Deep Learning MRI Reconstruction

M229 Advanced Topics in MRI

Shu-Fu Shih, Ph.D.

5/22/2025

Outline

- (1) General deep learning concepts
- (2) Introduction to classic convolutional neural networks
- (3) Considerations for applying deep learning in MRI reconstruction
- (4) Challenges of deep learning MRI reconstruction

Part 1: General deep learning concepts

Deep learning (DL)

- Deep learning is a branch of machine learning that relies on artificial neural networks composed of multiple interconnected layers, enabling the system to automatically learn and represent intricate patterns and relationships within large and complex datasets.

Deep learning (DL)

Need a lot of layers and trainable parameters

- Deep learning is a branch of machine learning that relies on artificial neural networks composed of multiple interconnected layers, enabling the system to automatically learn and represent intricate patterns and relationships within large and complex datasets.

It can "learn" without a lot of manual intervention

Data have inherent structures or patterns

Usually require a lot of data for training

Deep learning (DL)

Need a lot of layers and trainable parameters

- Deep learning is a branch of machine learning that relies on artificial neural networks composed of multiple interconnected layers, enabling the system to automatically learn and represent intricate patterns and relationships within large and complex datasets.

It can "learn" without a lot of manual intervention

Data have inherent structures or patterns

Usually require a lot of data for training

- Factors that lead to the success of deep learning since 2010s
  - (1) Advances in high-performance computational power, especially GPUs
  - (2) Availability of large public datasets for training
  - (3) Improved network architecture designs and training strategies
  - (4) Accessibility of code and toolboxes for training deep neural networks

MRI reconstruction tasks

- Different MRI reconstruction tasks:
  - (1) Reconstruction from undersampled data
    - To recover images from sub-Nyquist sampled measurements (e.g., from uniform undersampling, variable density undersampling, k-t undersampling)
  - (2) Image enhancement
    - To reduce noise in the images or improve image sharpness
  - (3) Image super-resolution
    - To increase image resolutions
  - (4) Artifacts reduction
    - To reduce specific types of artifacts from hardware imperfections, MRI physics or physiological constraints (e.g., EPI artifacts, motion artifacts)
  - and more….

MRI reconstruction tasks

- Conventionally, these reconstruction tasks are carried out with a “hand-crafted” model (either by observations, experiments or assumptions)
  - Example 1:
    - Observe redundancy in multi-coil data -> Construct a model to for the under-determined inverse problem -> Develop parallel imaging algorithms
  - Example 2:
    - Make assumptions on the underlying noise model in the MRI images -> Construct a signal model that includes the noise term -> Develop algorithms to suppress the noise

Quick review on image reconstruction model

- General image acquisition model: $y = Ax + n$
  - y: the acquired data in the sensor domain (e.g., k-space in MRI)
  - x: the image
  - n: additive noise
  - A: an operator which is modality dependent
    - For computed tomography (CT): A is Radon transform
    - For fully sampled Cartesian MRI: A is Fourier transform
    - For undersampled Cartesian MRI: A includes subsampling and Fourier transform
    - For non-Cartesian MRI: A is non-uniform Fourier transform

Image reconstruction model

- To solve an underdetermined inverse problem (e.g., in the case of undersampled MRI), constrained reconstruction methods have been popular

Image model

$$y = Ax + n$$

Constrained reconstruction optimization problem

$$argmin_x \quad L(x) = \underbrace{F(Ax, y)}_{\text{Consistency with sensor-domain data}} + \lambda \underbrace{\Phi(W, x)}_{\text{Regularization term encoding prior information}}$$

Image reconstruction model

- Deep learning uses information from a large dataset to learn a non-linear mapping.
- In the task for MRI reconstruction from undersampled data:

Non-linear neural network

$$Q = \mathscr{F}(P)$$

Images with reduced artifacts or fully sampled images/k-space data

Images or k-space data from undersampled measurements

How can DL help MRI reconstruction?

- “Hand-crafted” models may not capture all the complex factors involved in MRI reconstruction, as some effects are too intricate or nonlinear to be accurately represented.

- Deep learning can complete reconstruction tasks without manually constructing a signal model

DL MRI reconstruction

- It’s impossible to cover all aspects of deep learning for MRI reconstruction because it’s an active and rapid-changing research field.
- There are a lot of online resources and UCLA lectures on deep learning. We will only briefly introduce the classic convolutional neural networks.
- In this lecture, we will focus on:
    - Special considerations of MRI reconstruction compared to other computer vision tasks
    - Popular approaches for DL MRI reconstruction
    - Challenges of DL MRI reconstruction

Part 2: Introduction to classic convolutional neural networks

Convolution Neural Networks (ConvNet)

- ConvNet is one of the most popular deep learning networks for imaging tasks

- We will introduce several key components in ConvNet and show how it can be trained
  - Convolution layer
  - Pooling layer
  - Activation function
  - Loss function
  - Optimizer
  - Regularization
  - Batch normalization
  - …

Where it all started…

- LeNet-5[1]: one of the very first ConvNet architectures with back-propagation for handwritten digit recognition

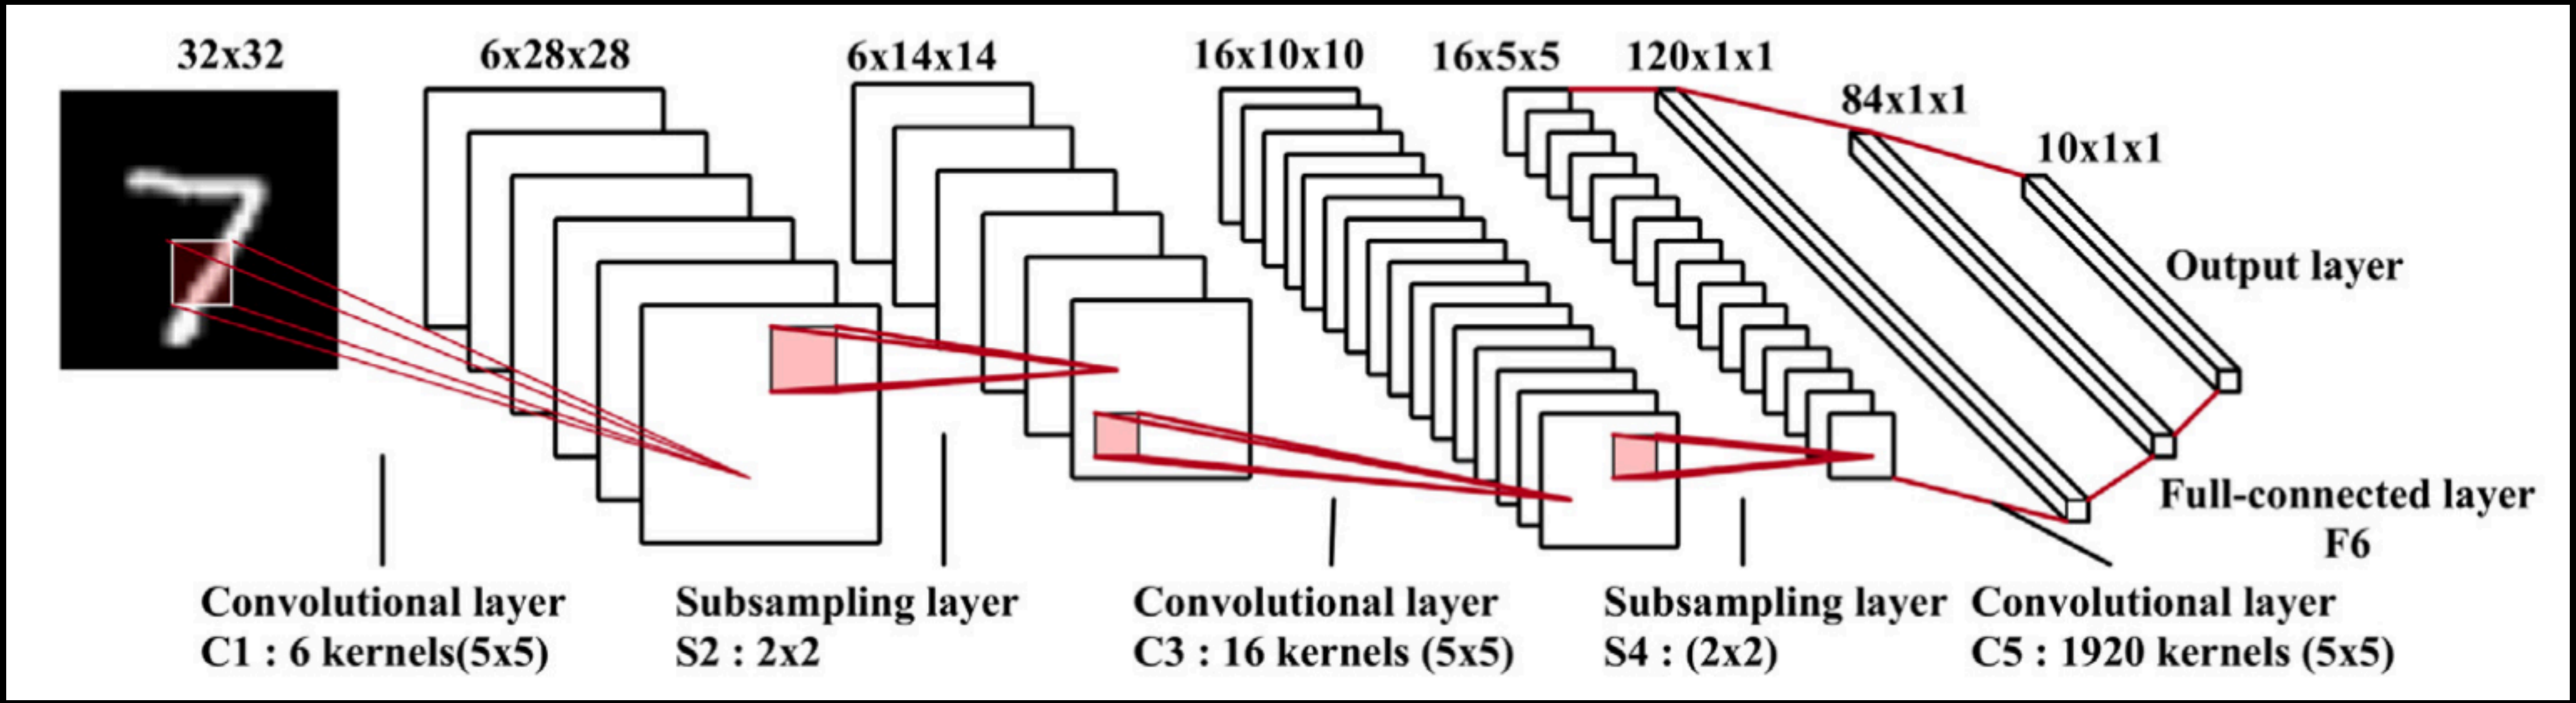

[1] LeCun et al., Proceedings of the IEEE, 1998

(Figure from: Gu et al., Pattern Recognition, 2018)

A glimpse of popular ConvNet models

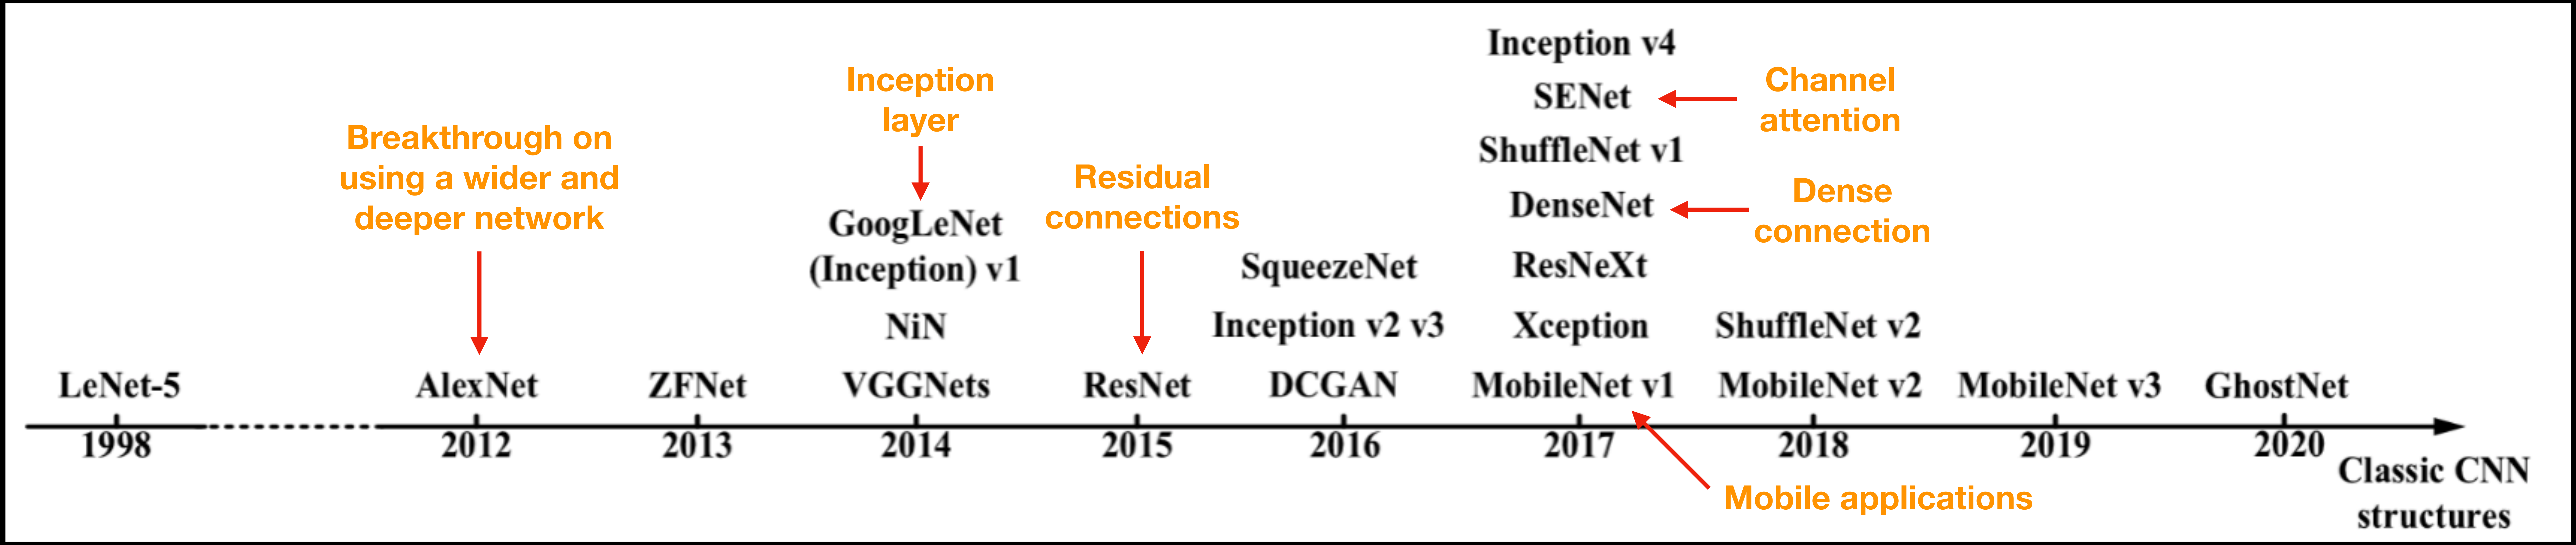

* Many of these ConvNet were first used in natural images (not medical images) and in a variety of tasks (e.g., classification, segmentation…)

(Figure from: Li et al., IEEE Trans Neural Netw Learn Syst 2022)

Popular ConvNet: U-Net

- The original U-Net was designed for medical image segmentation.
- It has been modified and applied in many DL-based MRI reconstruction tasks.

