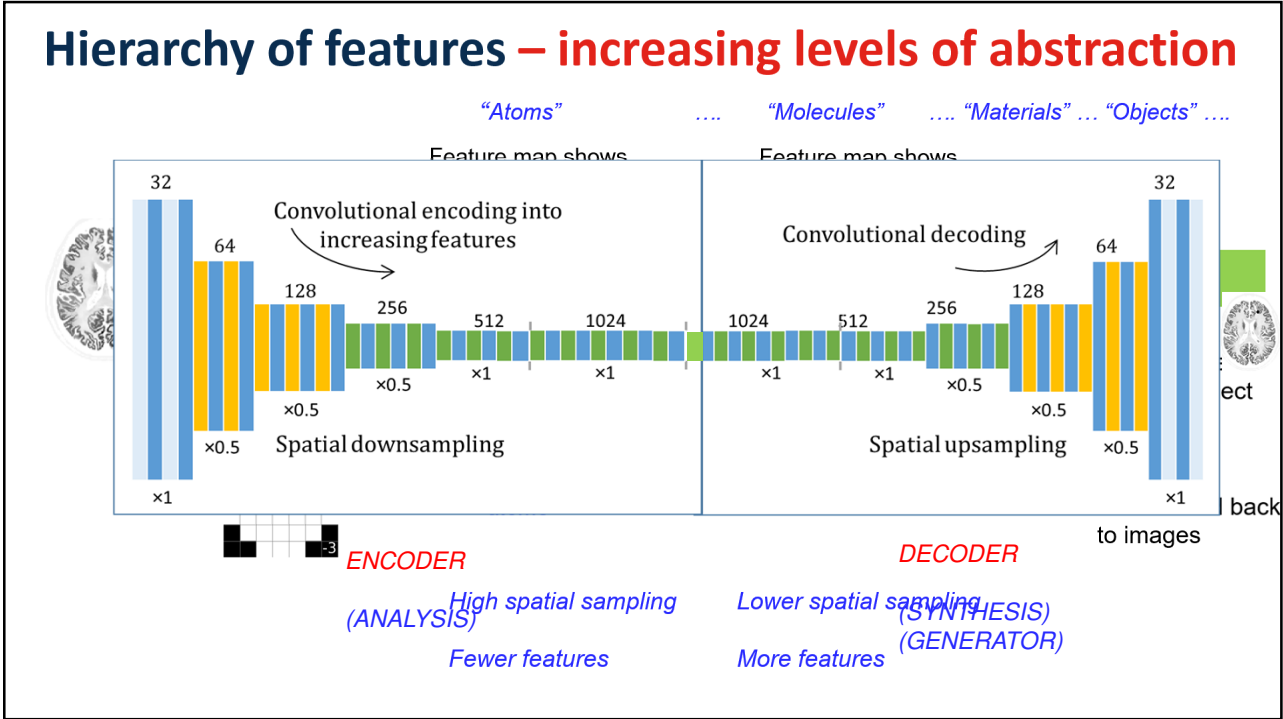
30



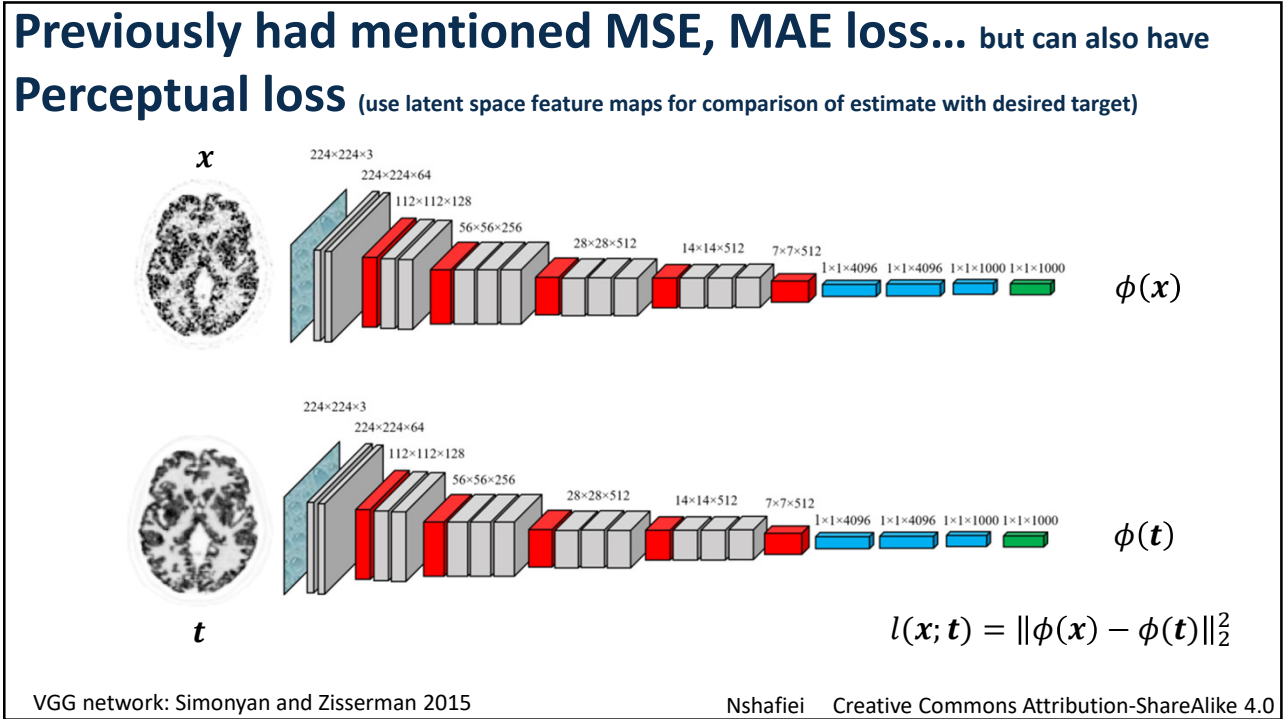
31



32




33



34

Direct DL reconstruction from sinograms: DeepPET



- Conv 7x7
- Conv 5x5
- Conv 3x3
- BN + ReLU
- Upsampling


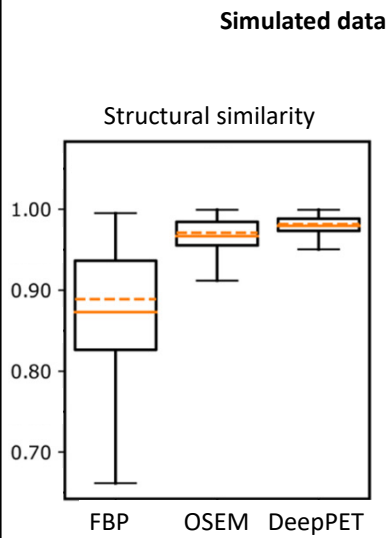
Convolutional encoder decoder:

- **200,000 simulated training data pairs** to learn **~60 million parameters** via MSE
- Note a simple linear mapping would need > 1 billion parameters


Figure represents method of Häggström *et al*, MIA 2019

37

DeepPET: results



✓ Within distribution



▪ Outside distribution

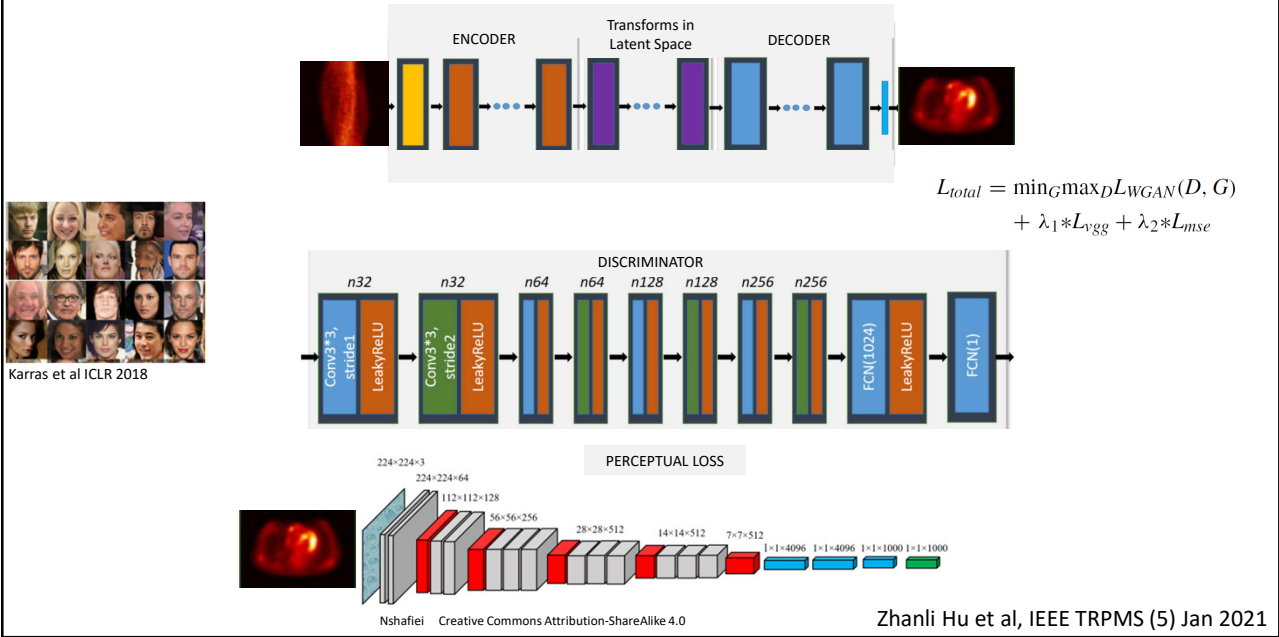
- ~10% lower noise (RMSE) than OSEM conventional reconstruction
- ~100x faster than OSEM

Häggström *et al*, MIA 2019

+discriminator, perceptual loss, DPIR-Net: Zhanli Hu *et al*, IEEE TRPMS (5) Jan 2021


38

Adding perceptual loss, and a discriminator: DPIR-Net



39

Adding a discriminator and perceptual loss - DPIR-Net



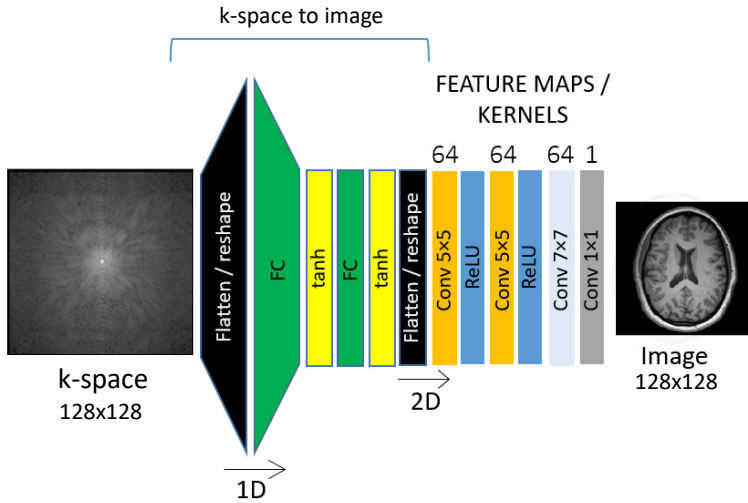
Method	50K		500K		5M	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
EM+ Gaussian	18.029	0.305	18.649	0.723	21.245	0.736
DeepPET	31.984	0.916	34.510	0.935	36.077	0.943
DPIR-Net	40.437	0.917	43.556	0.954	43.793	0.963

Zhanli Hu et al, IEEE TRPMS (5) Jan 2021

40

20

MRI reconstruction from k-space: AUTOMAP



Training on MR brain data only yielded best results (compared to natural images, or random noise as ground truth examples)

Reduced noise and artefacts compared to conventional reconstruction (RMSE nearly halved)

Faster reconstruction!

Demonstrated for PET data also.

Uses ~ 50,000 training data pairs to learn ~800 million parameters

Figure represents method of Zhu *et al*, Nature 2018

41

COMPARISON OF DIRECT RECONSTRUCTION METHODS

	NAME	ARCHITECTURE	LOSS FUNCTION	DATA SIZES	NUMBER OF TRAINING PAIRS
PET	DeepPET Häggström et al. MIA 2019	CNN (CED) [>60 million parameters]	MSE	269x288 to 128x128 (2D)	~200,000
	DPIR-Net Hu et al. TRPMS 2020	As above + discriminator [>60 million parameters]	MSE, perceptual loss, discriminator	269x288 to 128x128 (2D)	~40,000
MR	AUTOMAP Zhu et al. Nature 2018	FC layers, CNN [>800 million parameters]	MSE with L1 penalty	128x128 to 128x128 (2D)	~50,000
CT	iCT-Net Li et al, IEEE TMI 2019	CNN+FC [~ <1 million parameters]	MSE	512x512 (2D)	58 real scans [millions of simulated samples for pre-training]
	DirectPET Whiteley et al. MIC 2019 J. Med. Imag. 2020	FC layers, CNN [~350 million parameters]	MAE and perceptual loss and MS-SSIM	400x168x16 to 400x400x16	~2,000

42

Direct DL reconstruction summary

Examples: DeepPET, DPIR-Net, AUTOMAP, ...

Hägström *et al*, MIA 2019, Zhanli Hu *et al*, IEEE TRPMS (5) Jan 2021, Zhu *et al* Nature 2018


- ✓ Few model assumptions (less modeling error)
- ✓ Data driven, just the network's inductive prior
- ✓ Fast reconstructions

- Slow training
- **Huge data needs** (>>10k data pairs)
- Relearns physics & statistics
- Huge network (10-100 million parameters, just for 2D)
- Applied mainly for **2D reconstruction**, not yet really for fully 3D
- Generalisation / outside training distribution?
- Stability? Antun *et al*, PNAS 2020


43

Fast PET (for TOF histo images)

Histo Image



Simply a U-Net type architecture, only 20 million parameters
 Works for full 3D volumes
 67x faster than OSEM
 Noise reductions