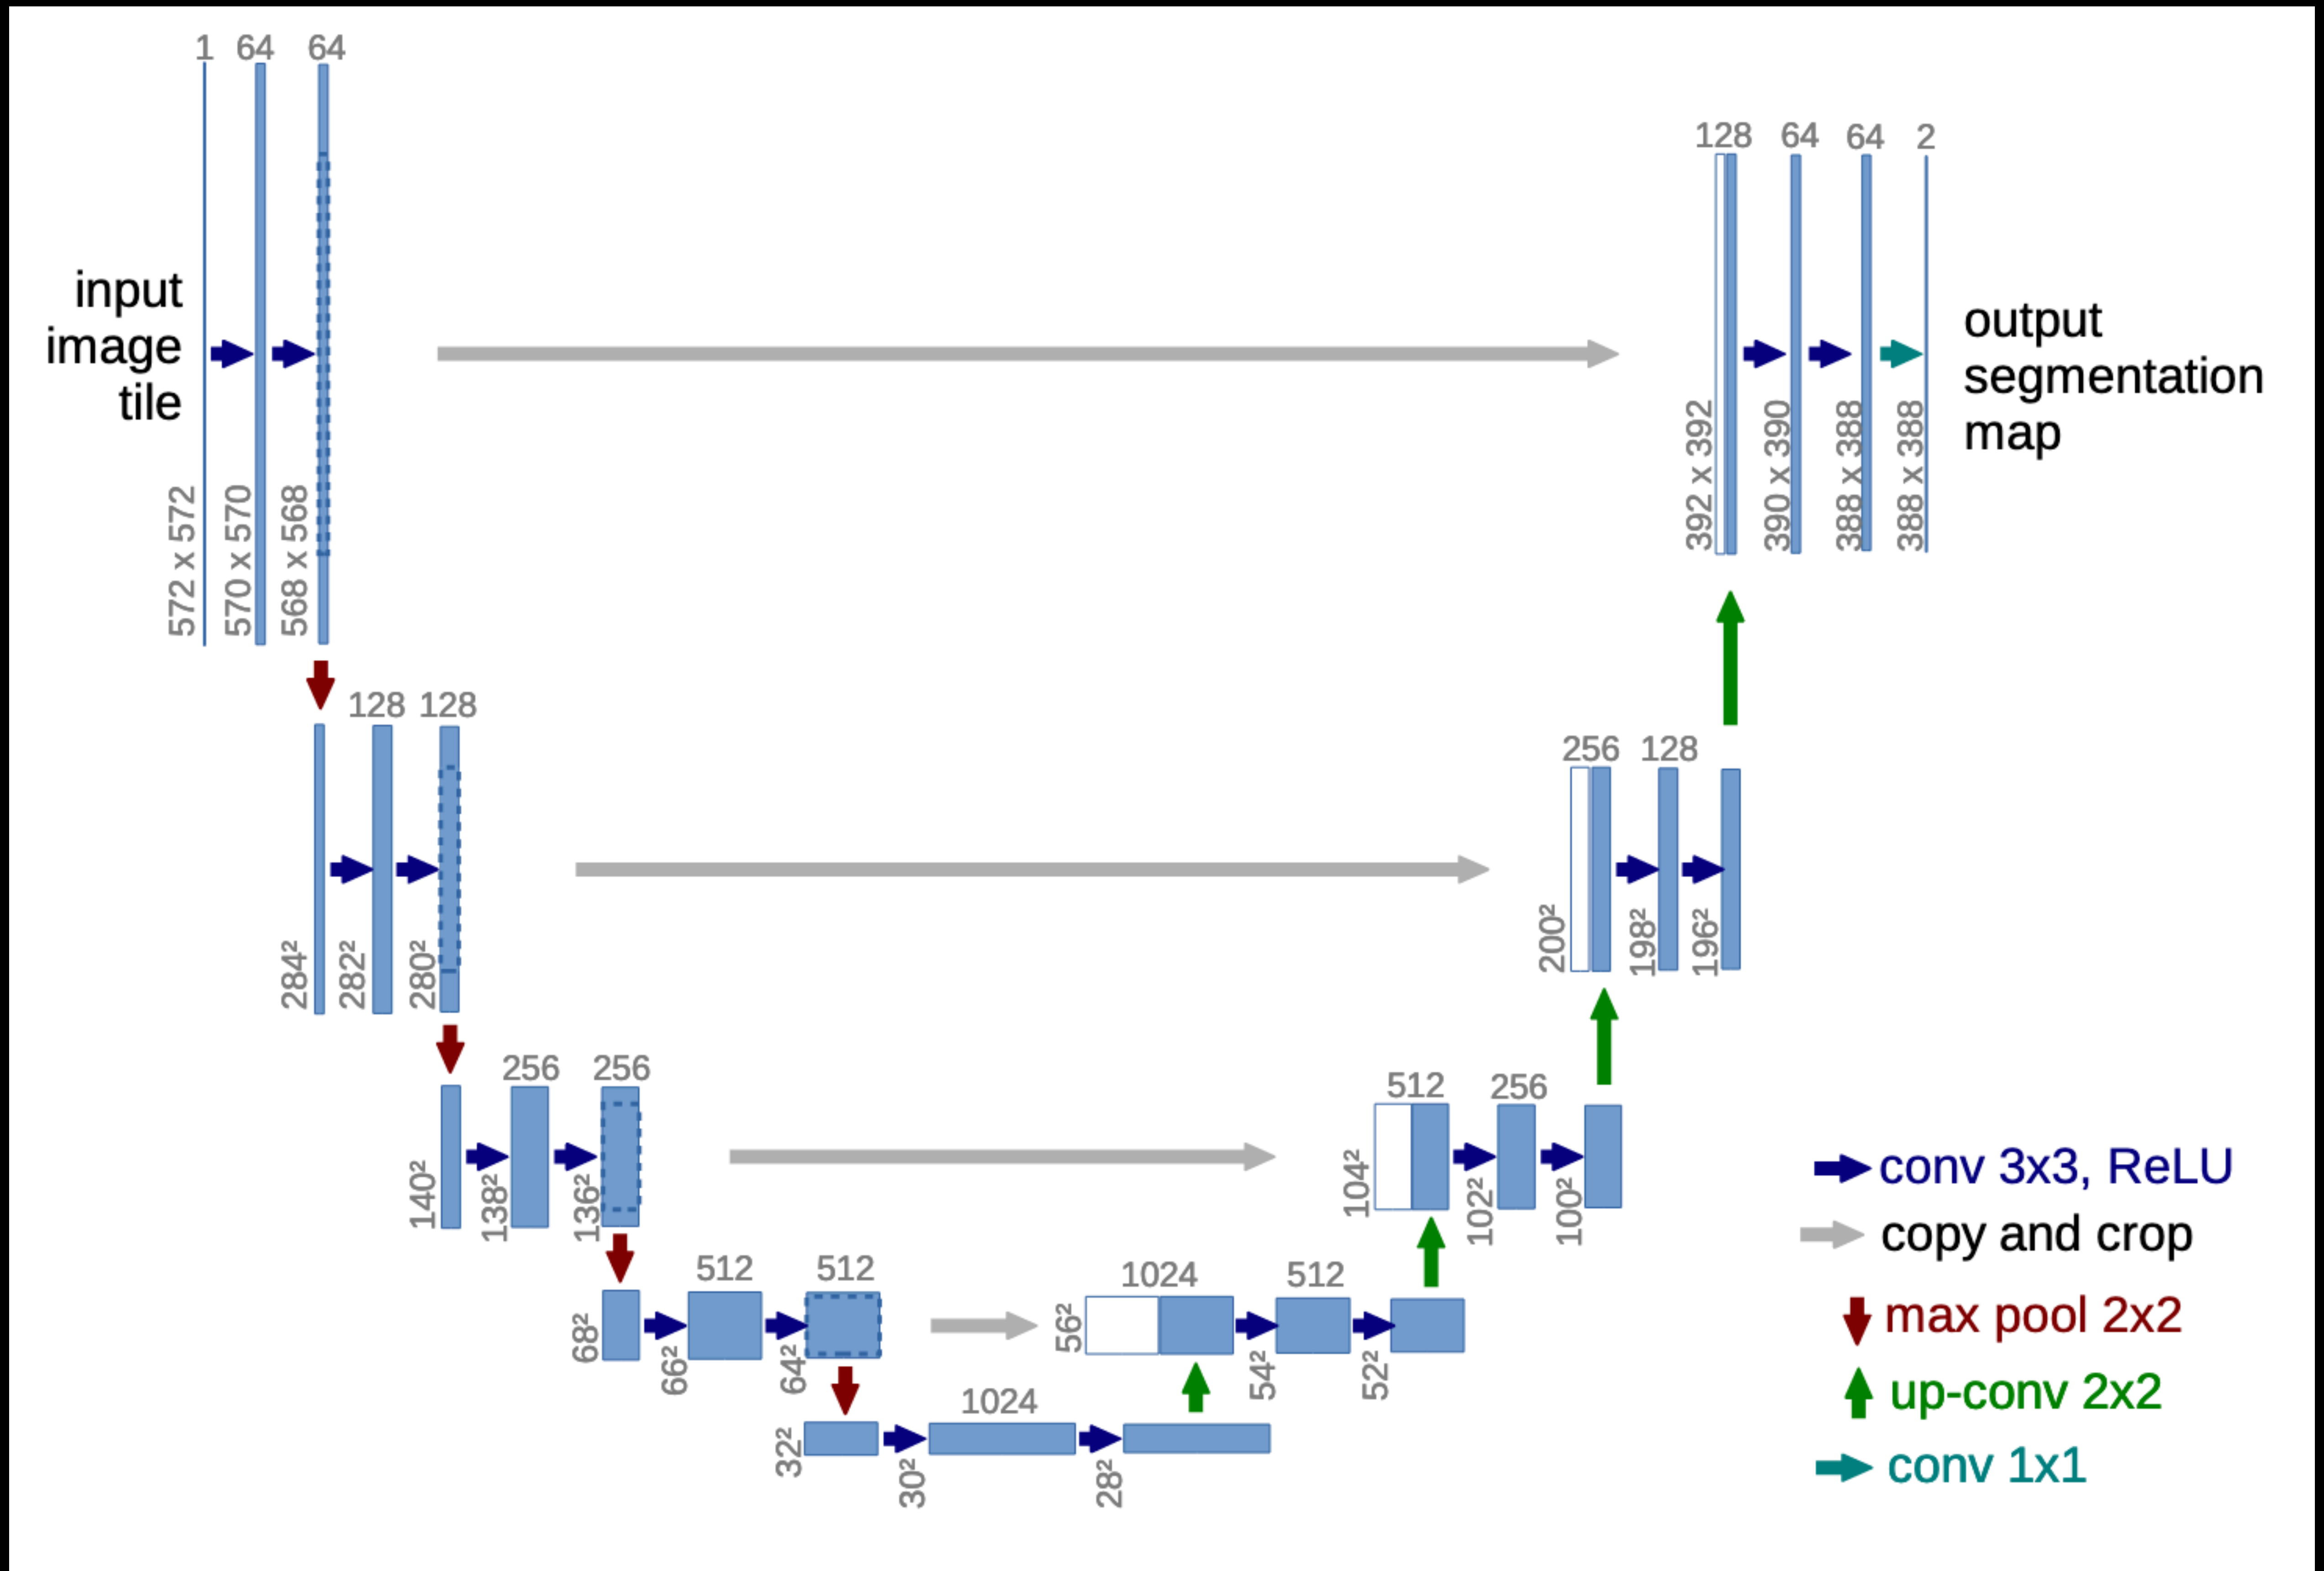

- Convolution at different levels
- Pooling layers
- Contracting and expansive paths
- Skipped connections

[1] Ronneberger et al., MICCAI, 2015

(Figure from: Ronneberger et al., MICCAI, 2015)