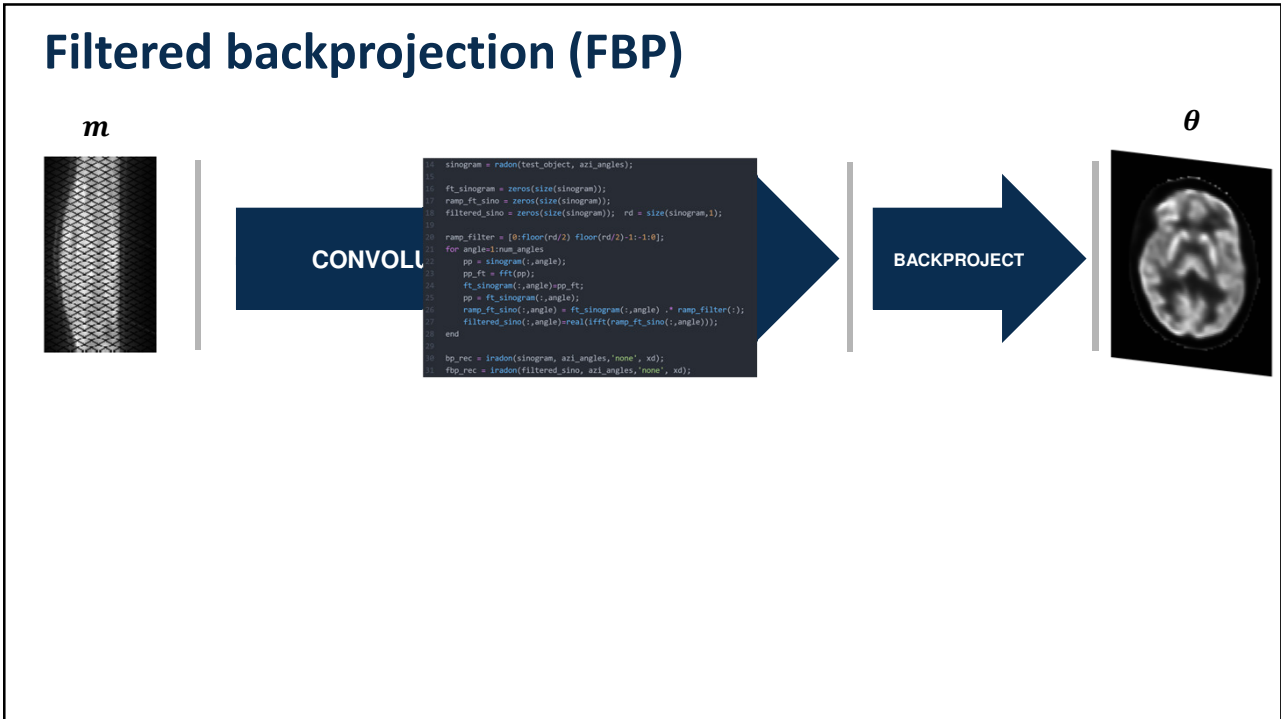


Whiteley et al. IEEE TRPMS 2021

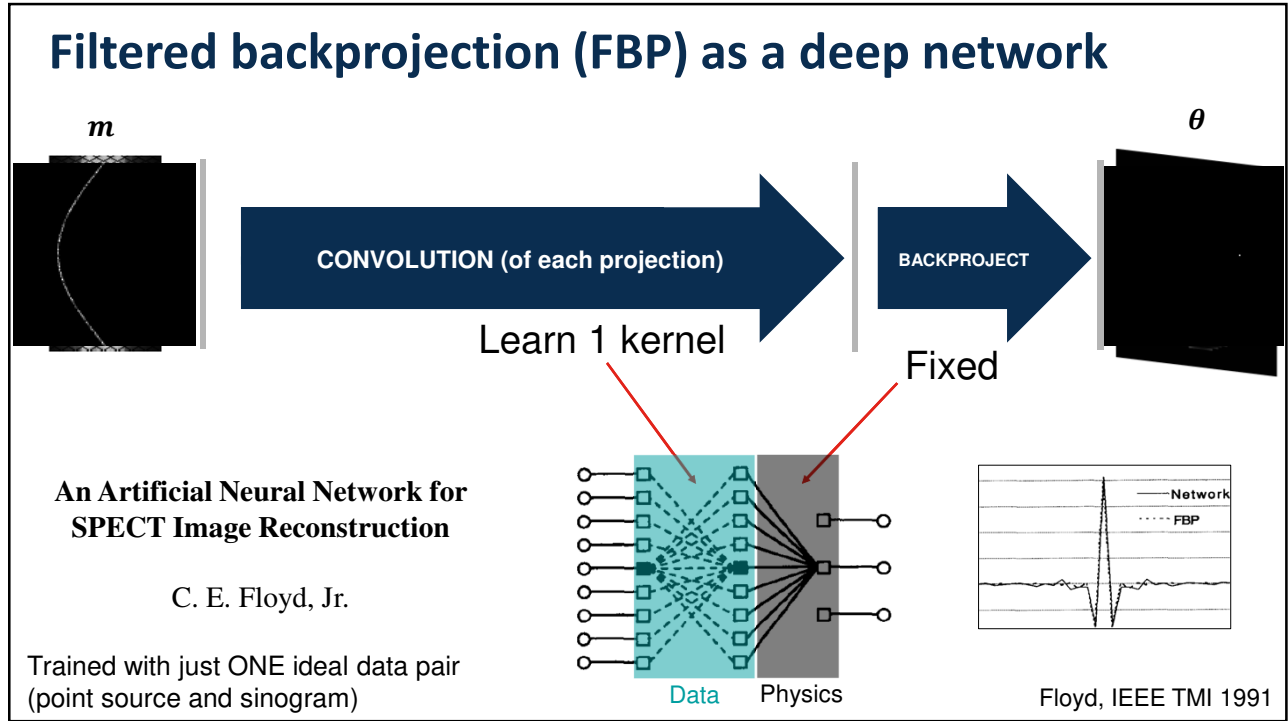
44




45



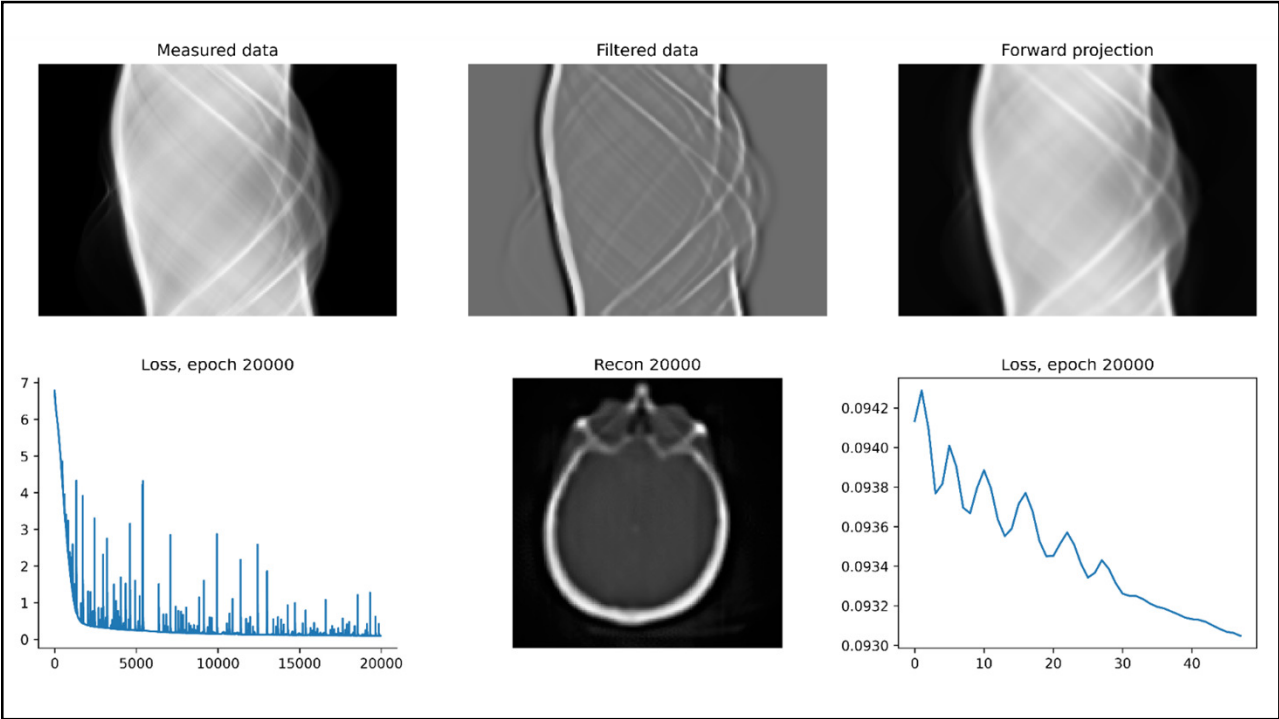
46



47



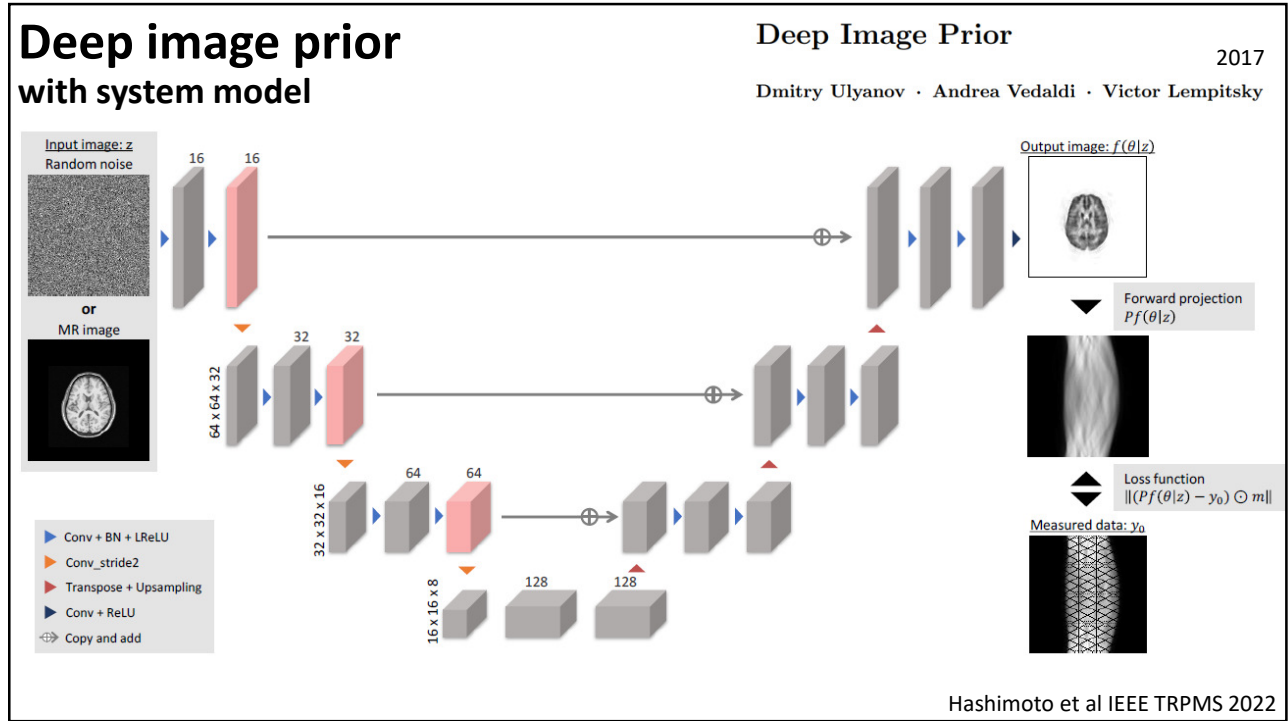
48



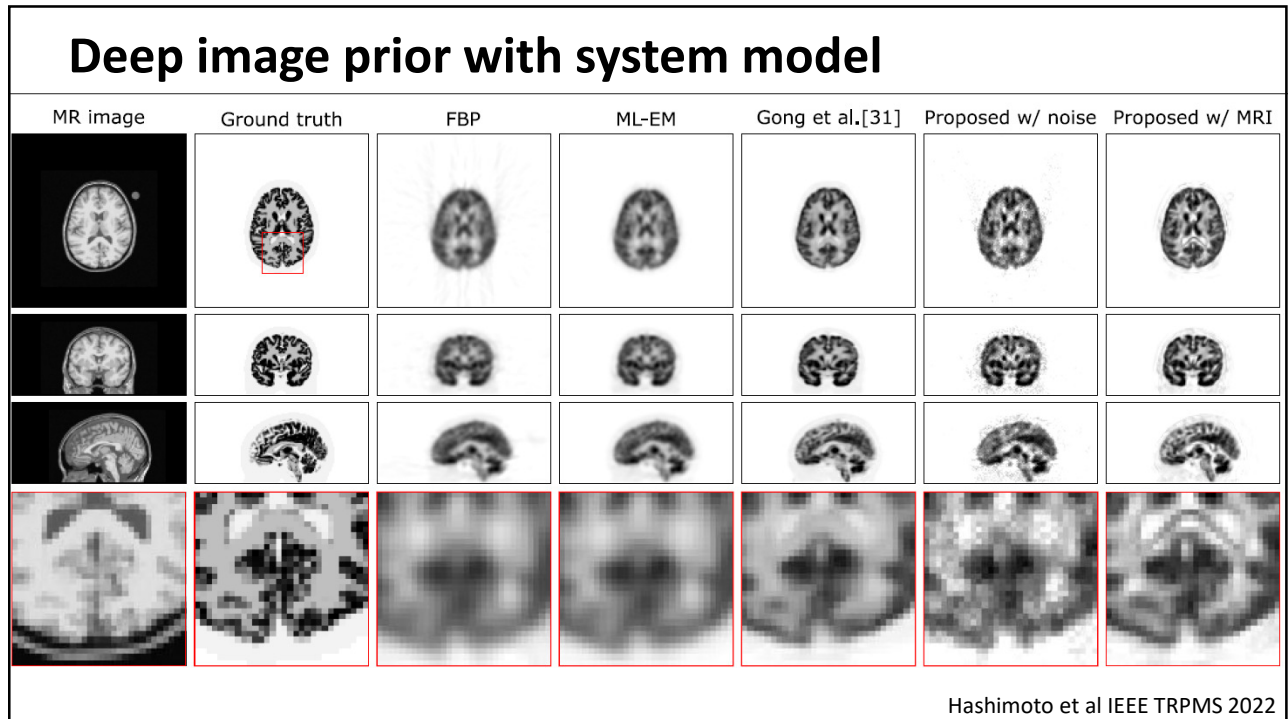
49

Deep Image Prior

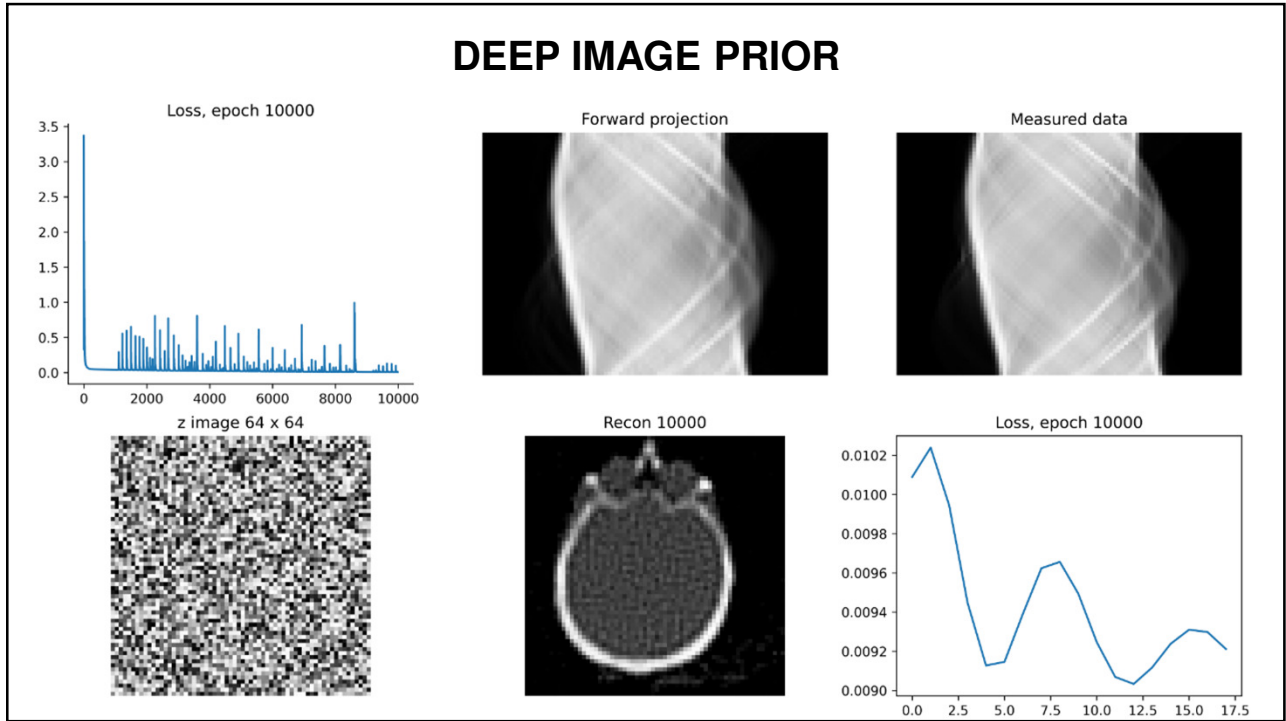
50



51



52



53

Iterative reconstruction with DL

54

AI IN PET RECONSTRUCTION AS SEEN BY



Journal of Nuclear Medicine October 2021, 62 (10) 1330-1333; DOI: <https://doi.org/10.2967/jnumed.121.262303>

HOT TOPICS

Artificial Intelligence for PET Image Reconstruction

Andrew J. Reader¹ and Georg Schramm²

¹School of Biomedical Engineering and Imaging Sciences, King's College, London, United Kingdom; and ²Division of Nuclear Medicine, Department of Imaging and Pathology, KU/UZ, Leuven, Belgium

55

Embedding deep learning into iterative reconstruction

Unrolled iterative methods:

- ✓ Iterative reconstruction uses physics and statistics modelling and theoretically convergent algorithms
- ✓ use DL for the regularisation (the prior, defined by the image manifold of the training data)

Compared to direct DL

- ✓ **Practical for 3D**
 - ✓ **Reduced training data** needs (~tens of 3D images)
 - ✓ Expect improved generalisation outside the training distribution
- **Examples**
 - Lim *et al* 2018 (BCD-Net for low count PET), TMI 2020 (Iterative NN)
 - Gong *et al* 2019 (MAPEM-Net)
 - Mehranian and Reader 2020 (FBSEM-Net)
 - Rui Hu, Huafeng Liu 2022 (TransEM)


56

28

Conventional ML-EM (and OSEM)