Convolutional layer

- Convolution operation: use a shared kernel to convolve with the entire image

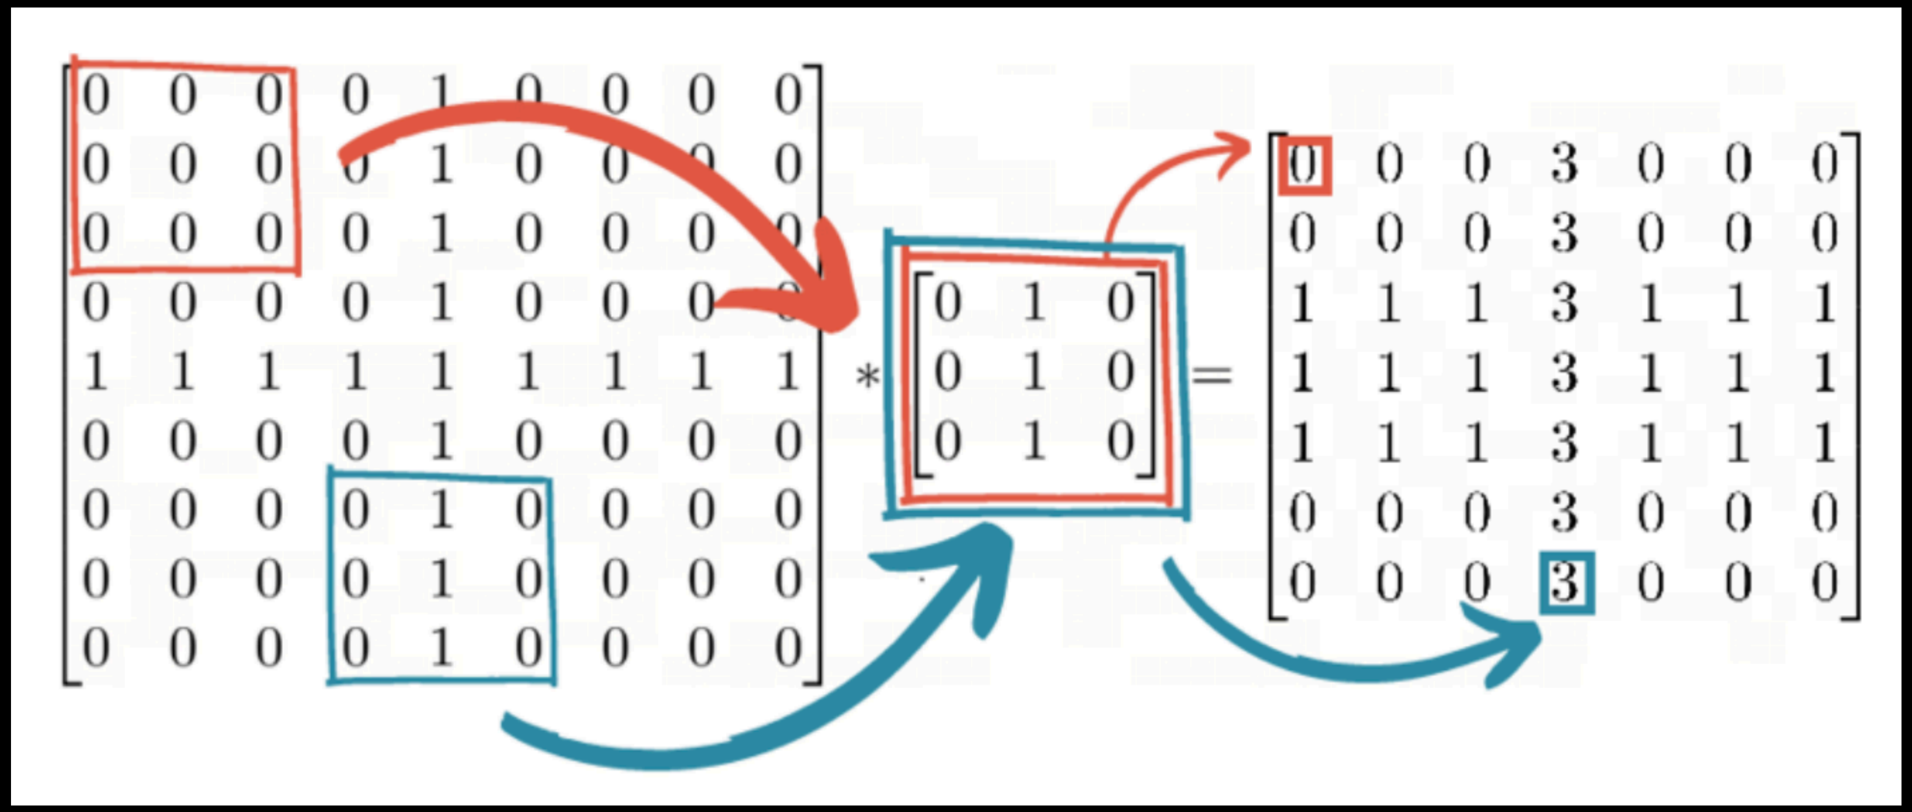

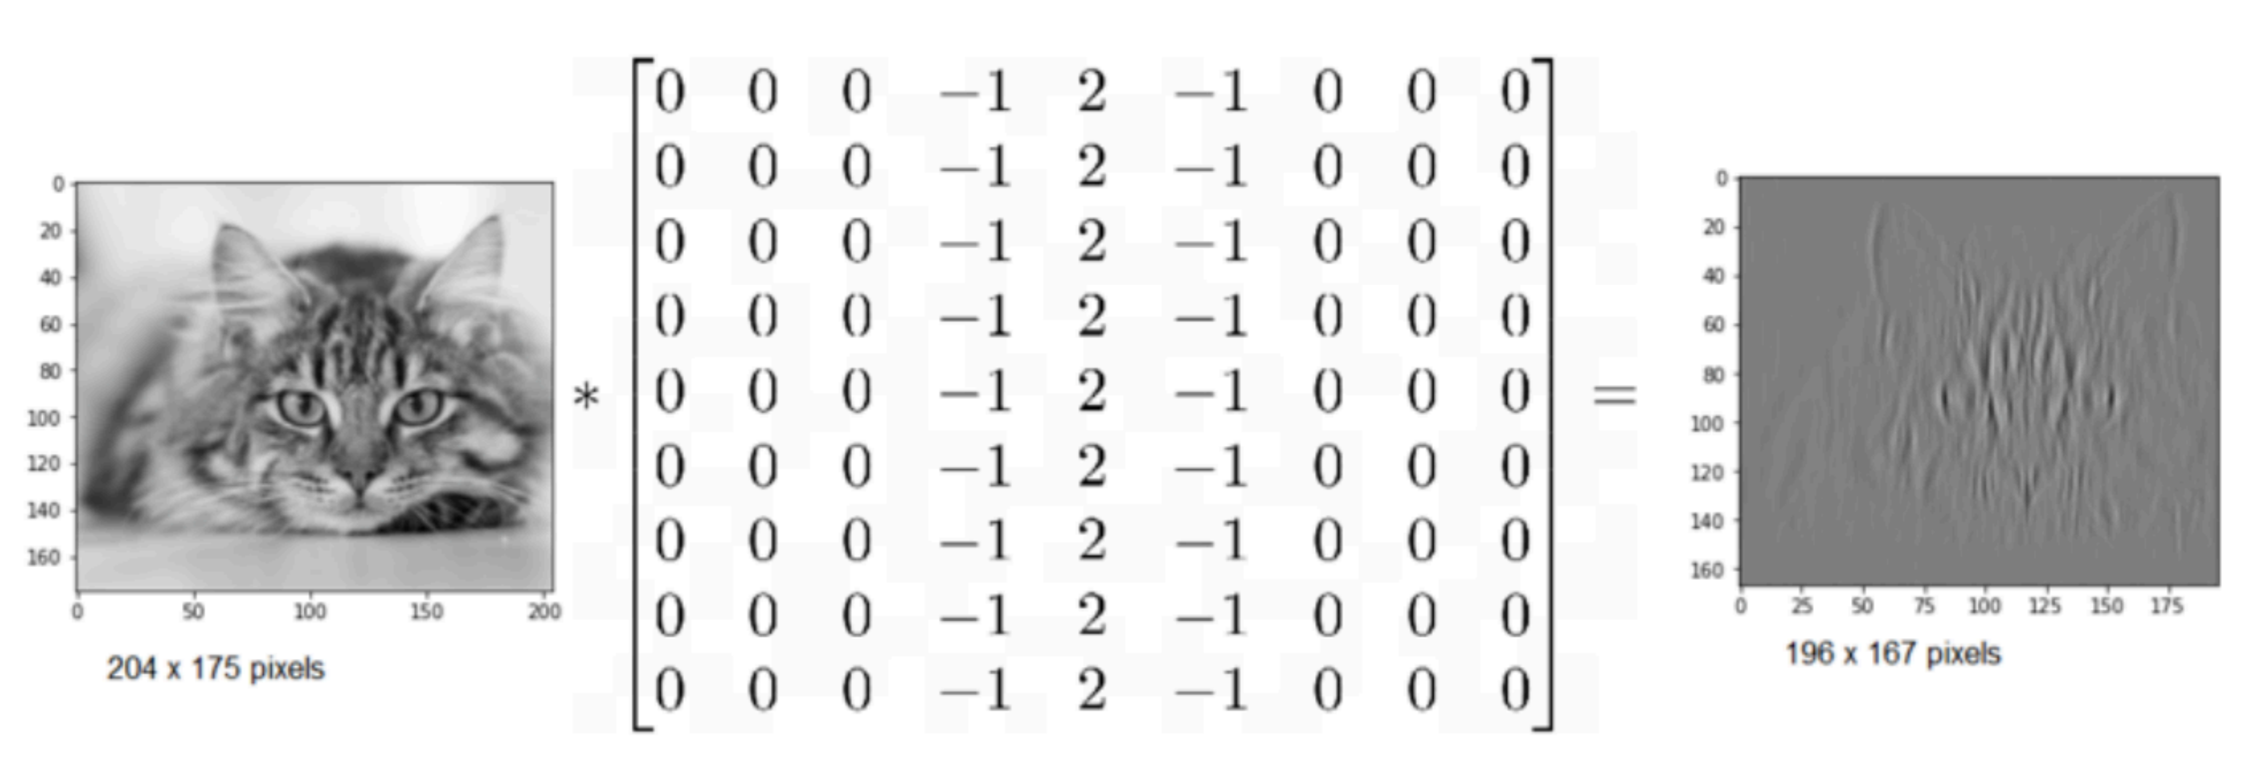

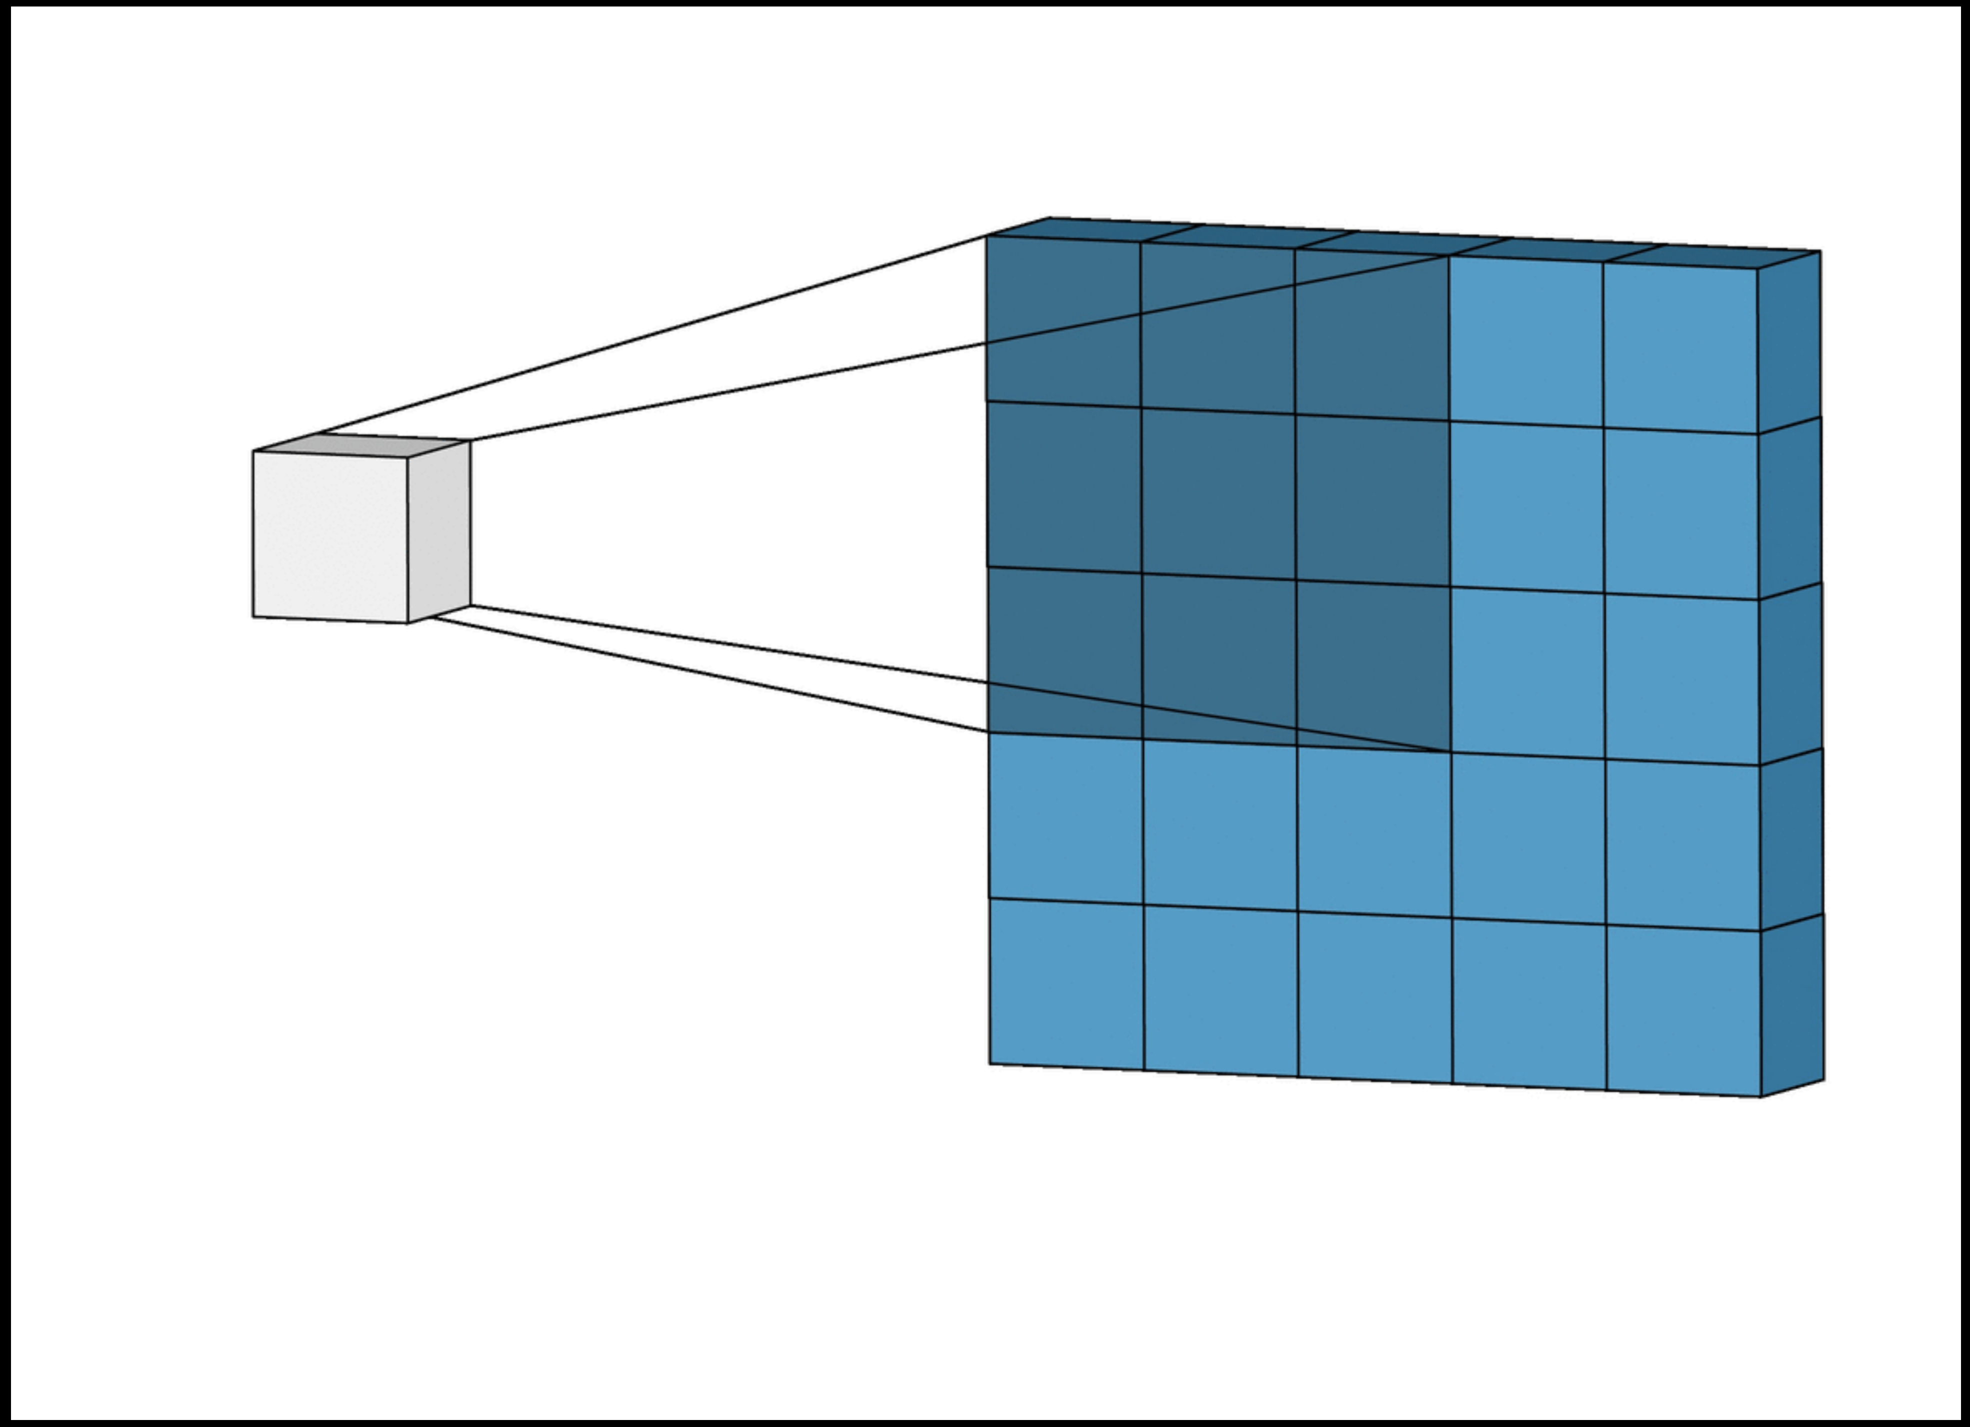

Figures from: https://deepai.org/machine-learning-glossary-and-terms/convolutional-neural-network

Convolutional layer

- Motivation of using a convolutional layer
- (1) Sparse interaction
    - Each pixel interacts with the kernel instead of all the other pixels.
- (2) Translational invariance
    - Some features are shared across the entire image.
    - The features do not change if the input is shifted.

Pooling layer

- Generate a summary of statistics with a reduced number of weights
  - Stride: the number of pixel shift for the next pooling operation