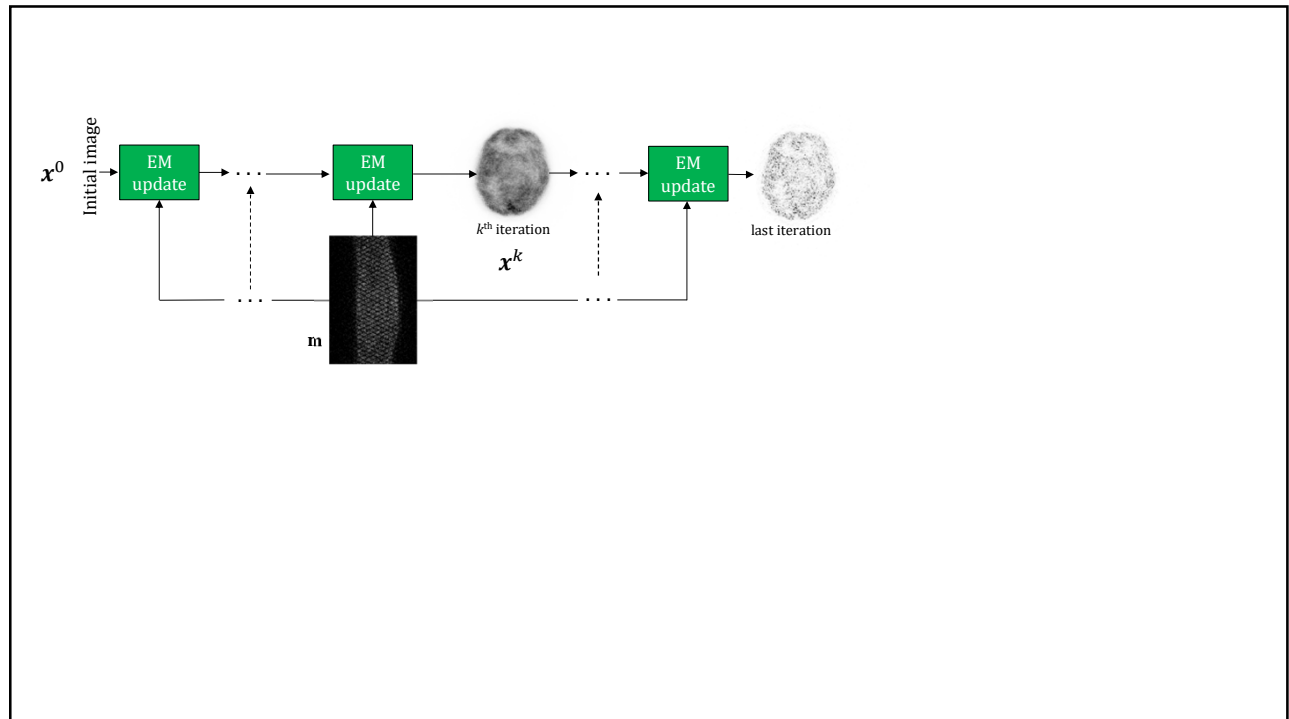
Iterative reconstruction unrolled into a **deep network**:

$$x^{n+1} = \frac{x^n}{A^T \mathbf{1}} A^T \left(\frac{m}{Ax^n} \right)$$

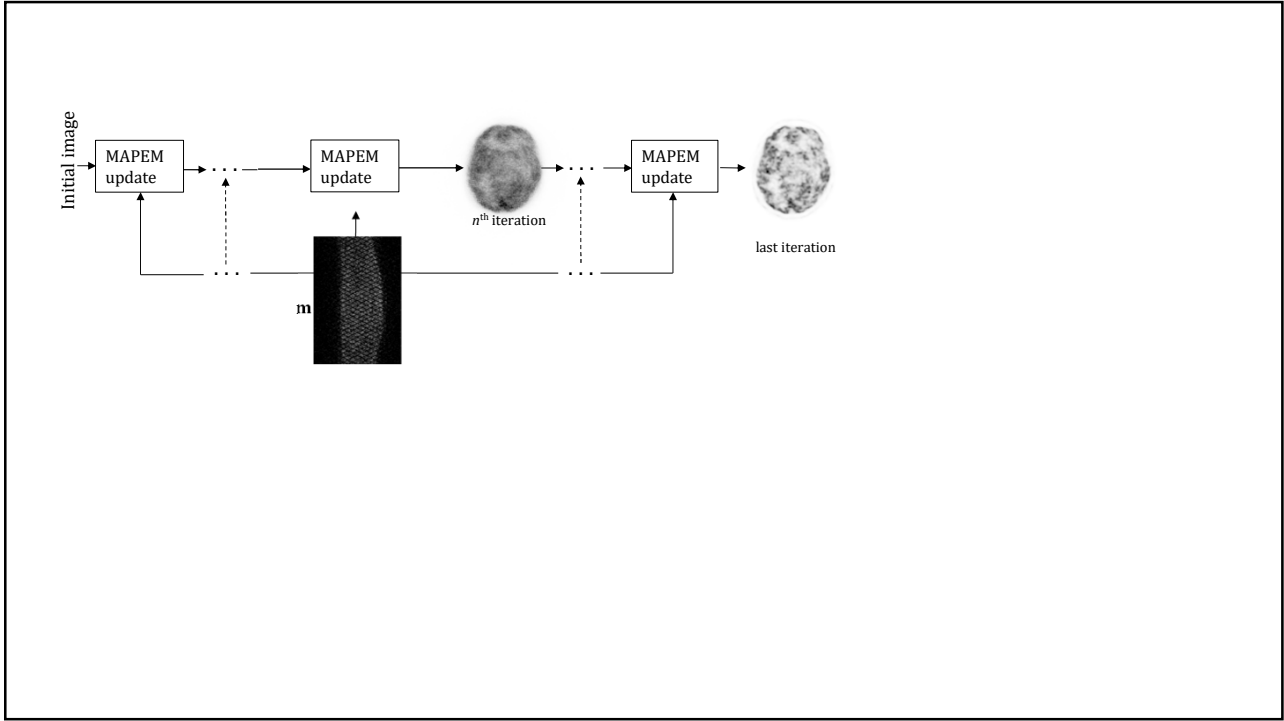


Unrolling iterative recon: Gregor & LeCun 2010

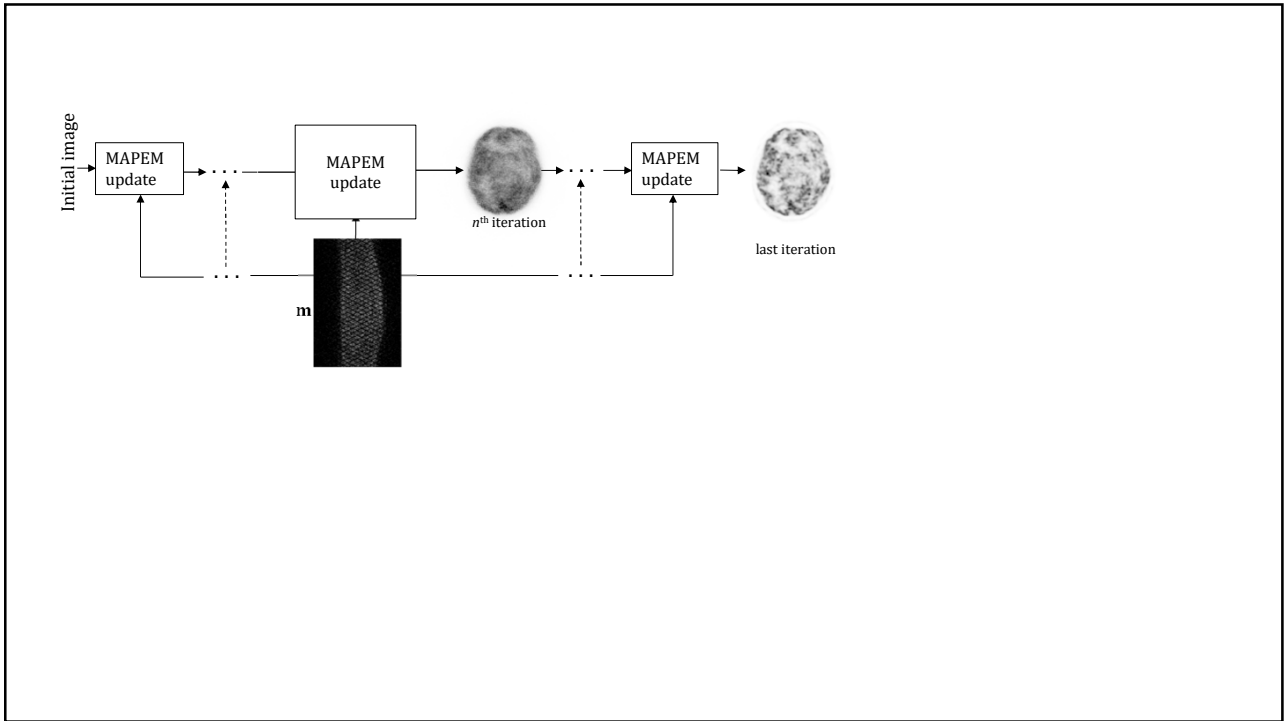
57



58



59



60

$$\hat{x} = \arg \max_x \sum_{i=1}^I (m_i \ln(Ax)_i - (Ax)_i) - \beta U(x)$$

$$x_j^{EM} = \frac{x_j^k}{\sum_{i=1}^I a_{ij}} \sum_{i=1}^I a_{ij} \frac{m_i}{(Ax^k)_i}$$

$$x_j^{SM} = \frac{1}{2 \sum_{l=1}^J w_{jl}} \sum_{l=1}^J w_{jl} (x_l^k + x_j^k)$$

De Pierro TMI 1995

$$v_j = \frac{\sum_{l=1}^J w_{jl}}{s_j}$$

$$s_j = (A^T \mathbf{1})_j$$

$$x_j^{k+1} = \frac{2x_j^{EM}}{(1 - \beta v_j x_j^{SM}) + \sqrt{(1 - \beta v_j x_j^{SM})^2 + 4\beta v_j x_j^{EM}}}$$

61

Example of PET regularisation

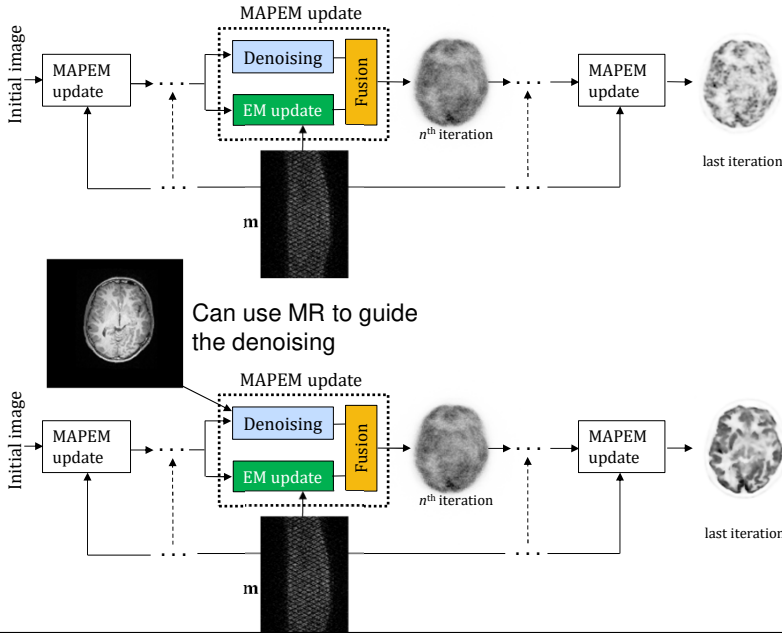
Mathematically convenient
(e.g. quadratic prior)

Unlikely to be optimal

De Pierro TMI 1995

62

Example of PET regularisation



Mathematically convenient (e.g. quadratic prior)

Unlikely to be optimal

De Pierro TMI 1995


How much MR guidance?

Unlikely to be optimal

...use deep learning

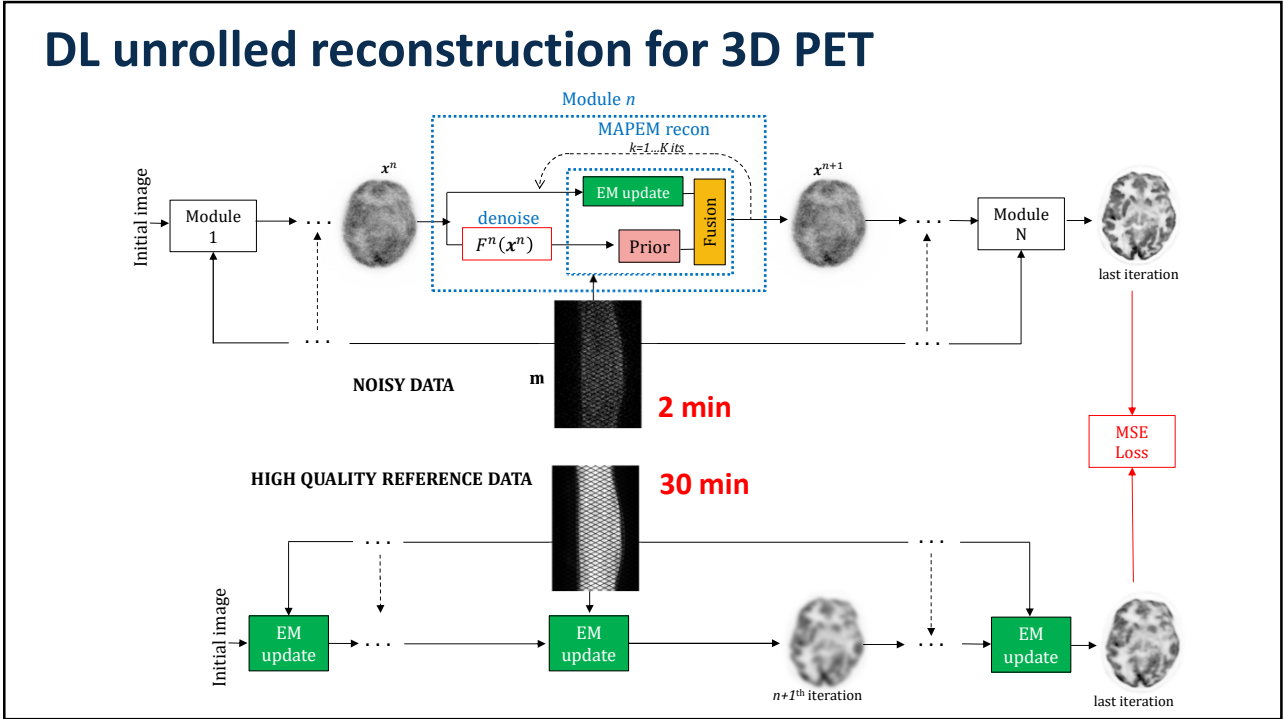
Bowsher IEEE NSS MIC 2004

63

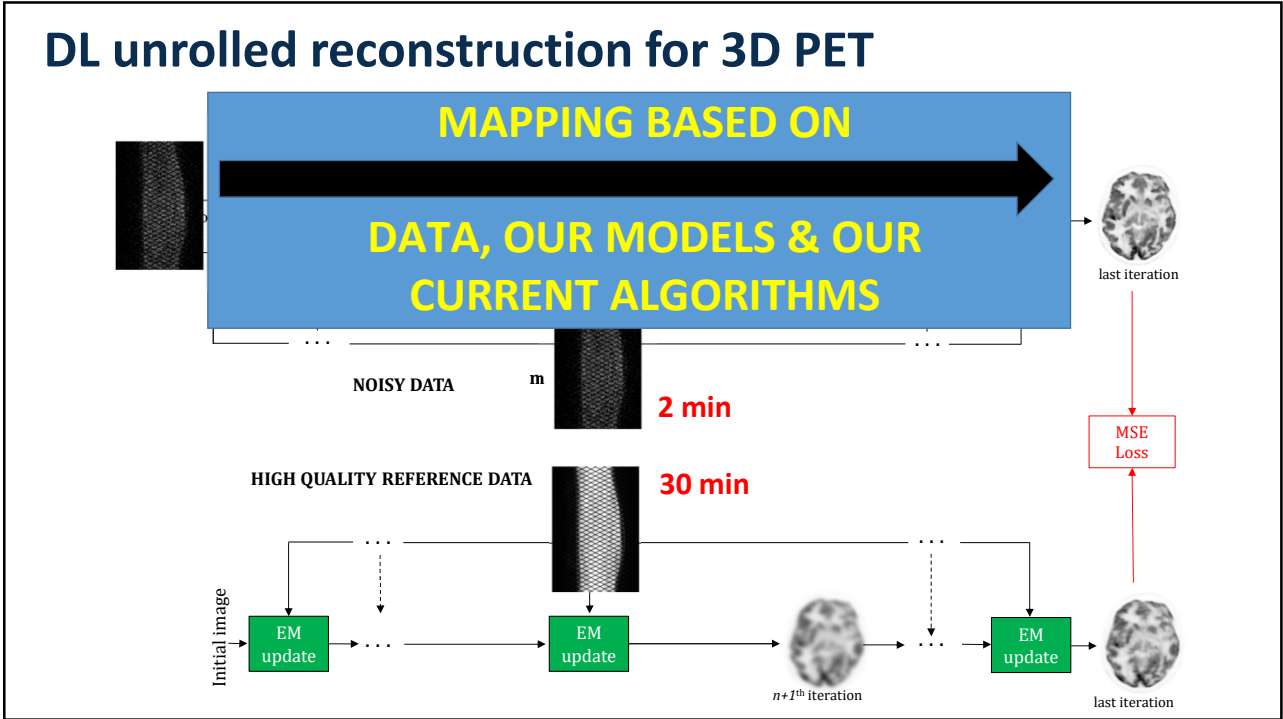


64

32



65



66

FBSEM-Net

Given current x^n :