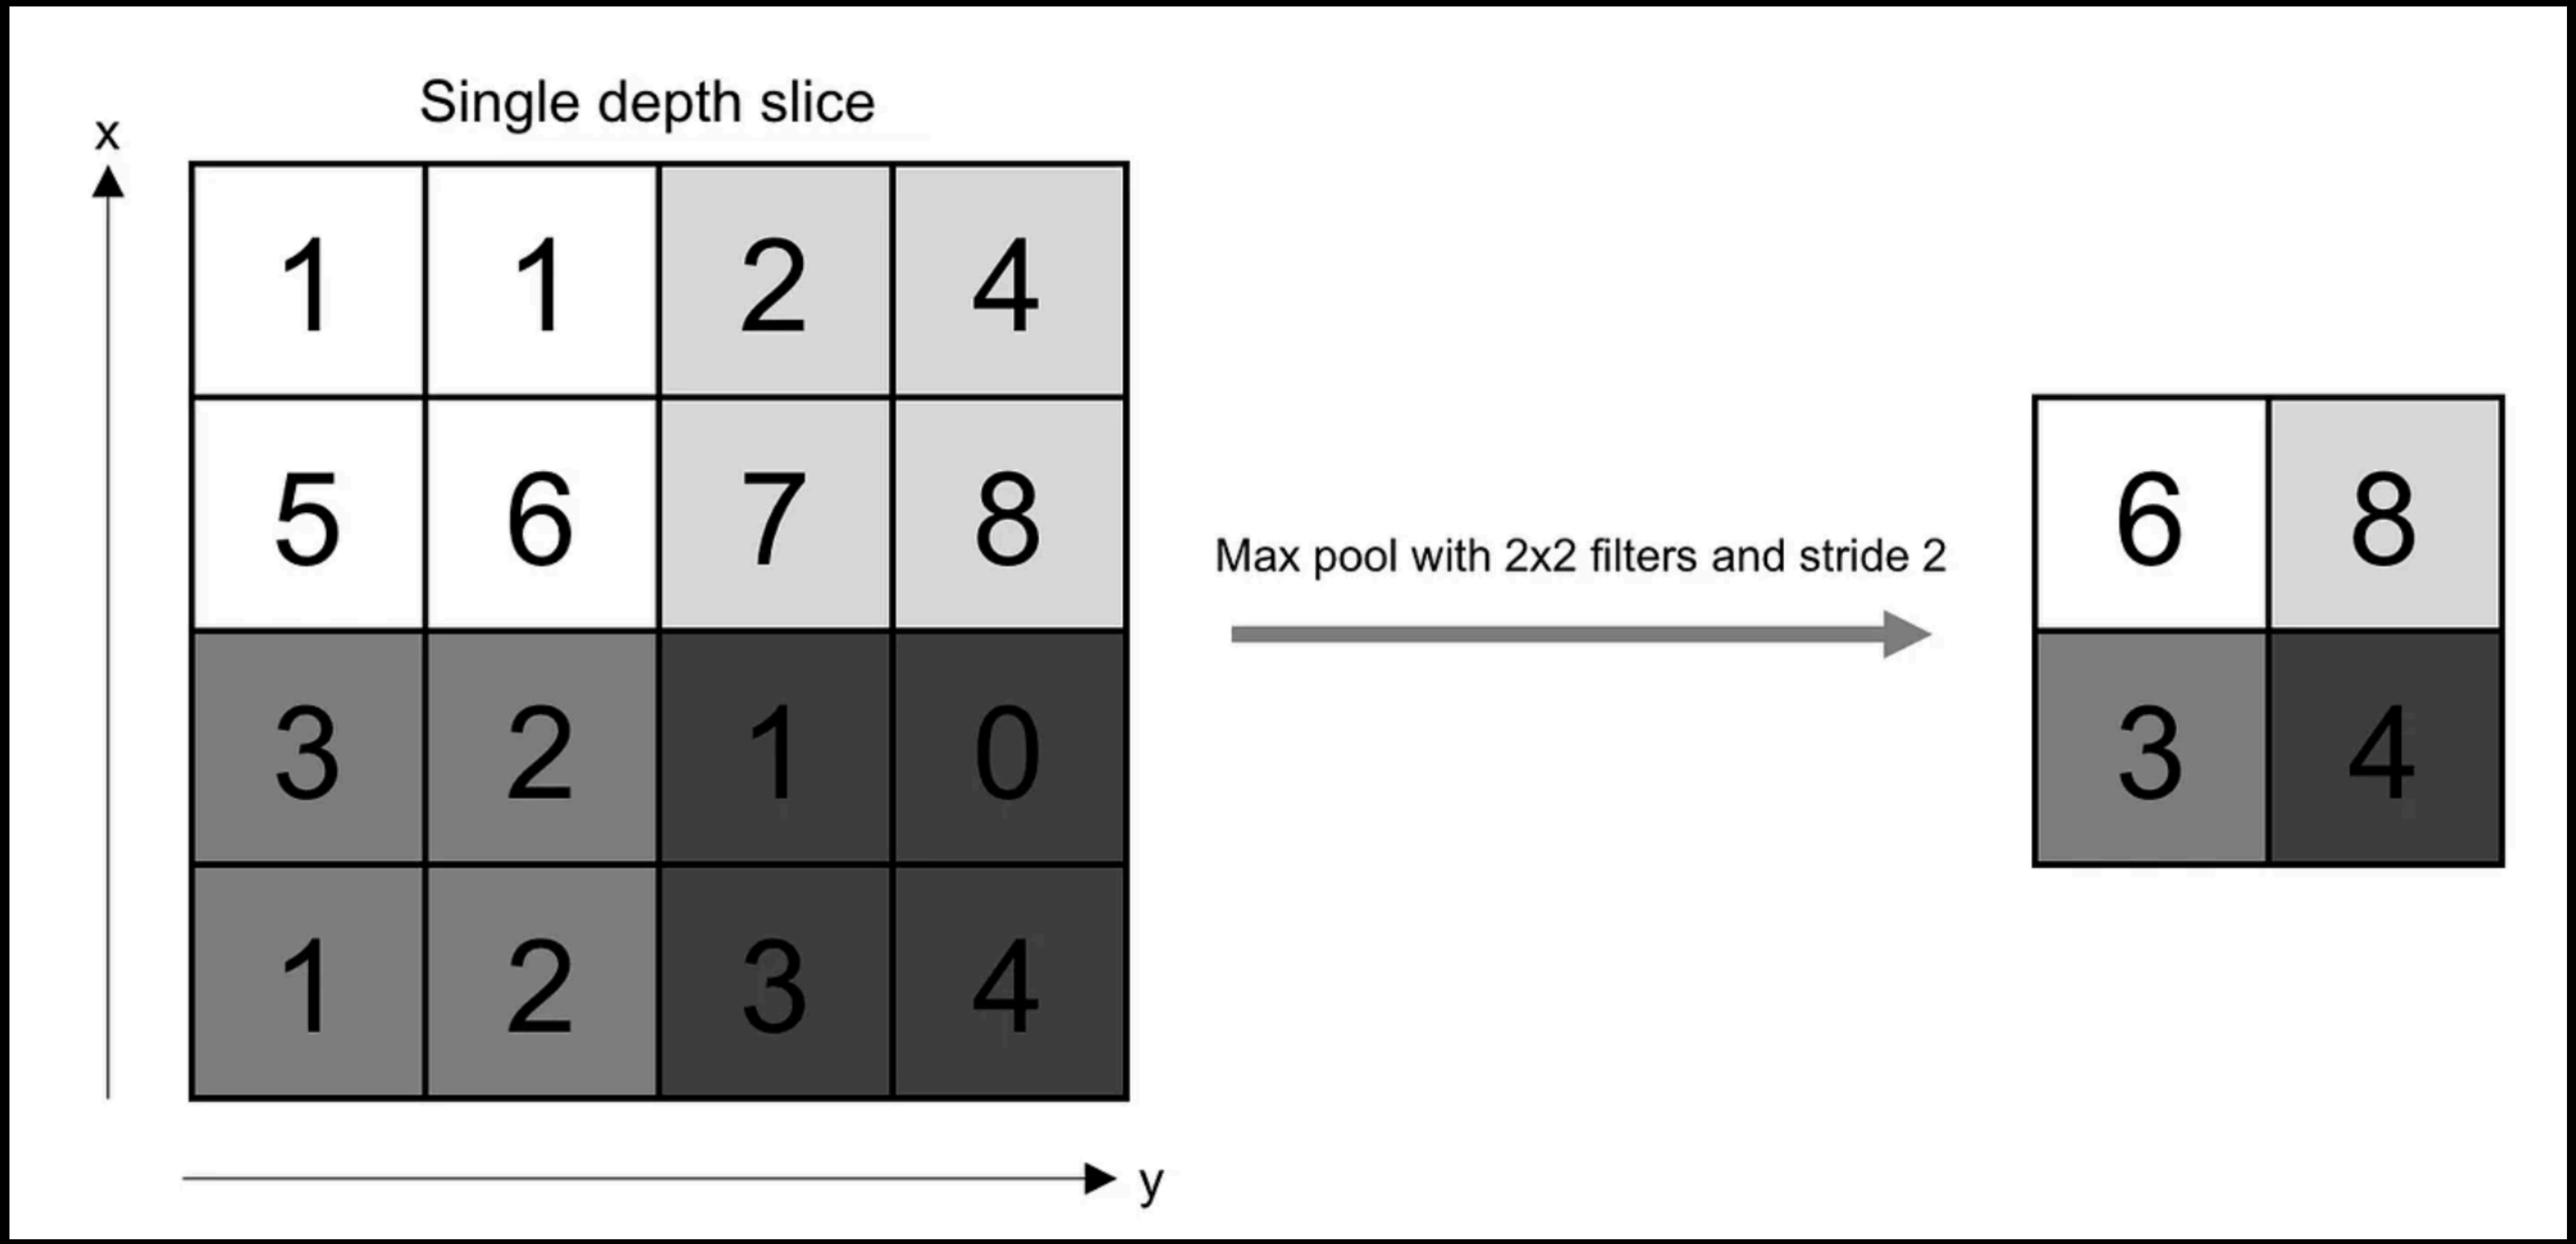

Activation function

- Convolution operation is linear. A stack of convolutional layers only generates a linear mapping process.
- Activation functions are used to introduce non-linearity to the network.
- ReLU (rectified linear unit): $f(a) = max(0,a)$

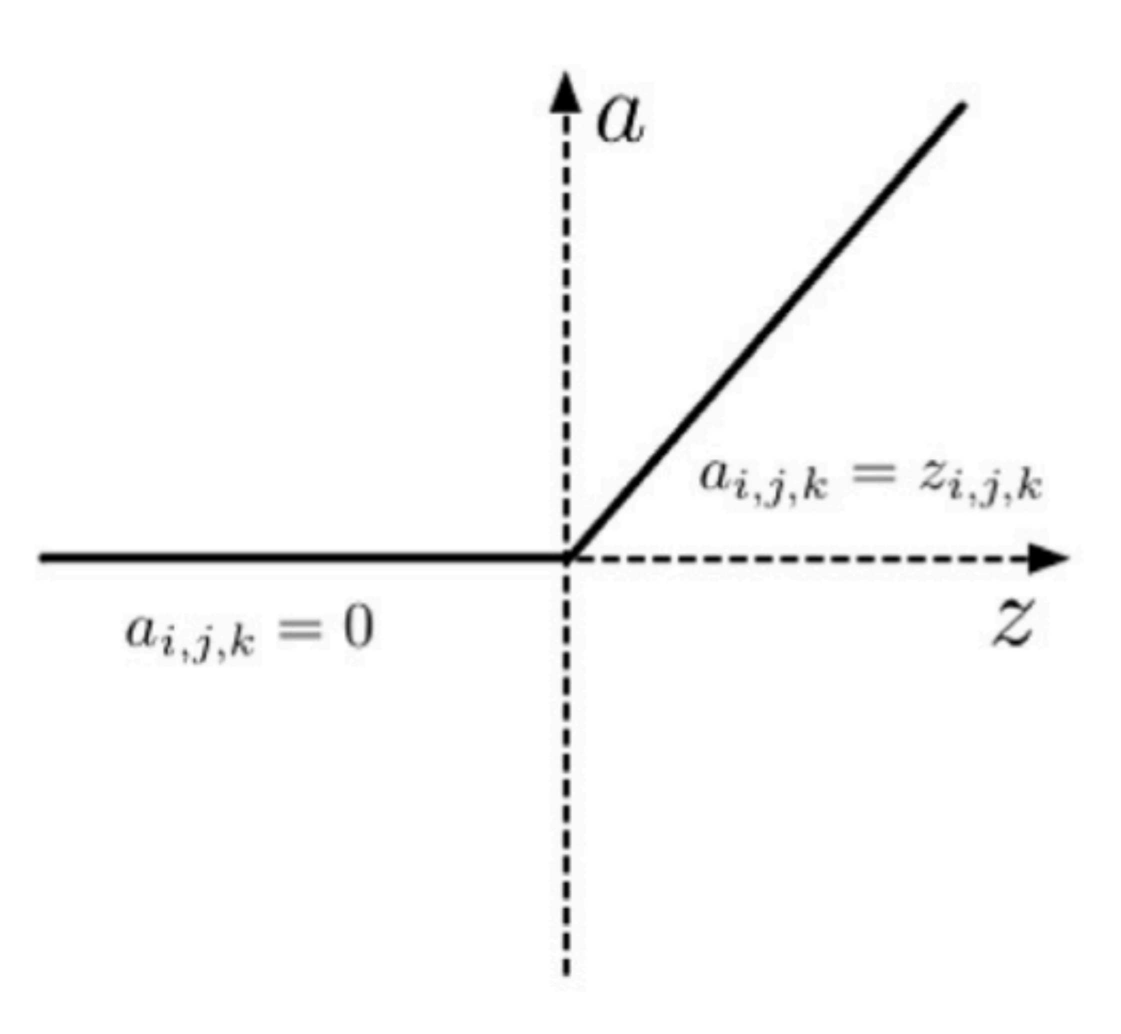

(a) ReLU

(Figures from: Gu et al., Pattern Recognition, 2018)

Improvements on activation functions

- ReLU has zero gradient when the node is not active
  - Different activation functions have been proposed to alleviate the problem

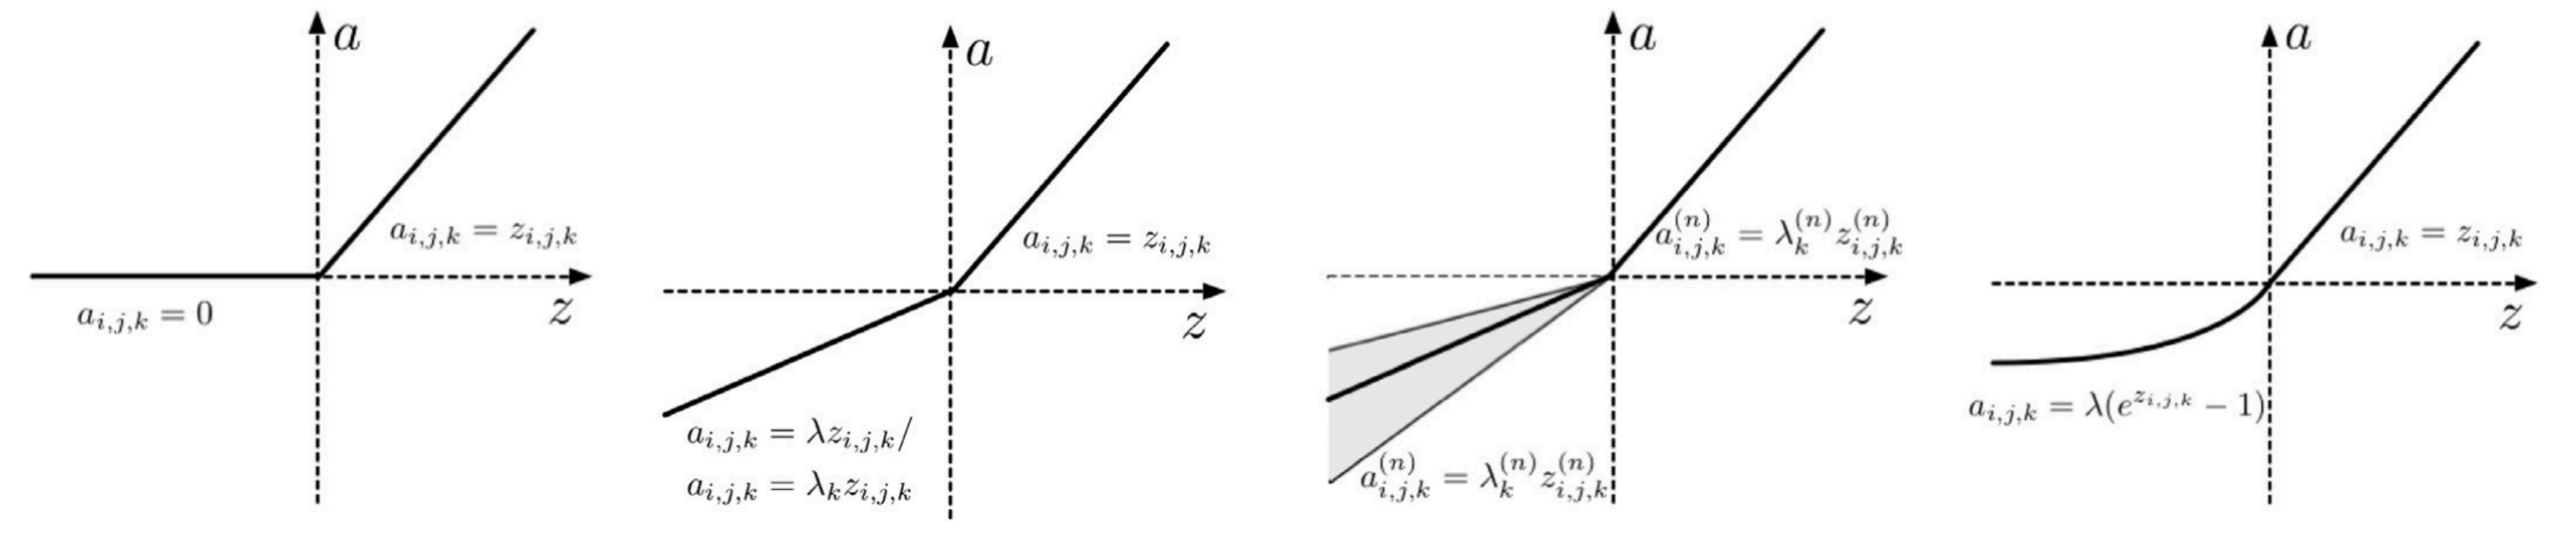

(a) ReLU (b) LReLU/PReLU (c) RReLU (d) ELU

(Figures from: Gu et al., Pattern Recognition, 2018)

Loss function

- We need an objective criteria to tell the network how well it performs.
- The overall network is trained to minimize the loss function.

- Loss functions for image reconstruction:
  - MSE loss / L2 loss
  - L1 loss
  - SSIM (structural similarity index measure) loss
  - perceptual loss
  - GAN (generative adversarial network) loss
  - …