$$\hat{x} = \arg \max_x \sum_{i=1}^I (m_i \ln(Ax)_i - (Ax)_i) - \beta U(x)$$

Proximity operator

$$x^{SM} = x^n - \gamma \beta \nabla U(x^n) = F(x^n) \quad \text{Learned regularisation}$$

$$x^{n+1} = \arg \max_x L(x|m) - \frac{1}{2\gamma} \|x - x^{SM}\|^2 \quad \text{Poisson log-likelihood with proximity to denoised update}$$

Complete data log-likelihood as a surrogate for $L(x|m)$, then can use EM update image

$$x^{EM} = \frac{x^n}{A^T \mathbf{1}} A^T \left(\frac{m}{Ax^n + \rho} \right)$$

$$\delta_j = \frac{1}{\gamma} \frac{1}{(A^T \mathbf{1})_j}$$

Simple quadratic penalty, so can use:

$$x_j^{n+1} = \frac{2x_j^{EM}}{(1 - \delta_j x_j^{SM}) + \sqrt{(1 - \delta_j x_j^{SM})^2 + 4\delta_j x_j^{EM}}}$$

Mehranian and Reader, IEEE TRPMS 2020

67

FBSEM-Net: real $[^{18}\text{F}]$ FDG data

Initial empty image

PET data
Low count (short scan, or low dose)

MRI

Reconstructed image of $[^{18}\text{F}]$ FDG

Learns ~77k parameters
 $N = 35$ training datasets

Uses MR, and the physics with DL

Mehranian and Reader, IEEE TRPMS 2020

68

FBSEM-Net: real $[^{18}\text{F}]\text{FDG}$ data

MRI	Reference	OSEM	FBSEM-Net

Uses MRI and AI
FBSEM-Net

Recent variations:
Sequential training [Corda d'Incan et al IEEE TRPMS 2021]
Using transformers [Rui Hu, Huafeng Liu 2022]

Mehranian and Reader, IEEE TRPMS 2020

30 min **2 min**

69

FBSEM-Net: real $[^{18}\text{F}]\text{FDG}$ 3D data

MRI	Reference	OSEM	U-Net	FBSEM-Net

← Uses MRI →

Mehranian and Reader, IEEE TRPMS 2020

30 min **2 min**

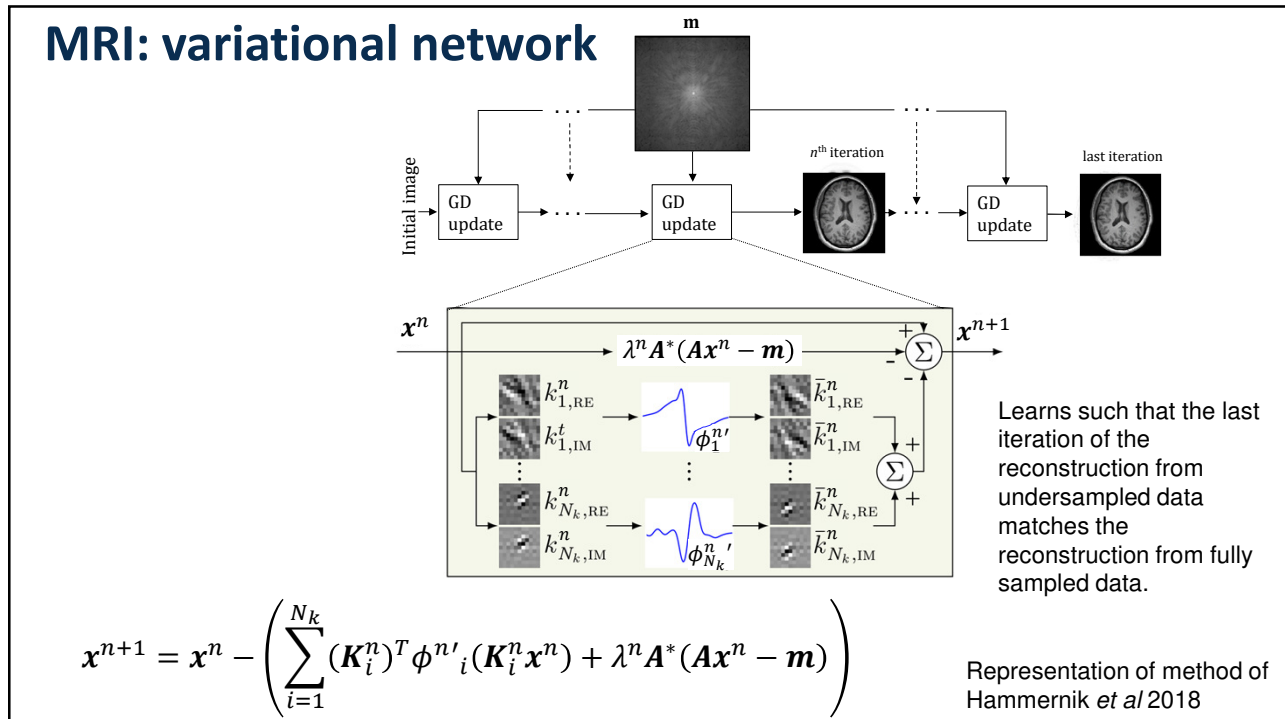
70

COMPARISON OF UNROLLED METHODS FOR PET

NAME	ARCHITECTURE (Reduced parameters)	LOSS FUNCTION	DATA SIZES 3D	NUMBER OF TRAINING PAIRS (LOW NUMBER)	BACKPROP
MAPEM-Net Gong et al 2019	CNN (U-Net) [>8 x 2 million parameters] Iteration/module dependent	MSE For end image	128x128x105	~18	Through all layers including EM update (Memory intense)
FBSEM-Net Mehranian & Reader 2020 *	CNN [77,000 parameters] Same for all iterations/modules	MSE For end image	114x114x128	~35	Through all layers, excluding EM update (Memory intense)
BCD-Net Lim et al 2018, INN Lim et al 2020	[10x4000 =40,000 parameters] Iteration/module dependent	MSE For current module compared to true/reference	200x200x112	~4	Training at iteration module only (Not demanding)

* New iteration-dependent target version: Corda-D'Incan et al TRPMS 2021
Sequential training also

71




72

Self-supervised MRI reconstruction

Supervised: images from fully-sampled data used for training


Self-supervised: half data used for k-space to reconstruct from, other half used for loss function



R=4



73



74

Towards Learned Primal Dual (LPD) – physics + data!



Figure Adapted from Massimiliano Colarieti-Tosti
Guazzo & Colarieti-Tosti J Imaging. 2021

Original method: Adler & Oktem IEEE TMI 2018

75

Towards Learned Primal Dual (LPD)

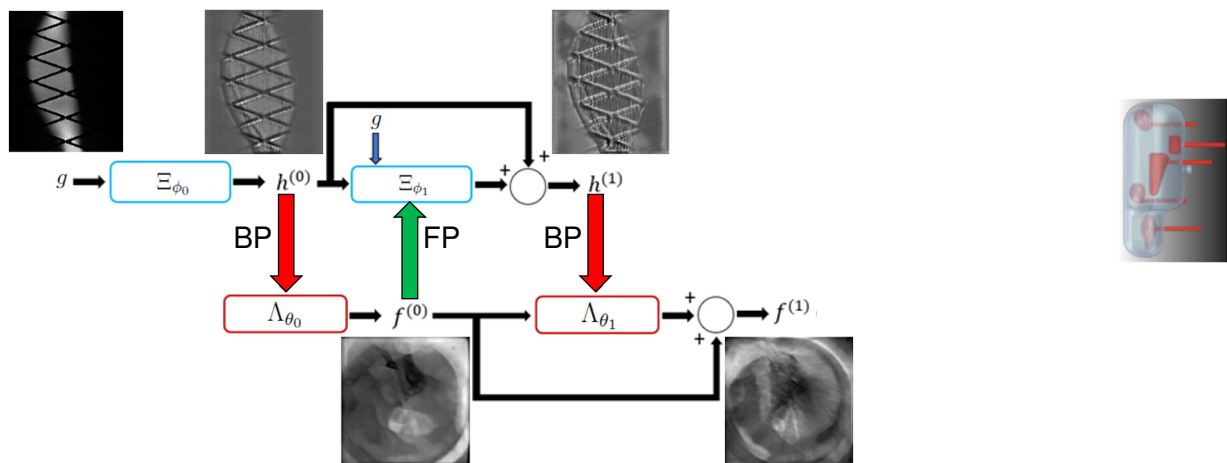
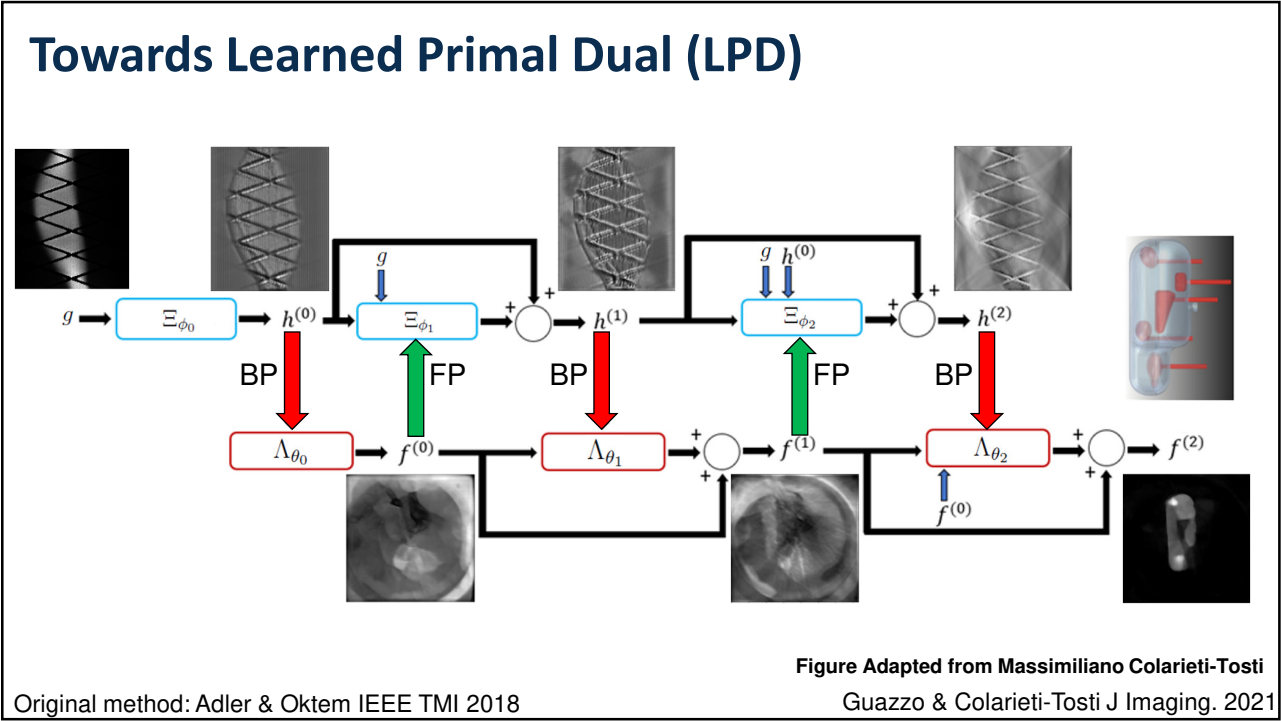


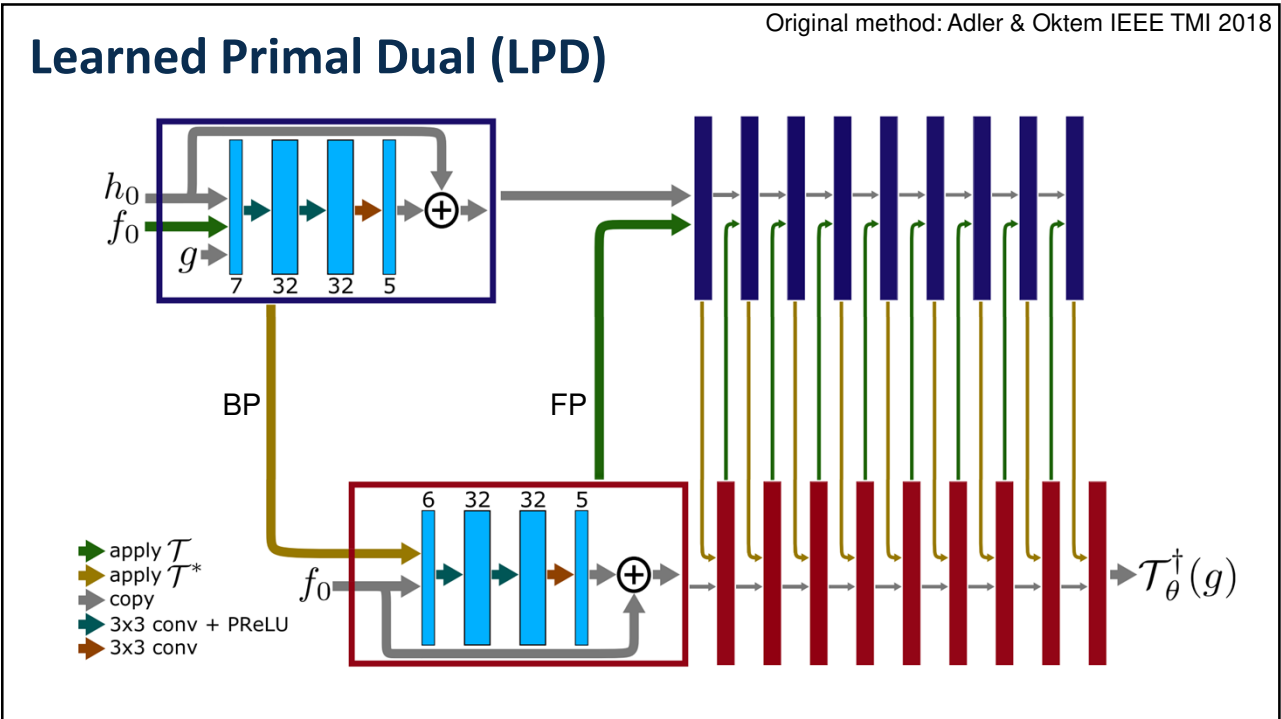
Figure Adapted from Massimiliano Colarieti-Tosti
Guazzo & Colarieti-Tosti J Imaging. 2021

Original method: Adler & Oktem IEEE TMI 2018

76




77



78


Learned Primal Dual (LPD)



Adler & Oktem IEEE TMI 2018

79

Learned Primal Dual (LPD)



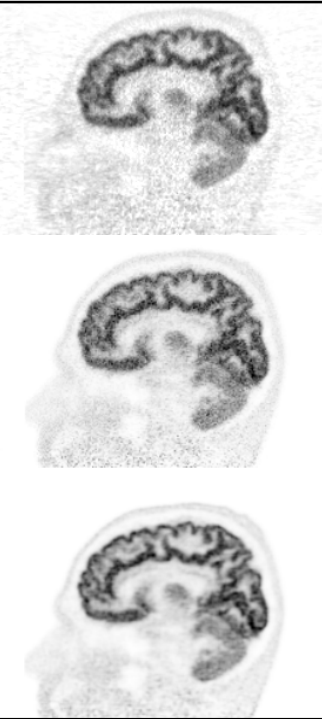
Method	PSNR	SSIM	Runtime	Parameters
FBP	33.65	0.830	423	1
TV	37.48	0.946	64 371	1
FBP + U-Net denoising	41.92	0.941	463	10^7
Learned Primal-Dual, linear	44.11	0.969	620	$2.4 \cdot 10^5$

➤ **2k training data pairs to learn ~240k parameters via MSE**

Adler & Oktem IEEE TMI 2018

80

40



School of Biomedical Engineering and Imaging Sciences

REVIEW

END OF PRESENTATION



81