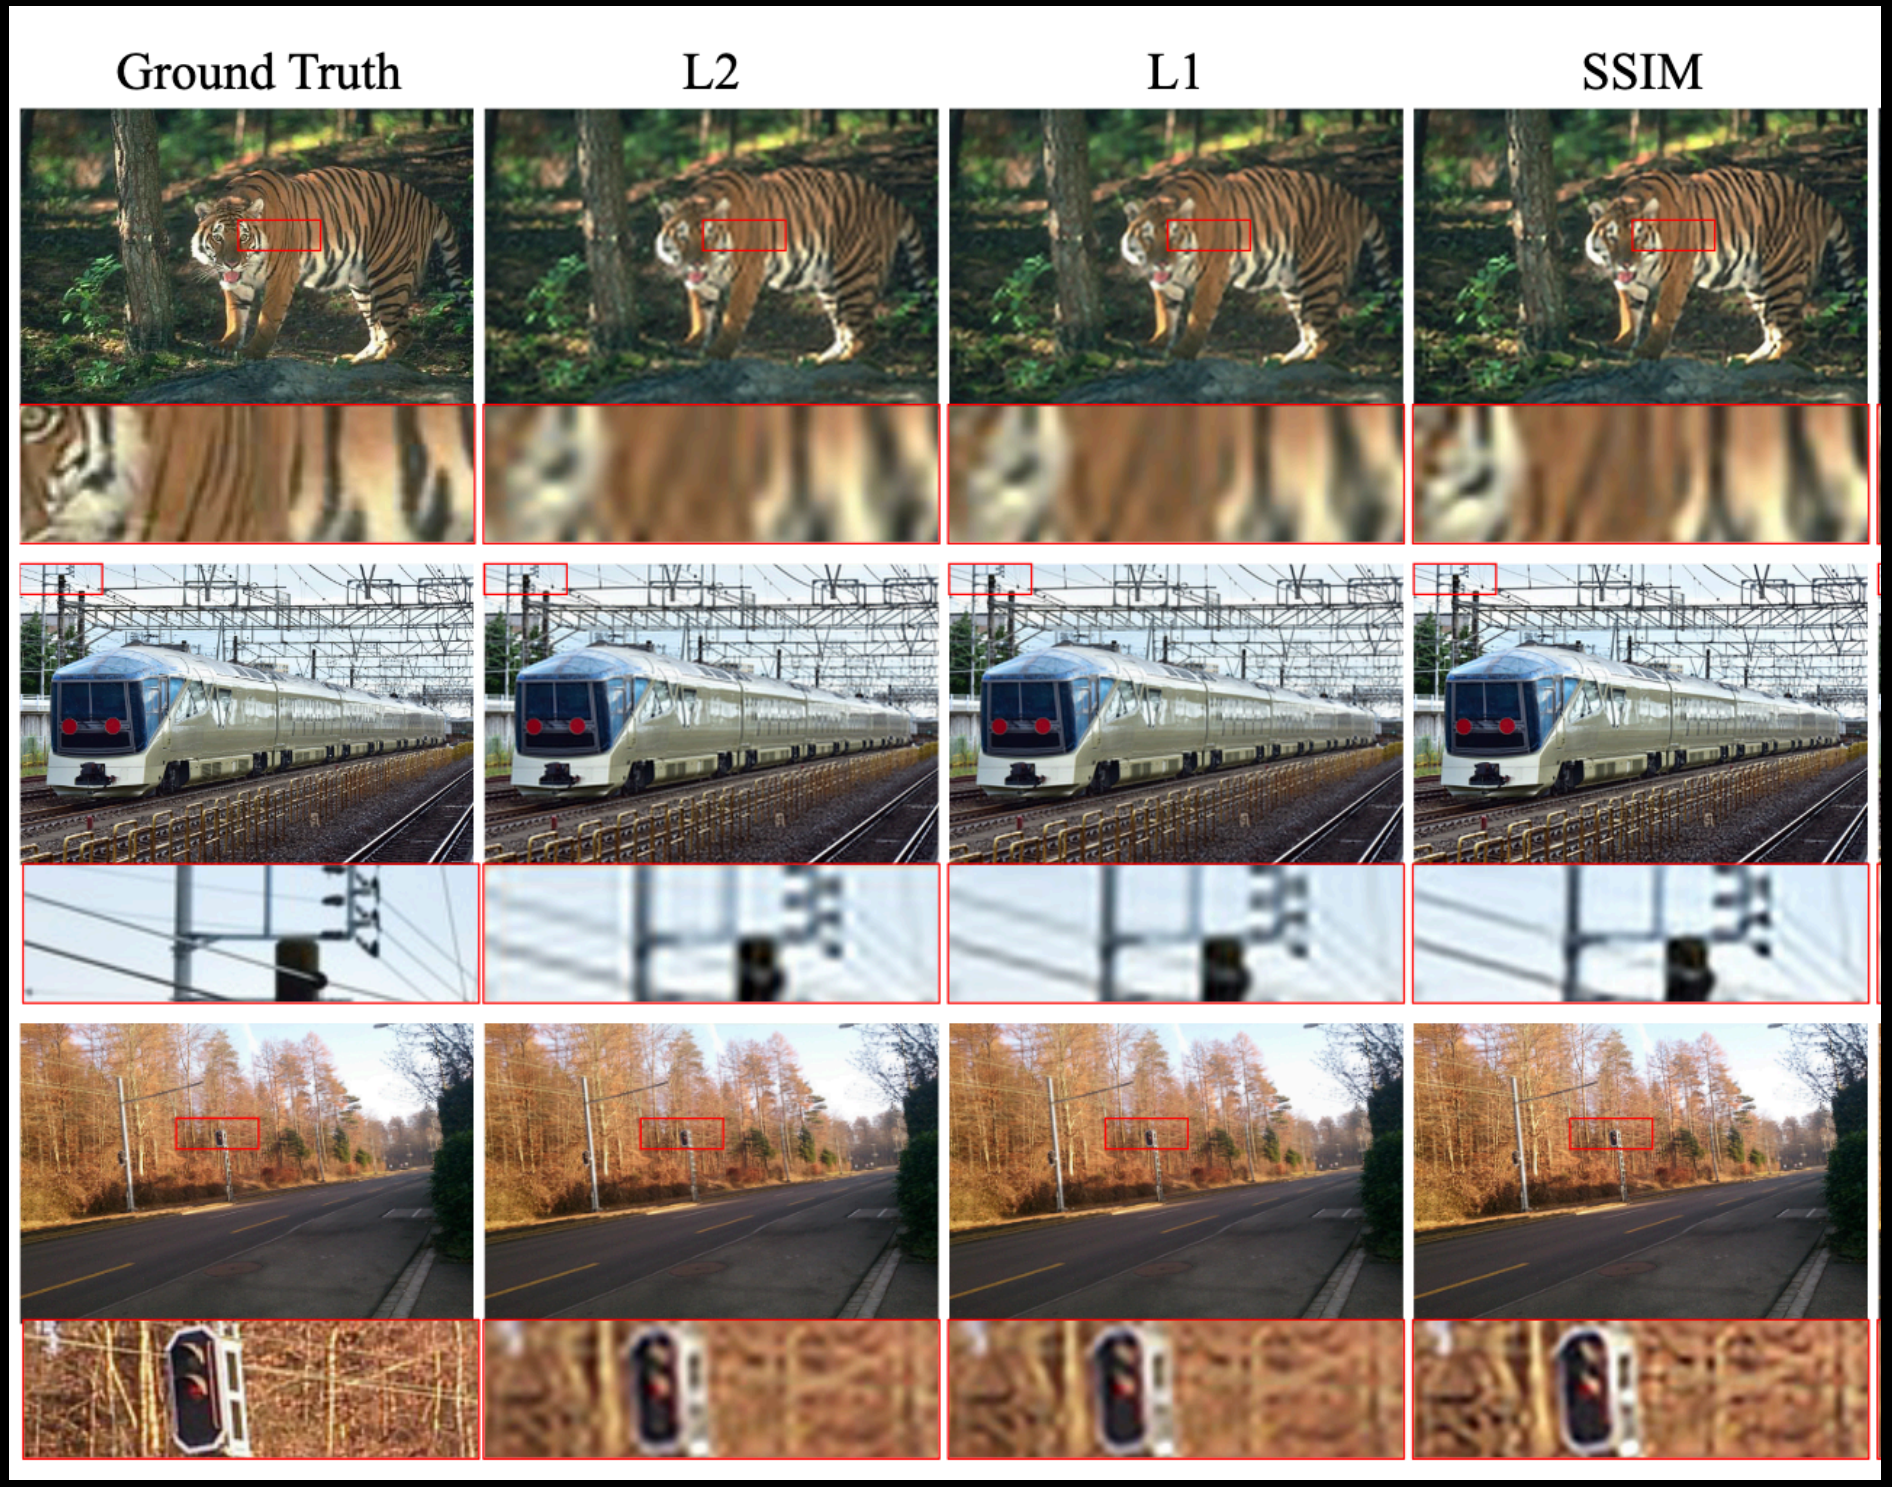

(Figure from: Mustafa et al., WACV, 2022)

Optimizer

- Algorithms used to update network parameters for loss minimization
  - Gradient descent
  - Stochastic gradient descent
    - Replace the actual gradient calculation from the entire dataset by using a randomly selected subset
    - “Batch size” can be used to refer to the number of training samples in one forward/backward pass

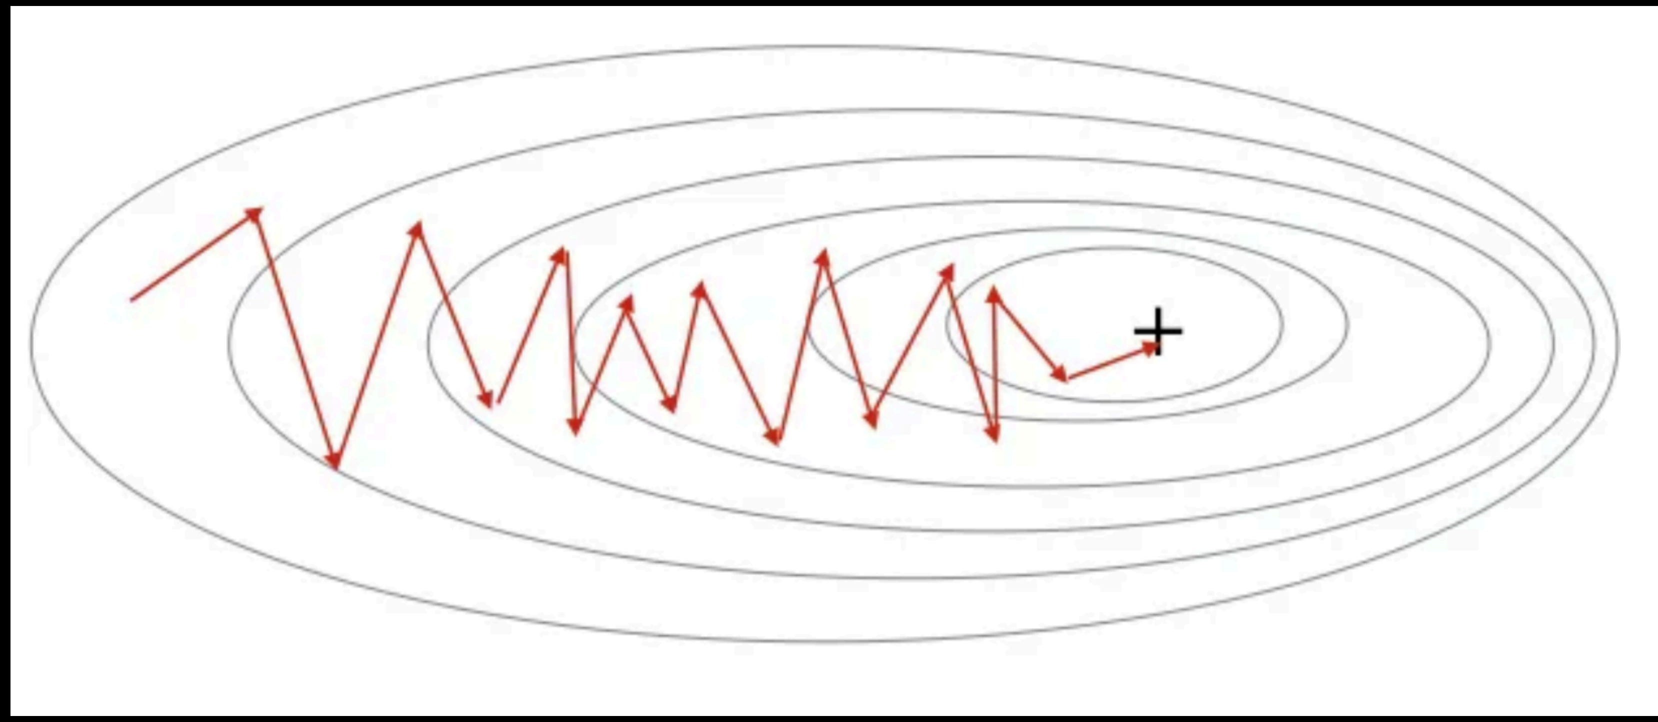

Stochastic gradient descent

(Figure from: https://medium.com/mlearning-ai/optimizers-in-deep-learning-7bf81fed78a0)

Optimizer

- To avoid local minimum problems, there are more adaptive optimizers that incorporate a "momentum" idea that use previous gradient information
    - Adagrad
    - RMSProp
    - Adam
    - …

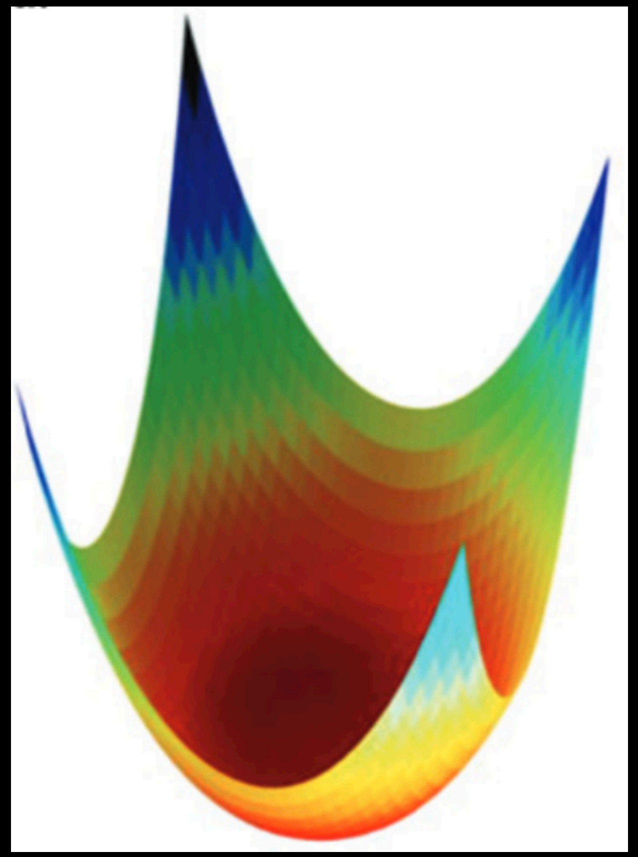

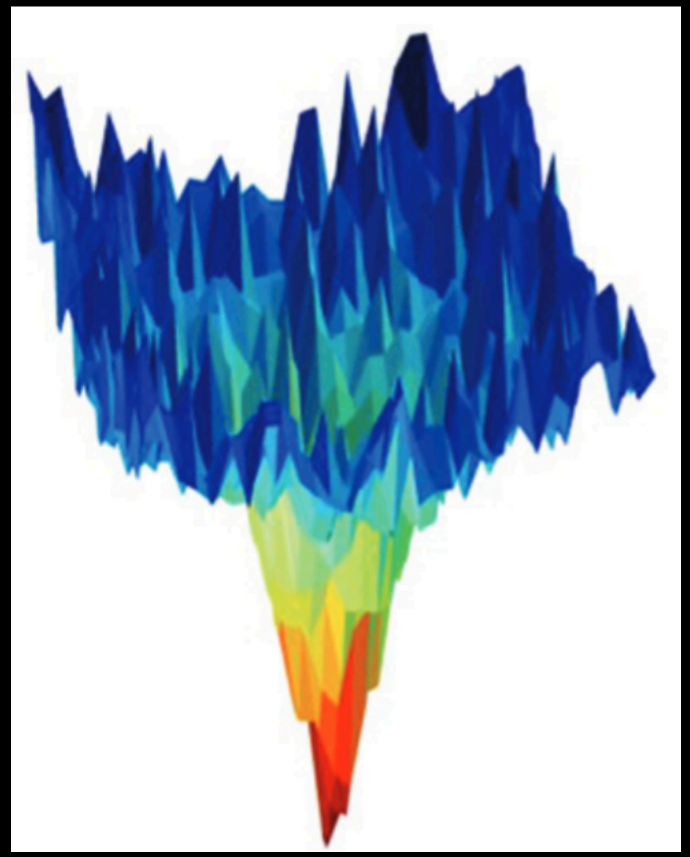

- Luckily, there are many optimizers already implemented in popular deep learning frameworks (PyTorch, TensorFlow…)

(Figures from: Cheng et al., RSNA, 2021)

Optimizer

- Find a suitable learning rate

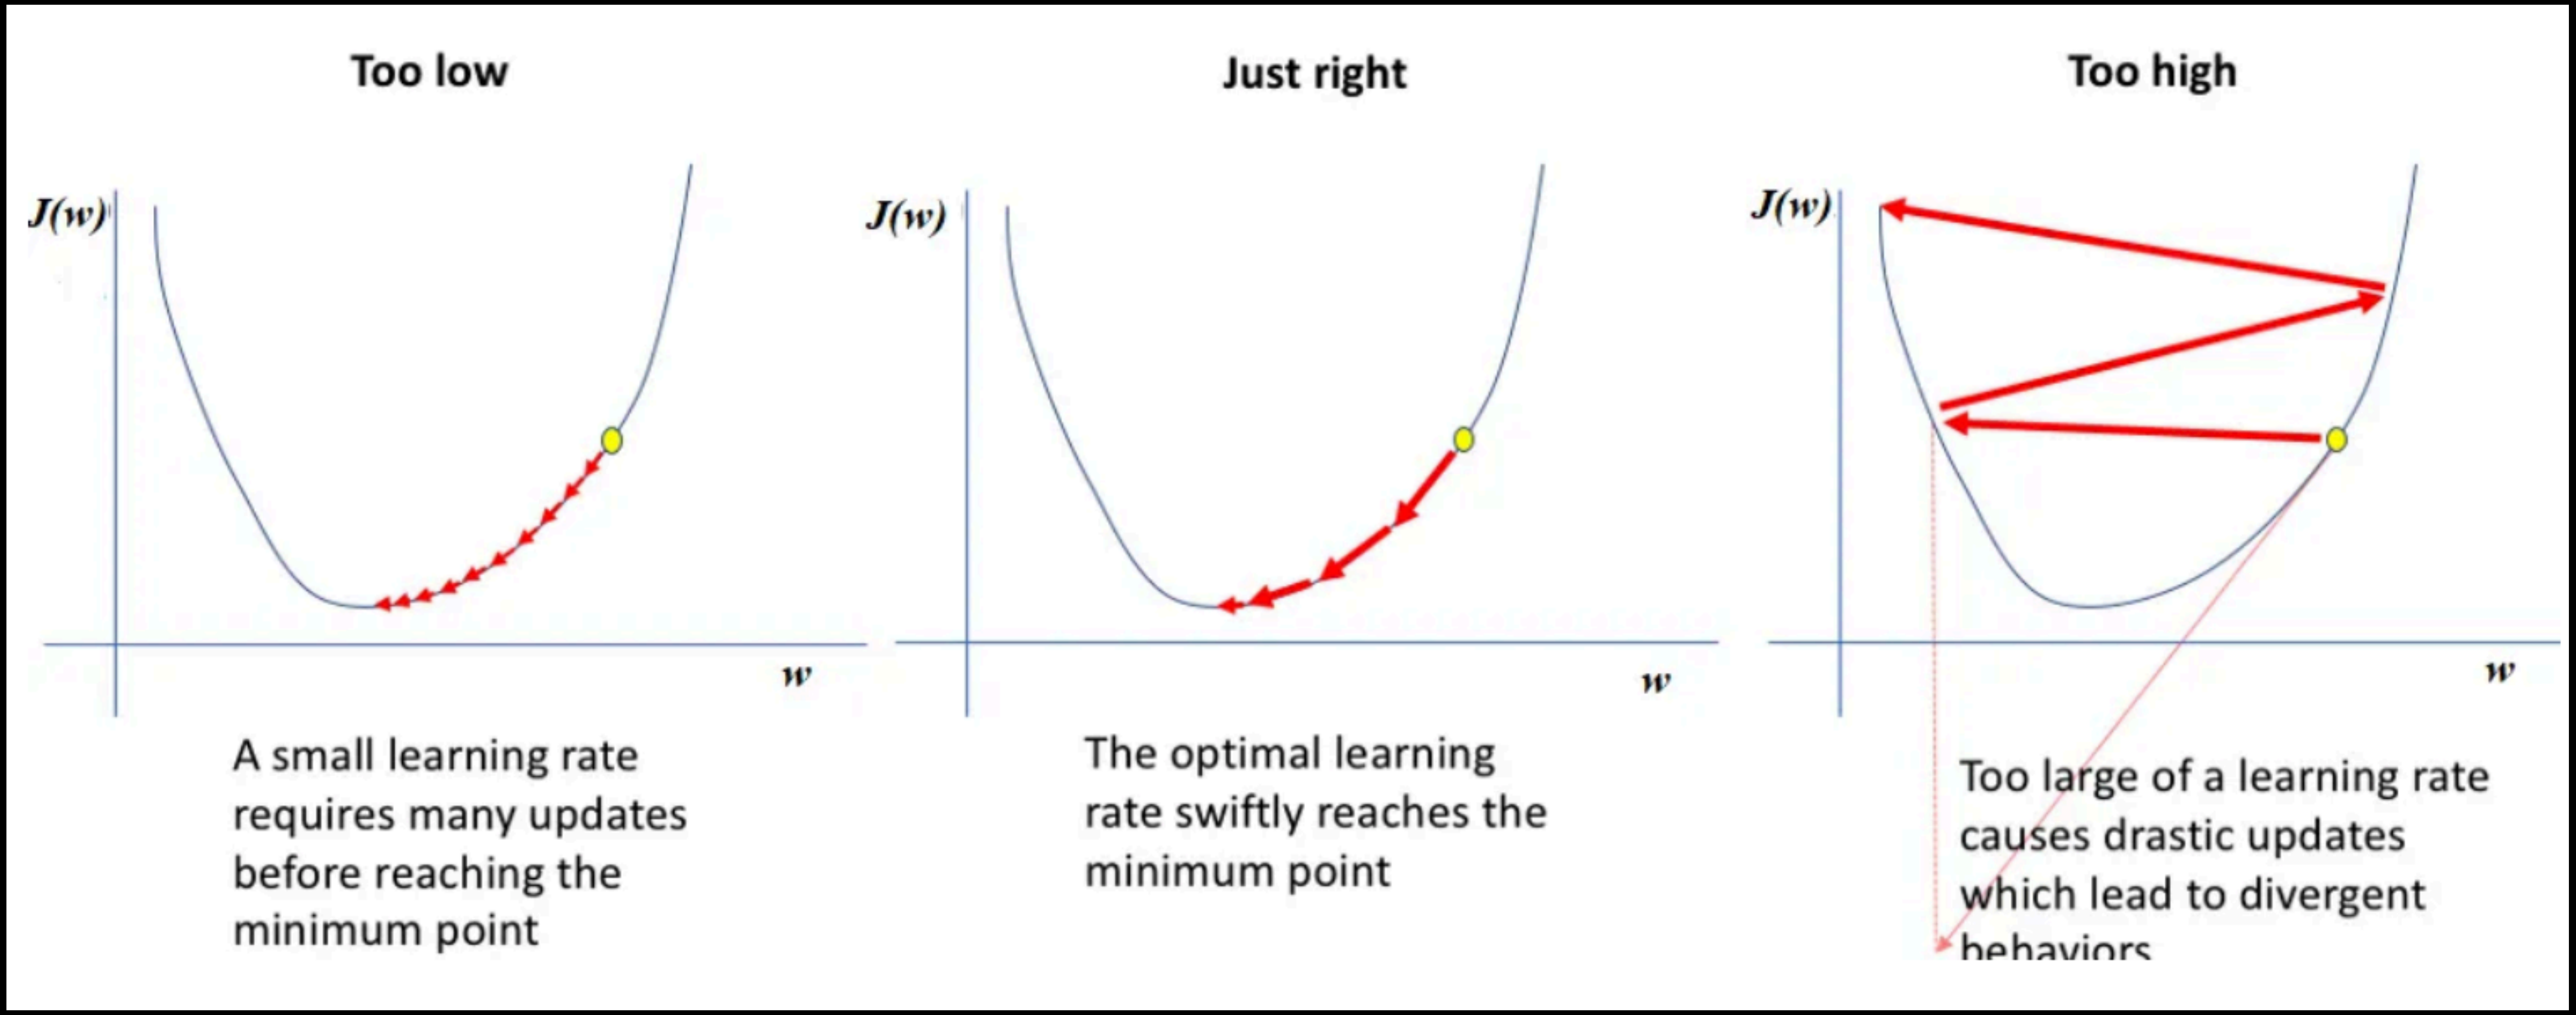

(Figure from: https://towardsdatascience.com/hyper-parameter-tuning-techniques-in-deep-learning-4dad592c63c8)

Back-propagation

- Once we know about the gradient, back-propagation is usually used as an efficient way to update the network's trainable parameters.

Back-propagation

One layer:

$$Q = \overset{\text{Nonlinear function}}{\phi}(\overset{\text{Weight}}{w}\overset{\text{Input image}}{P} + \overset{\text{Bias}}{b})$$

Network with deep layers:

$$Q = \mathscr{F}(P) = \phi(w_n \dots \underbrace{\phi(w_2 \underbrace{\phi(w_1 P + b_1)}_{\text{First layer}} + b_2)}_{\text{Second layer}} \dots + b_n)$$

Using chain rule
To calculate derivatives

$$Q = \mathscr{F}(g(P)) \longrightarrow \frac{\partial Q}{\partial P} = \frac{\partial \mathscr{F}(P)}{\partial P} = \mathscr{F}'(g(x))\,.\,g'(x)$$

- Luckily, back-propagation can be done easily using popular deep learning frameworks (PyTorch, TensorFlow…)

Regularization

- Regularization is any modification we make to a learning algorithm that is intended to reduce its generalization error but not its training error[1].

- Examples:
  - Include prior knowledge
  - Apply some constraints on the parameters in the loss function
  - Data augmentation: image flipping, rotation…
  - Dropout
  - …

[1] Goodfellow et al., *Deep learning*. MIT press, 2016

Regularization

- Dropout[1]
  - Randomly “turn off” some of the weights during the training process.

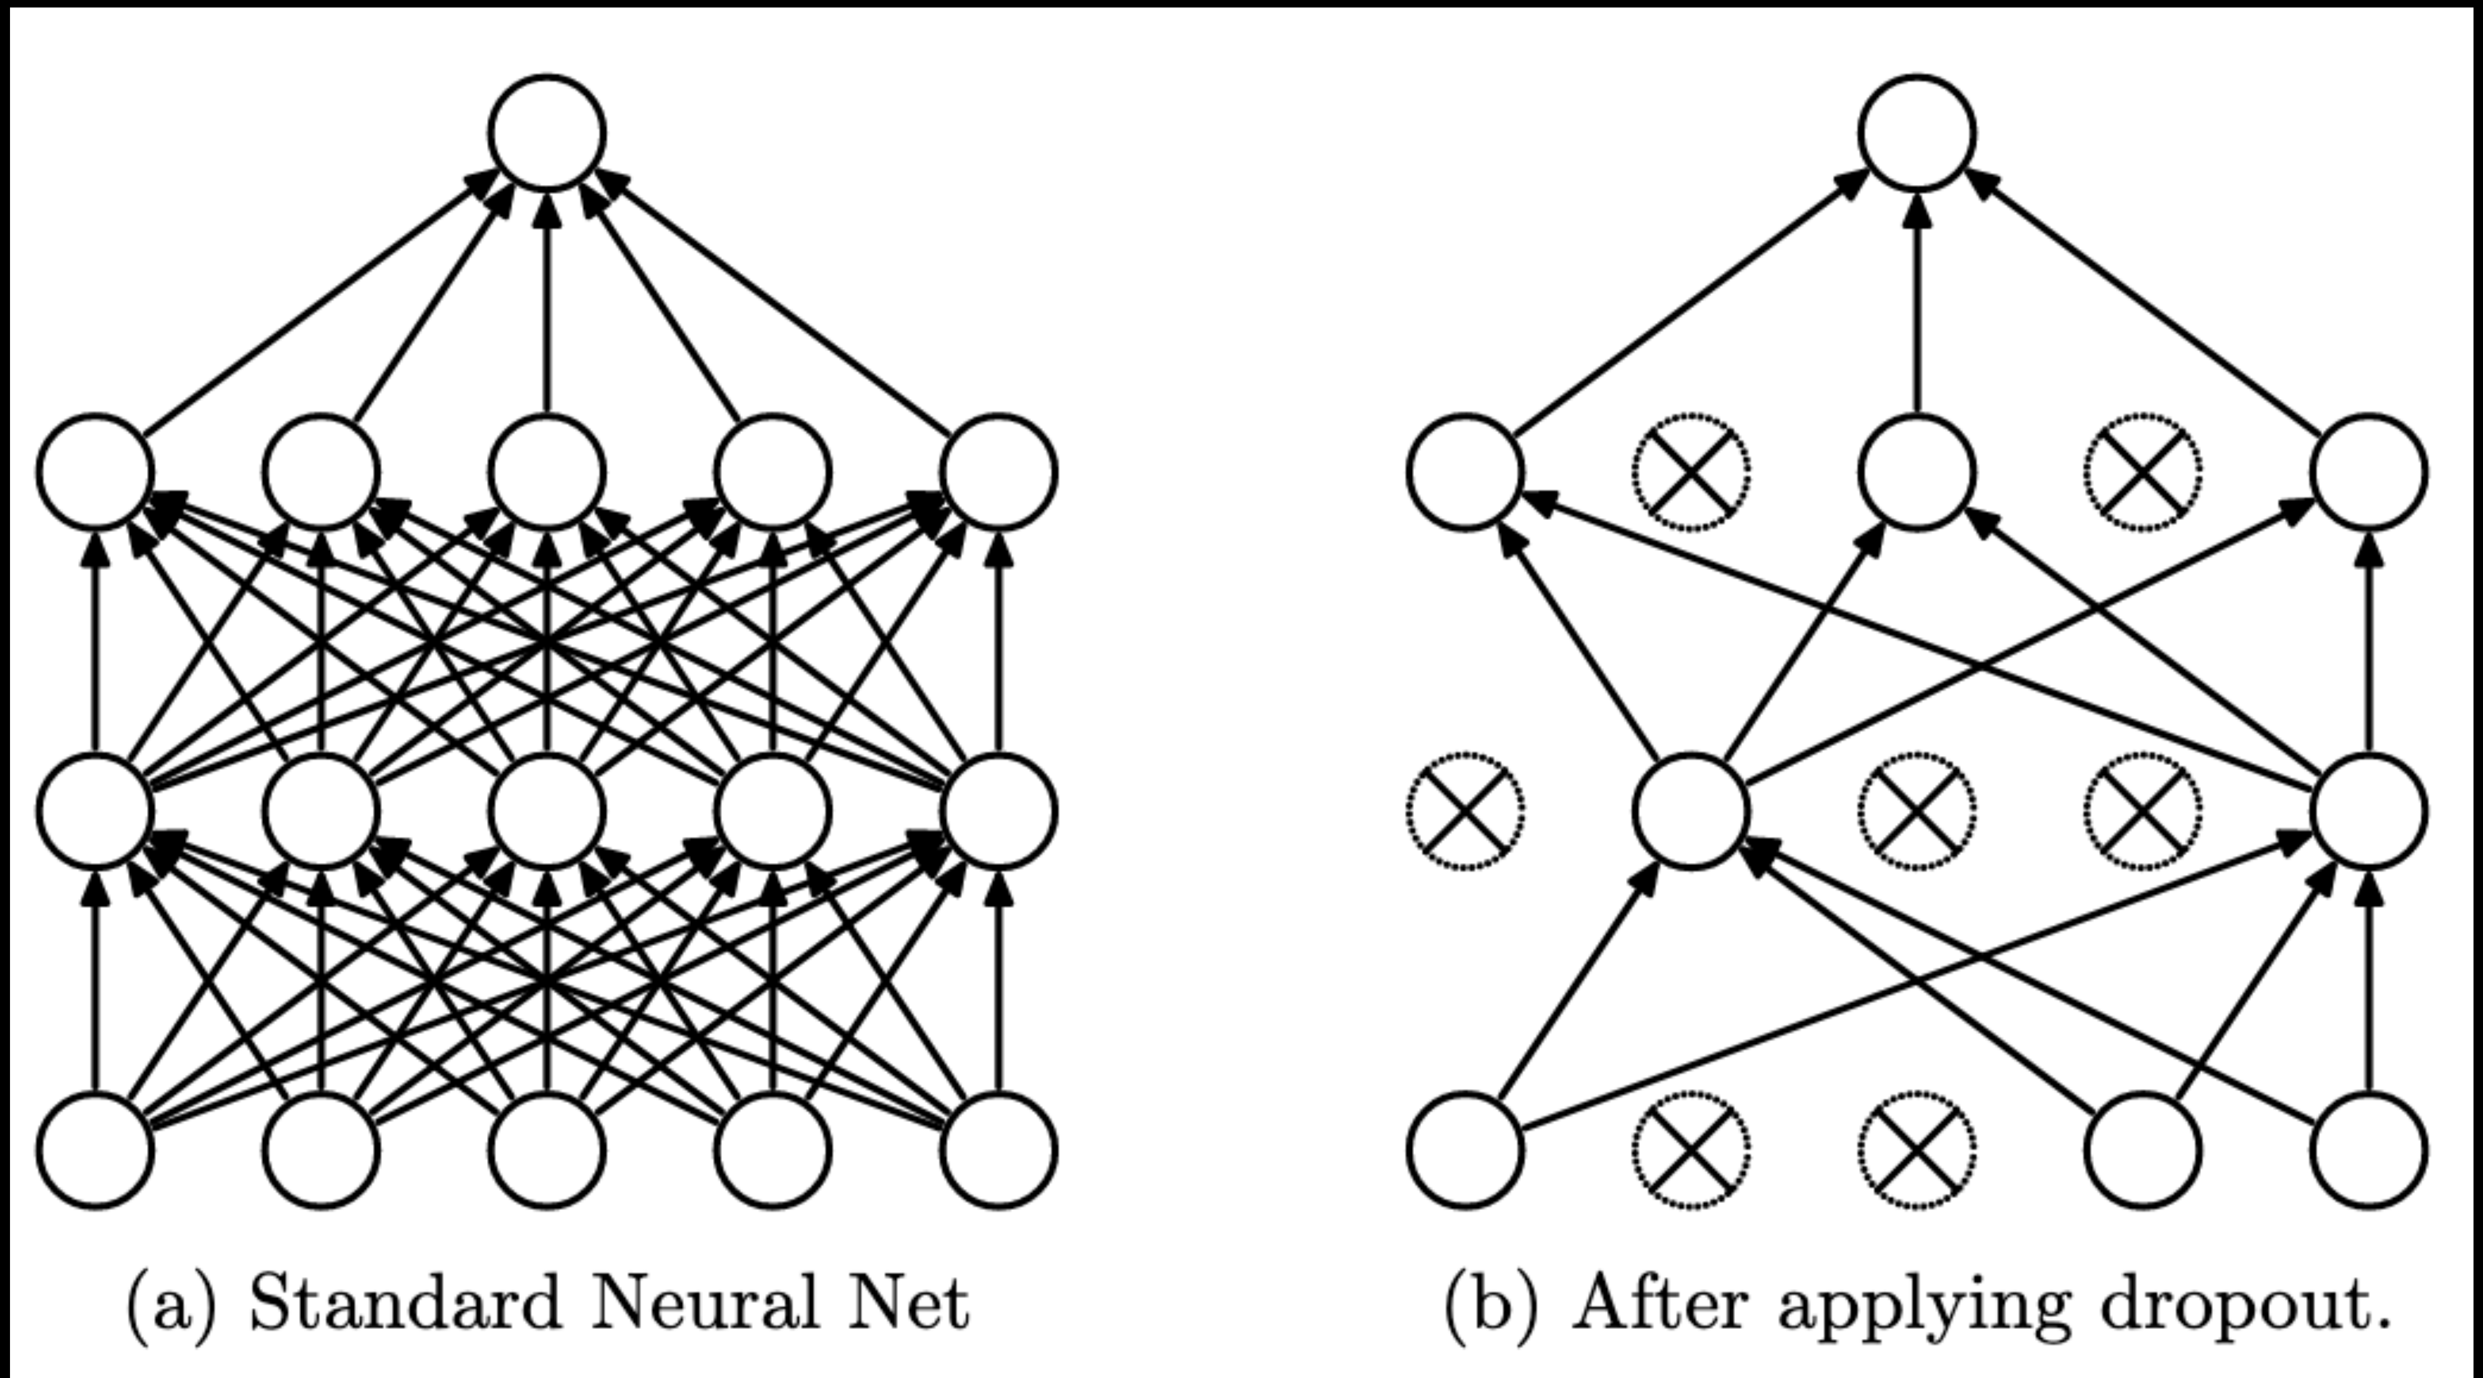

[1] Srivastava et al., JMLR, 2014

Batch normalization

- Internal covariance shift[1]
  - The distribution of the inputs in each layer changes as learning occurs in previous layers.
- Batch normalization[1] normalizes output of the previous layer by subtracting the batch mean, and then dividing by the batch's standard deviation (i.e., normalizing the previous output)

[1] Ioffe et al., PMLR, 2015

Data stratification

- A proper data stratification ensures that training and evaluation data is representative of the distributions in the population.

- Things to consider in MRI applications:
  - Subject demographics (sex, age,…)
  - Patients/Healthy volunteers
  - Different diseases
  - Sequence acquisition parameters
  - …

Validation

- Different validation methods
  - Train/test split
  - k-Fold cross validation
  - Leave-one-out cross validation
  - …

Validation

- k-fold cross validation

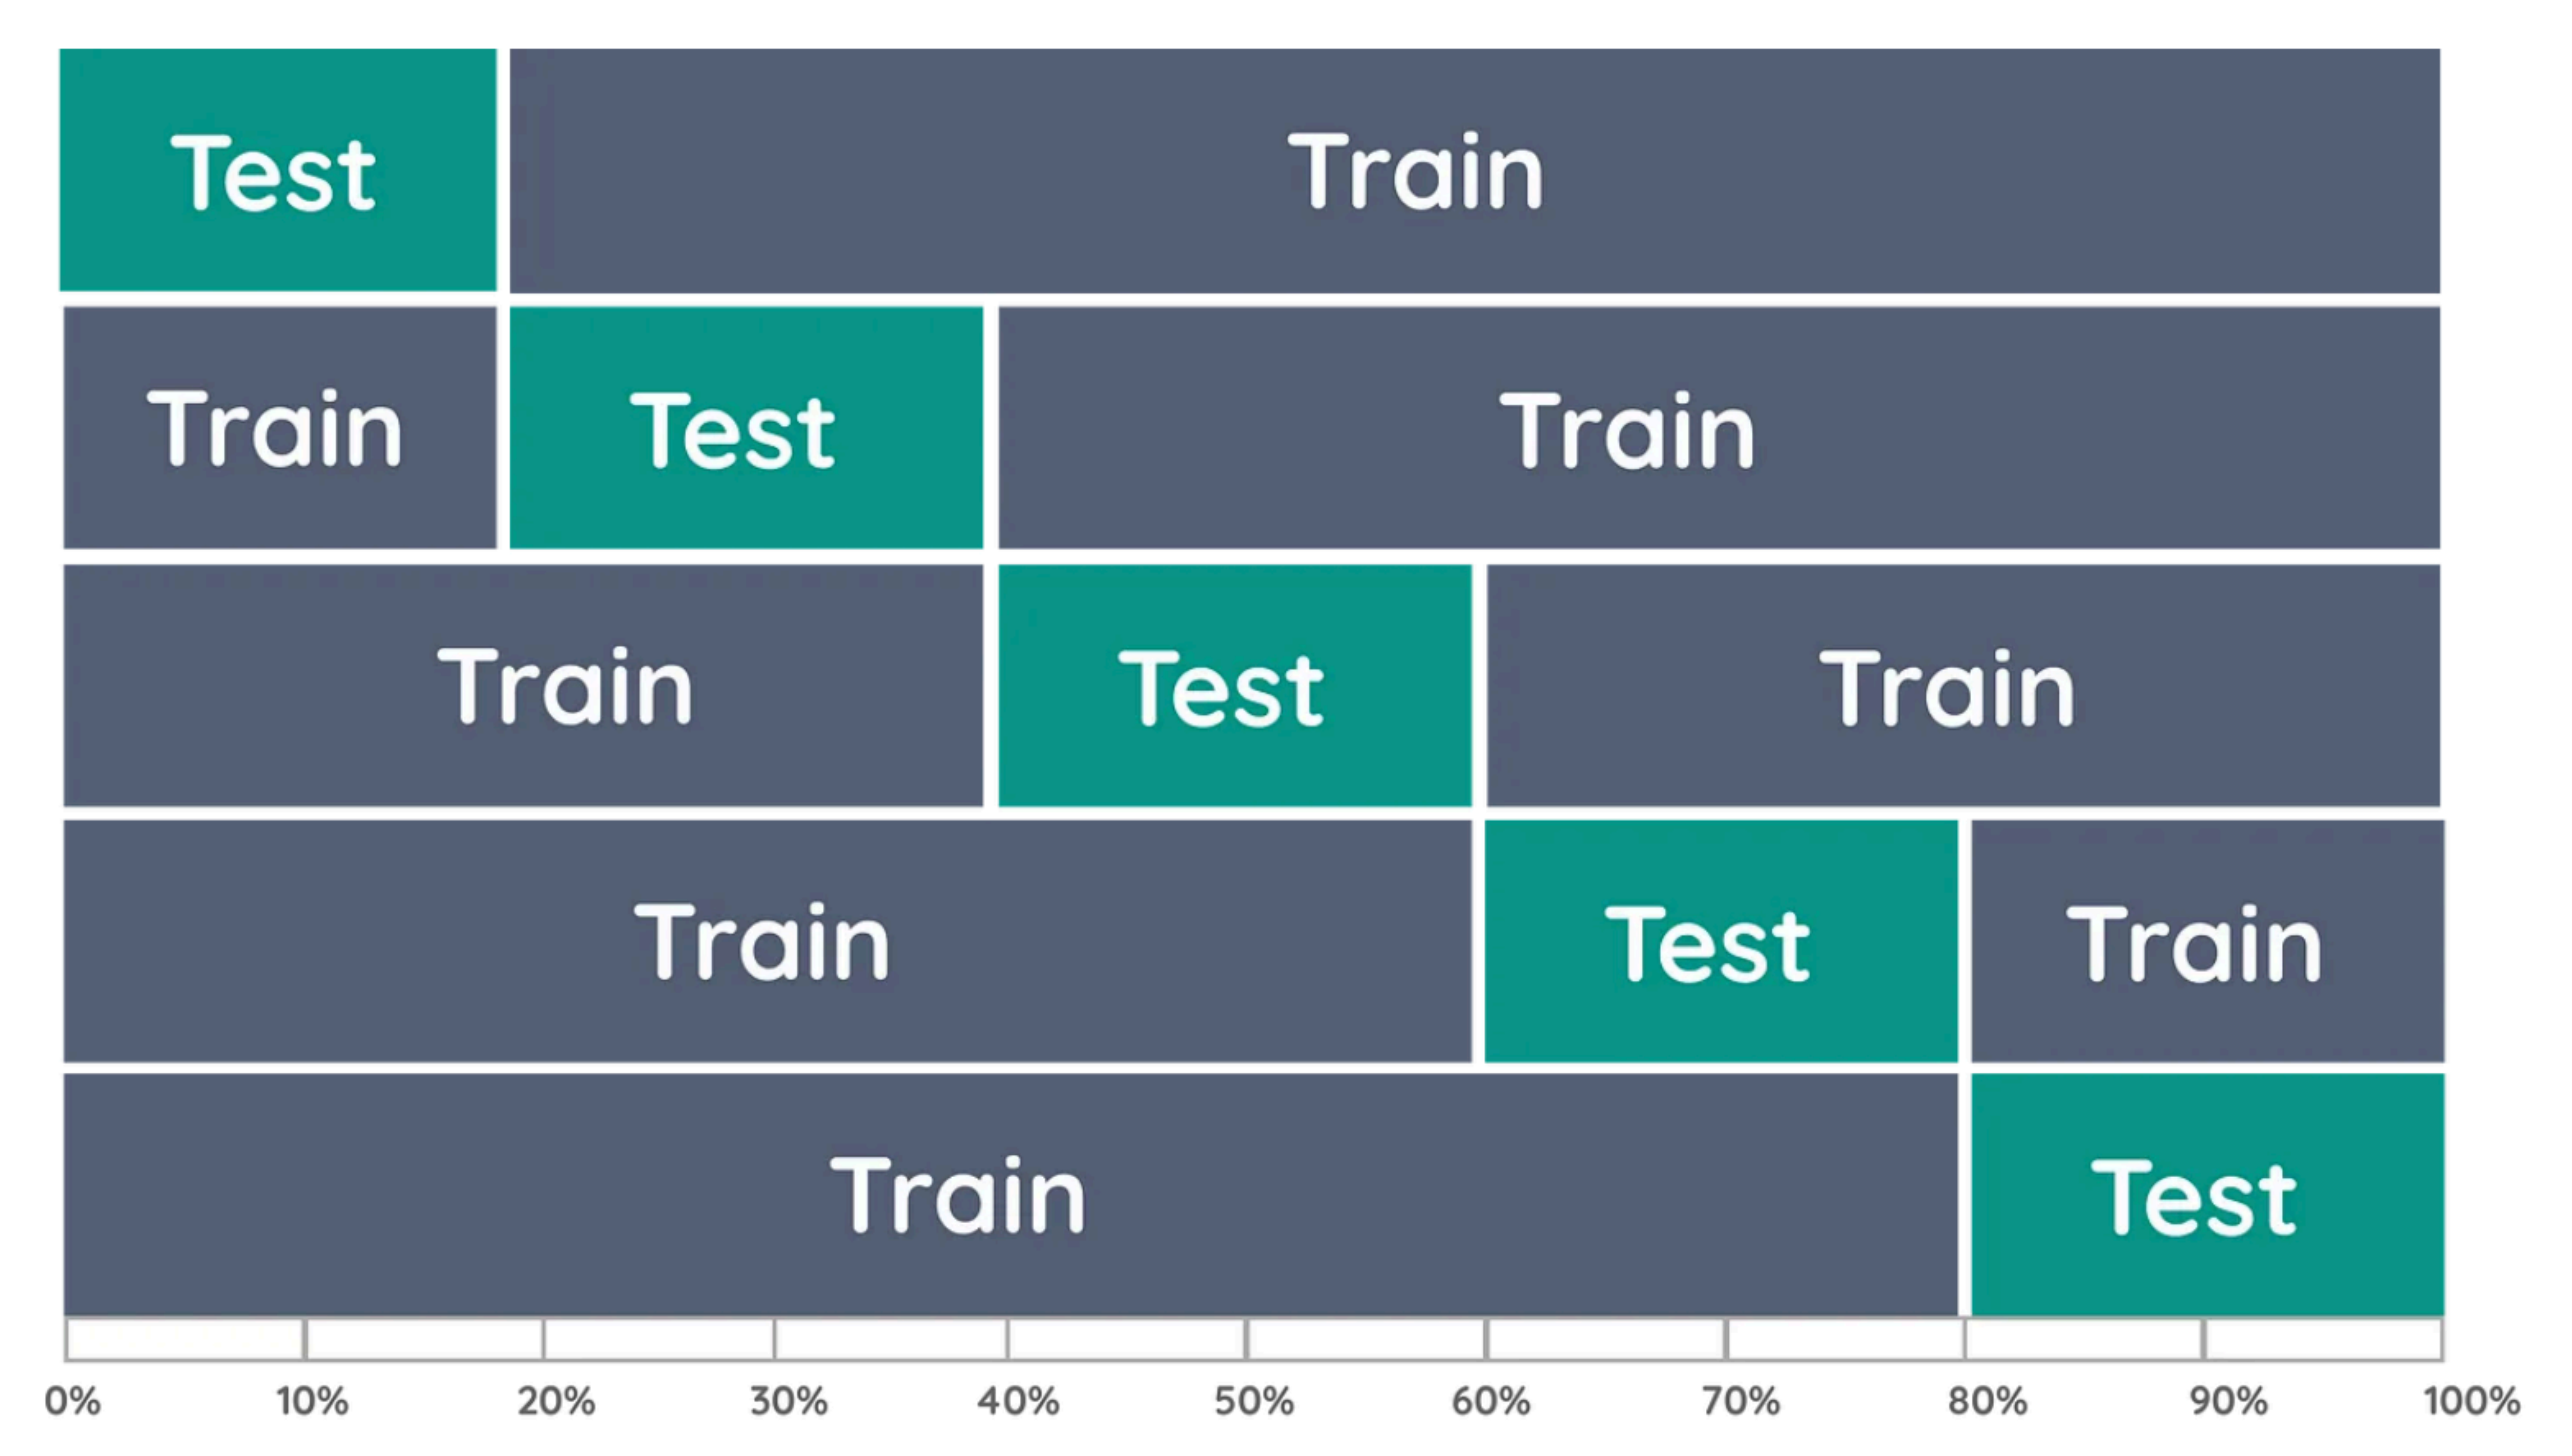

(Figure from: https://towardsdatascience.com/validating-your-machine-learning-model-25b4c8643fb7)

Hyperparameter tuning

- There are many hyperparameters in deep learning networks
  - Learning rate
  - Batch size
  - Architecture design: number of layers, numbers of channels
  - …

- Approaches for hyperparameter tuning
  - Grid search
  - Random search

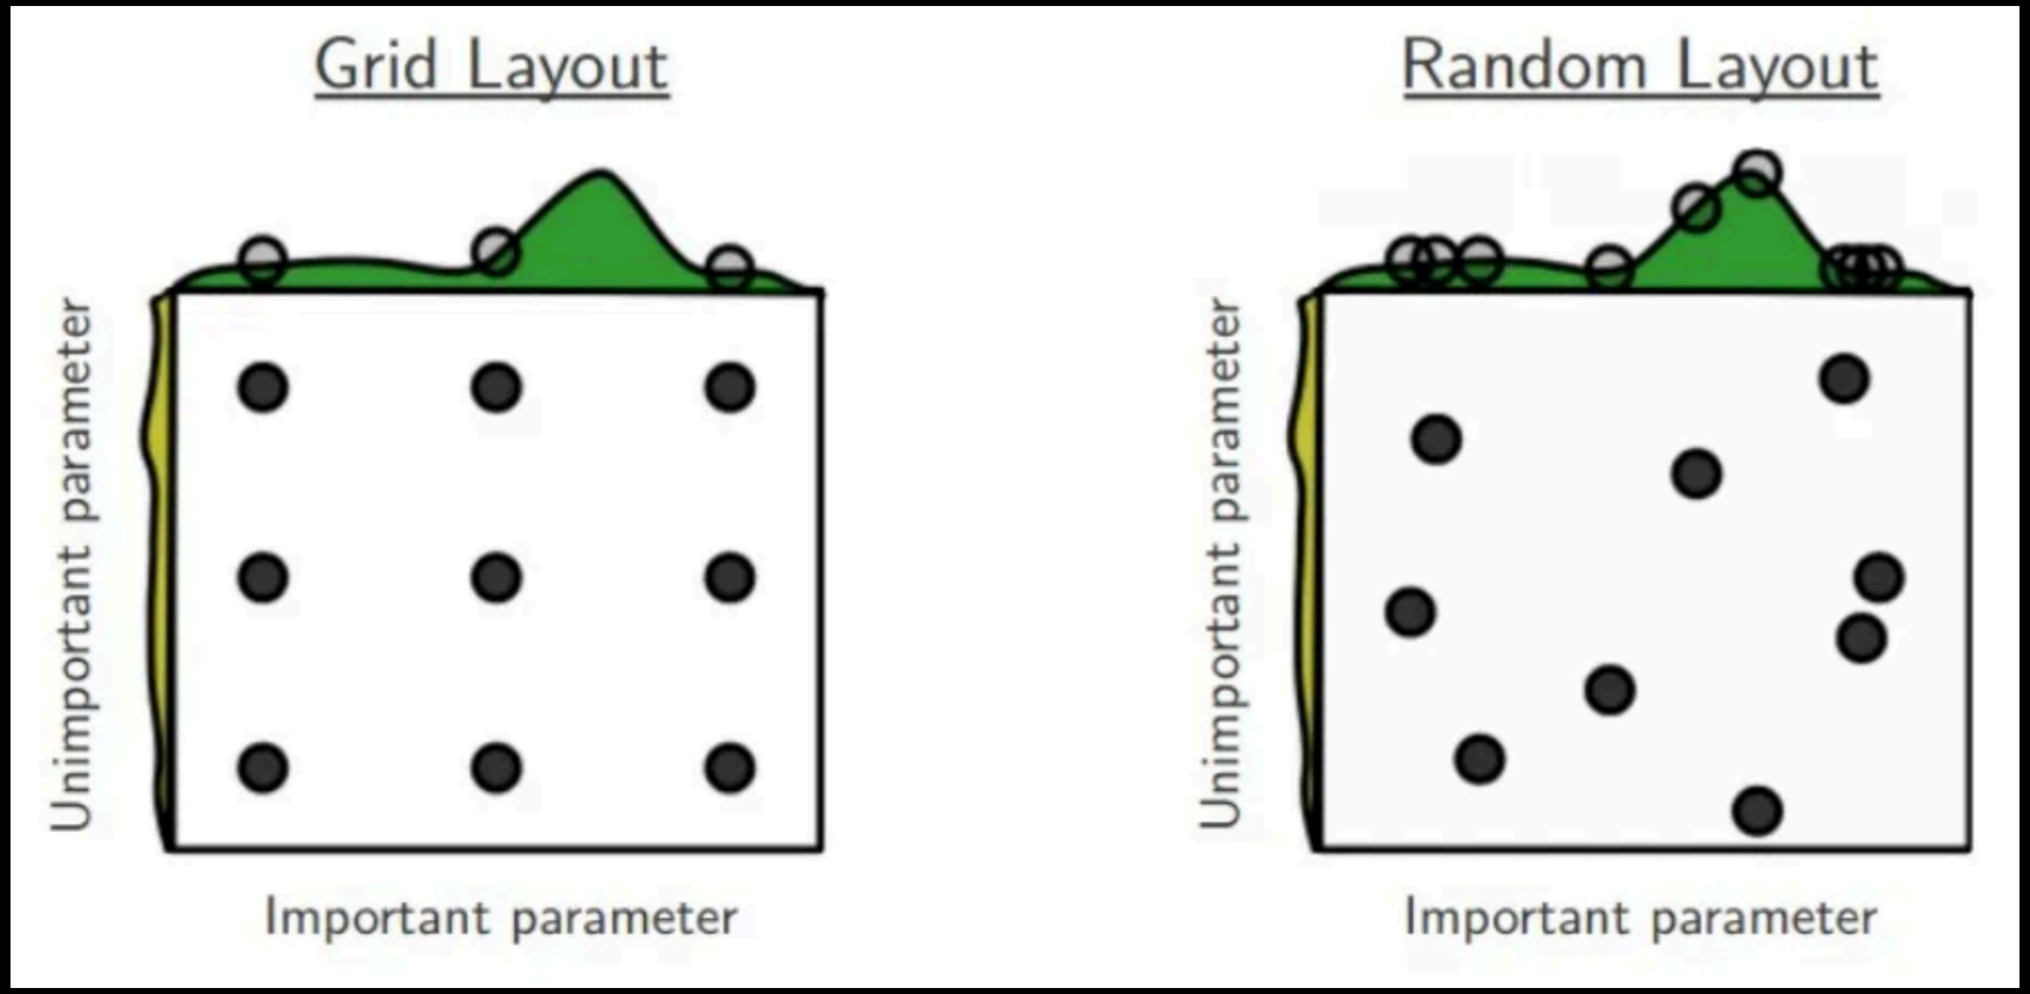

https://towardsdatascience.com/hyper-parameter-tuning-techniques-in-deep-learning-4dad592c63c8

Hyperparameter tuning

- Monitor validation loss for hyperparameter tuning
- Pay attention to signs of underfitting and overfitting

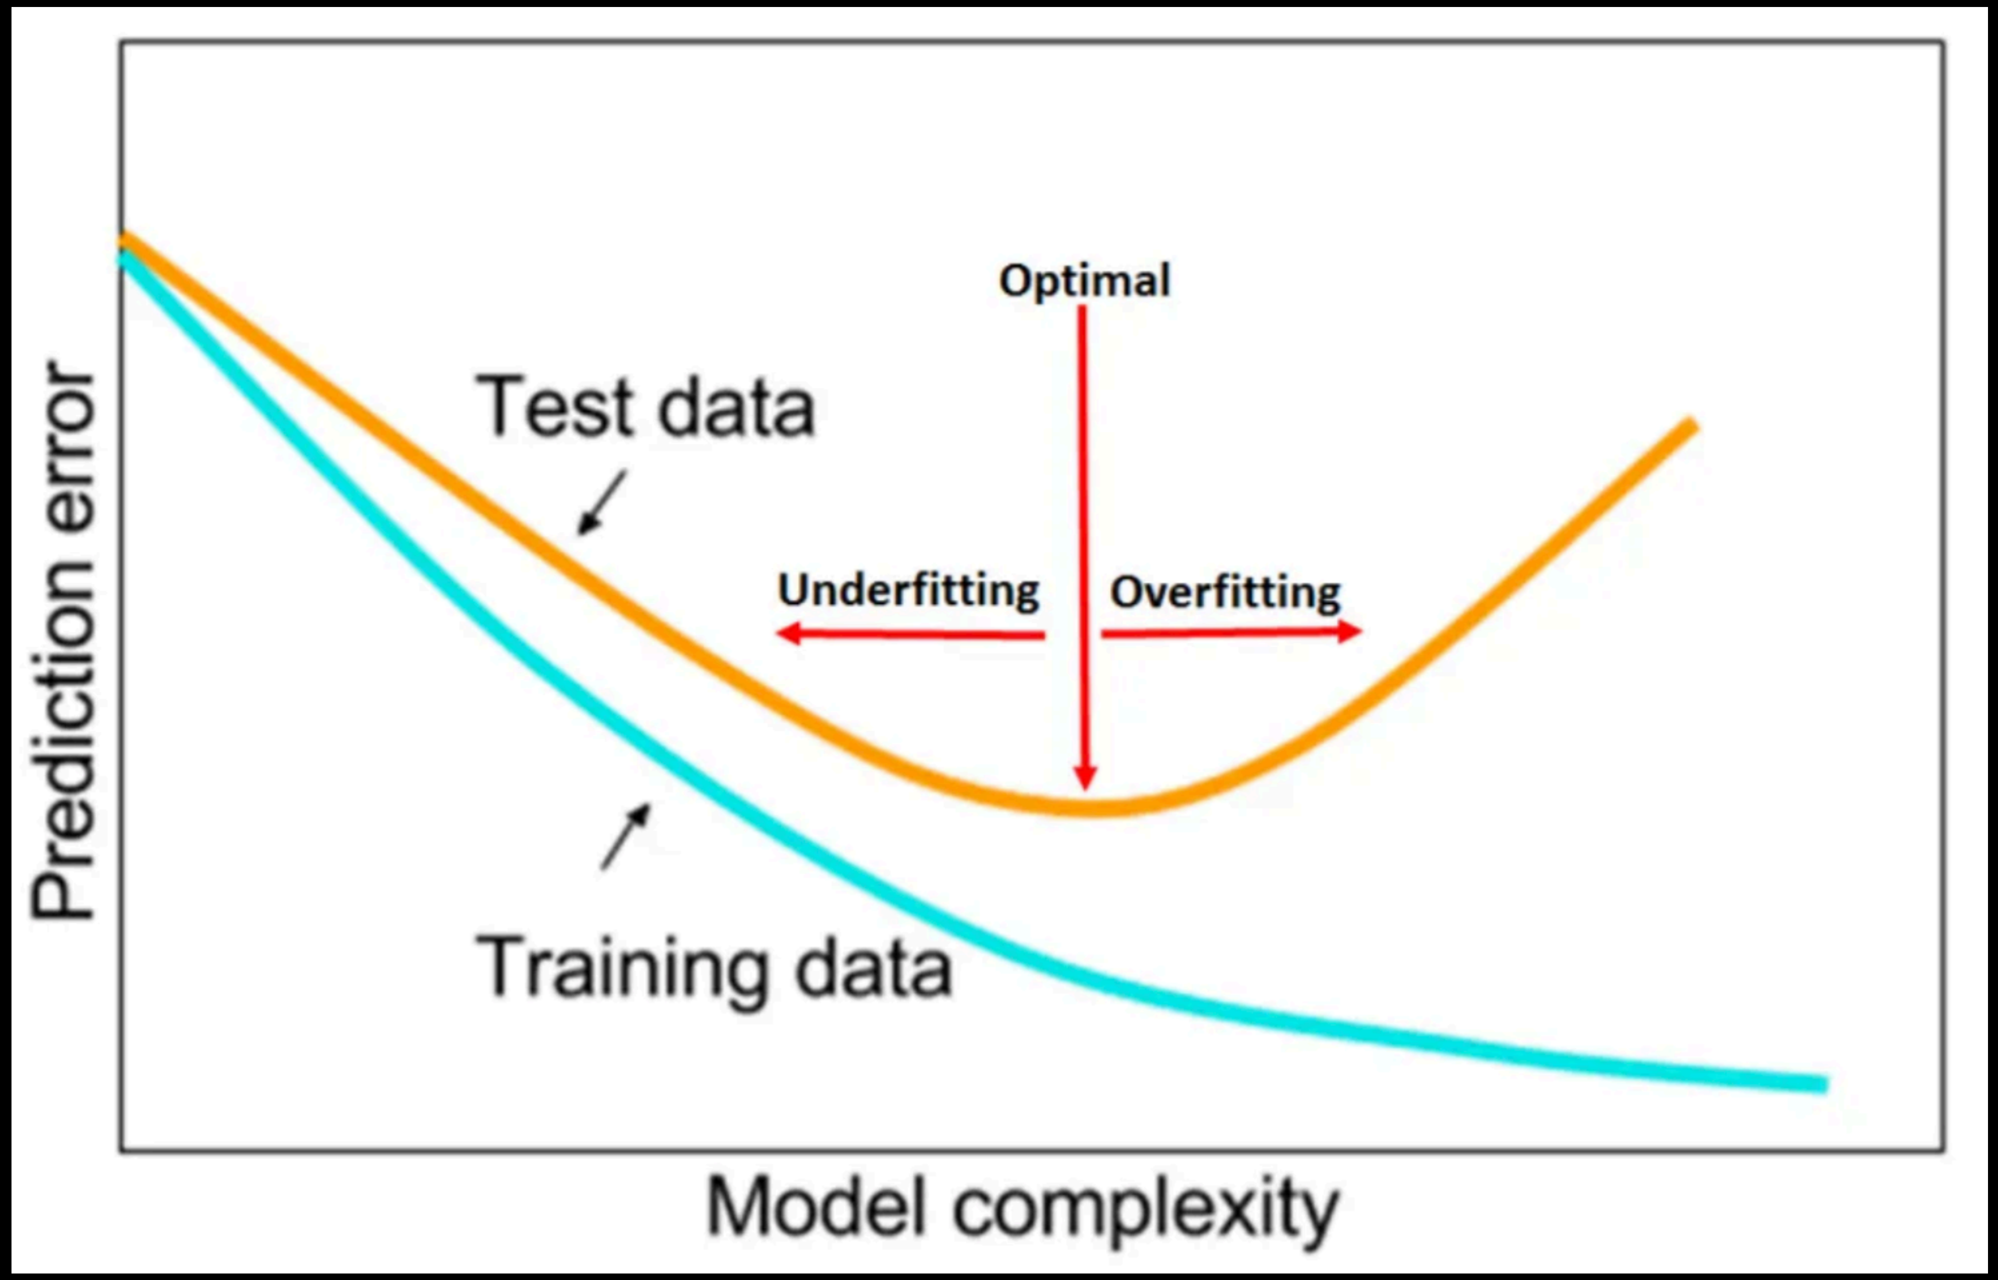

https://towardsdatascience.com/hyper-parameter-tuning-techniques-in-deep-learning-4dad592c63c8

Ablation study

- Ablation study investigates the performance of a neural network by removing one or several components at a time to understand the contribution from each component to the entire network.

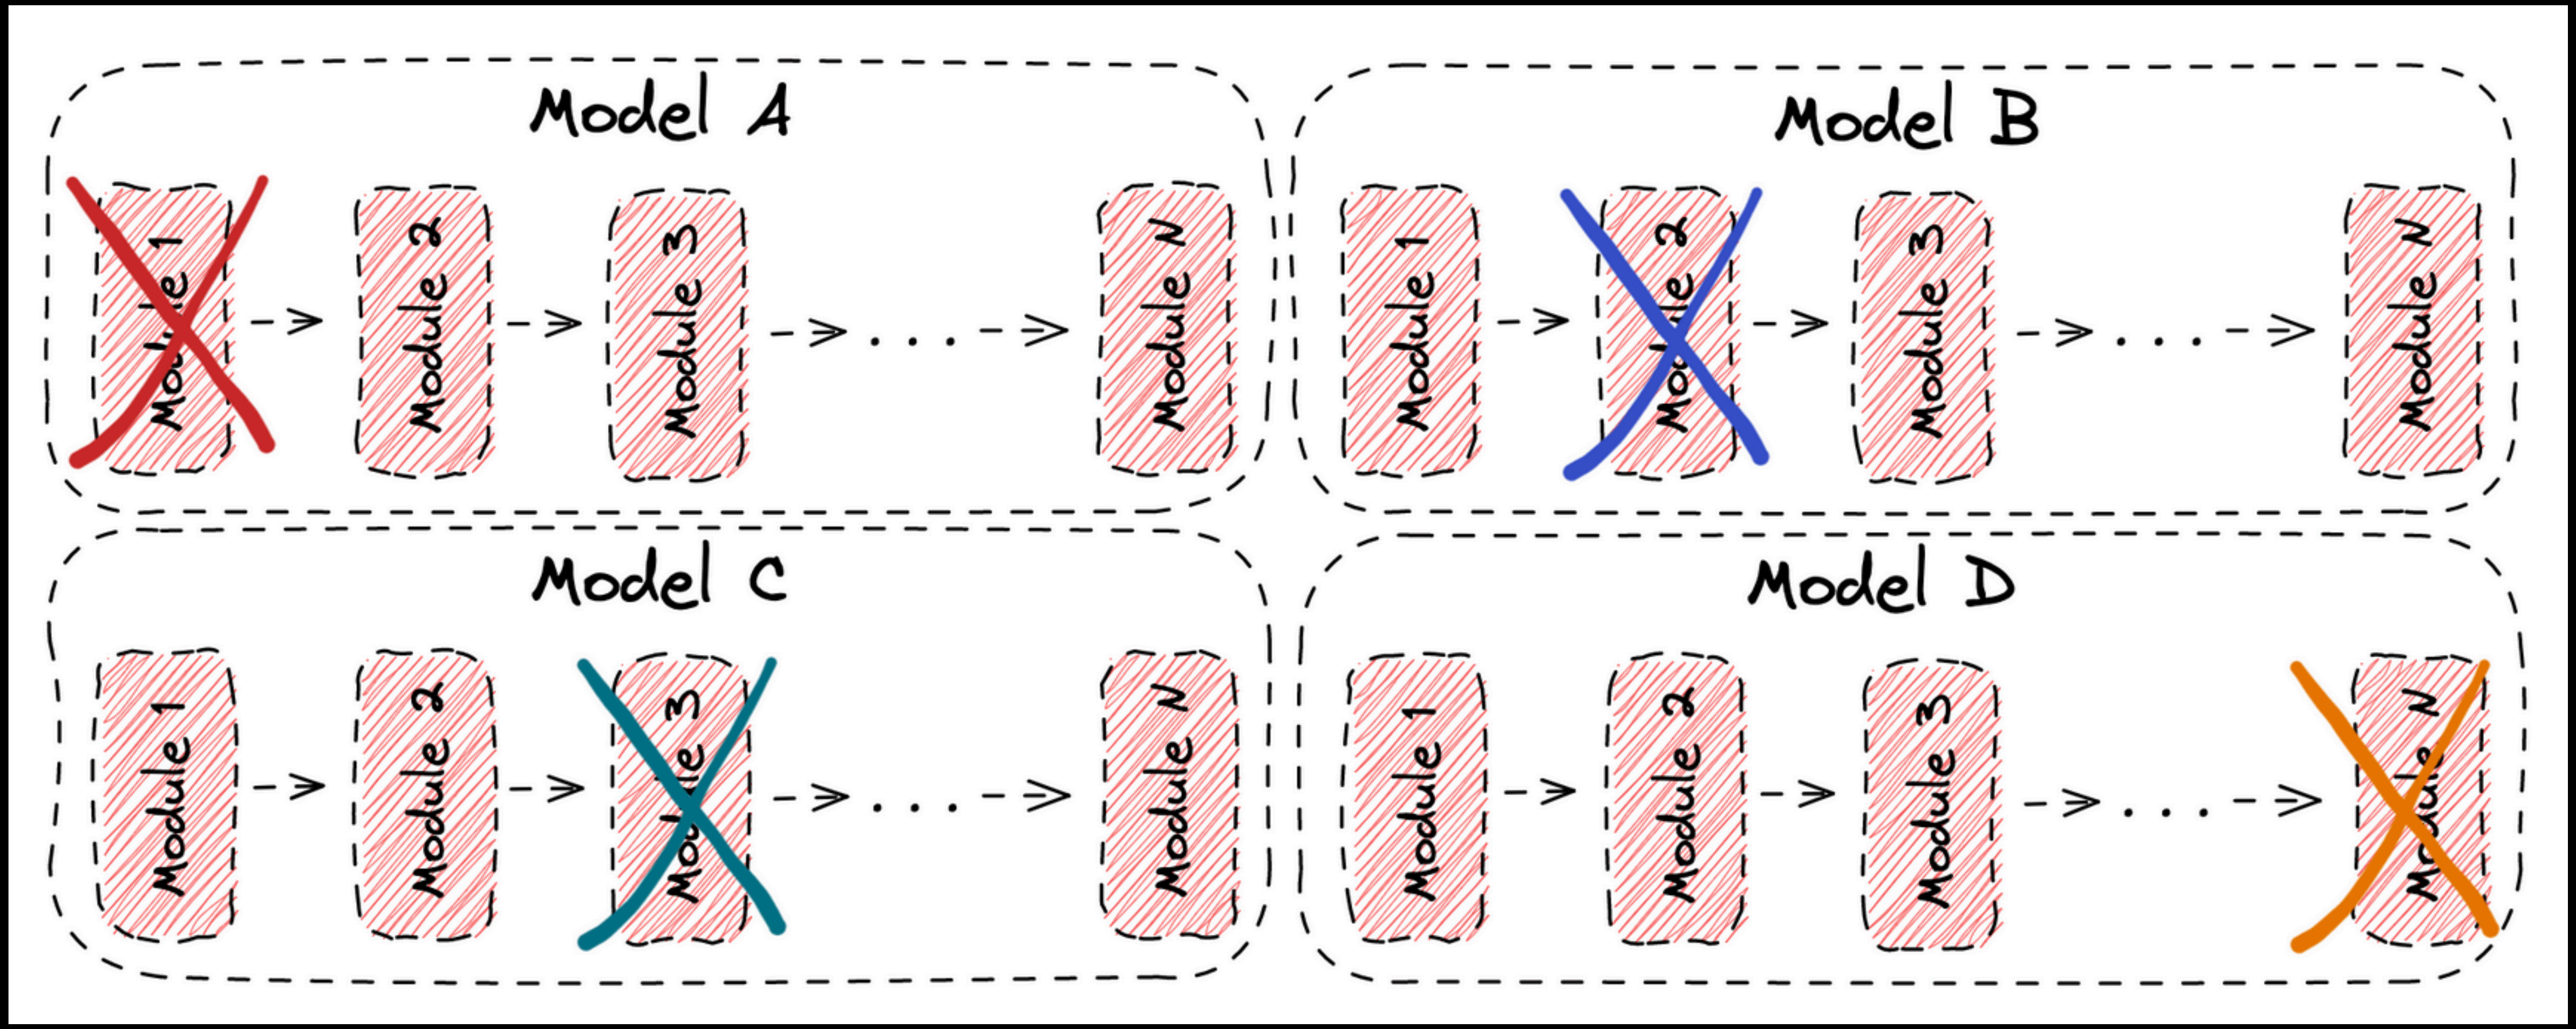

(Figure from: https://www.baeldung.com/cs/ml-ablation-study)

Image quality evaluation

- Quantitative image quality metrics
  - NRMSE, PSNR, SSIM…
- (For medical imaging applications) Radiology scoring
  - Experienced radiologists review and rate the image quality
- Statistical analysis

Part 3: Considerations for applying deep learning in MRI reconstruction

Considerations for MRI DL reconstruction

- Distinct differences between MRI recon versus other computer vision tasks:
  - (1) Data acquisition: MRI data acquired in the k-space domain, not in the image domain, and are inherently complex-valued.
  - (2) MRI physics: There is MRI physics behind the formation of the images.
  - (3) Availability of multi-contrast images: There can be multiple contrasts (e.g., different coils, different T1/T2 weightings) in the MRI dataset.
  - (4) Clinical workflow compatibility: Developing DL applications in MRI needs to consider whether it can be compatible with the clinical workflow.

- Let's see how MRI researchers apply deep learning with considerations of MRI data characteristics…

Deep learning MRI reconstruction

- Different approaches:

Image-domain learning

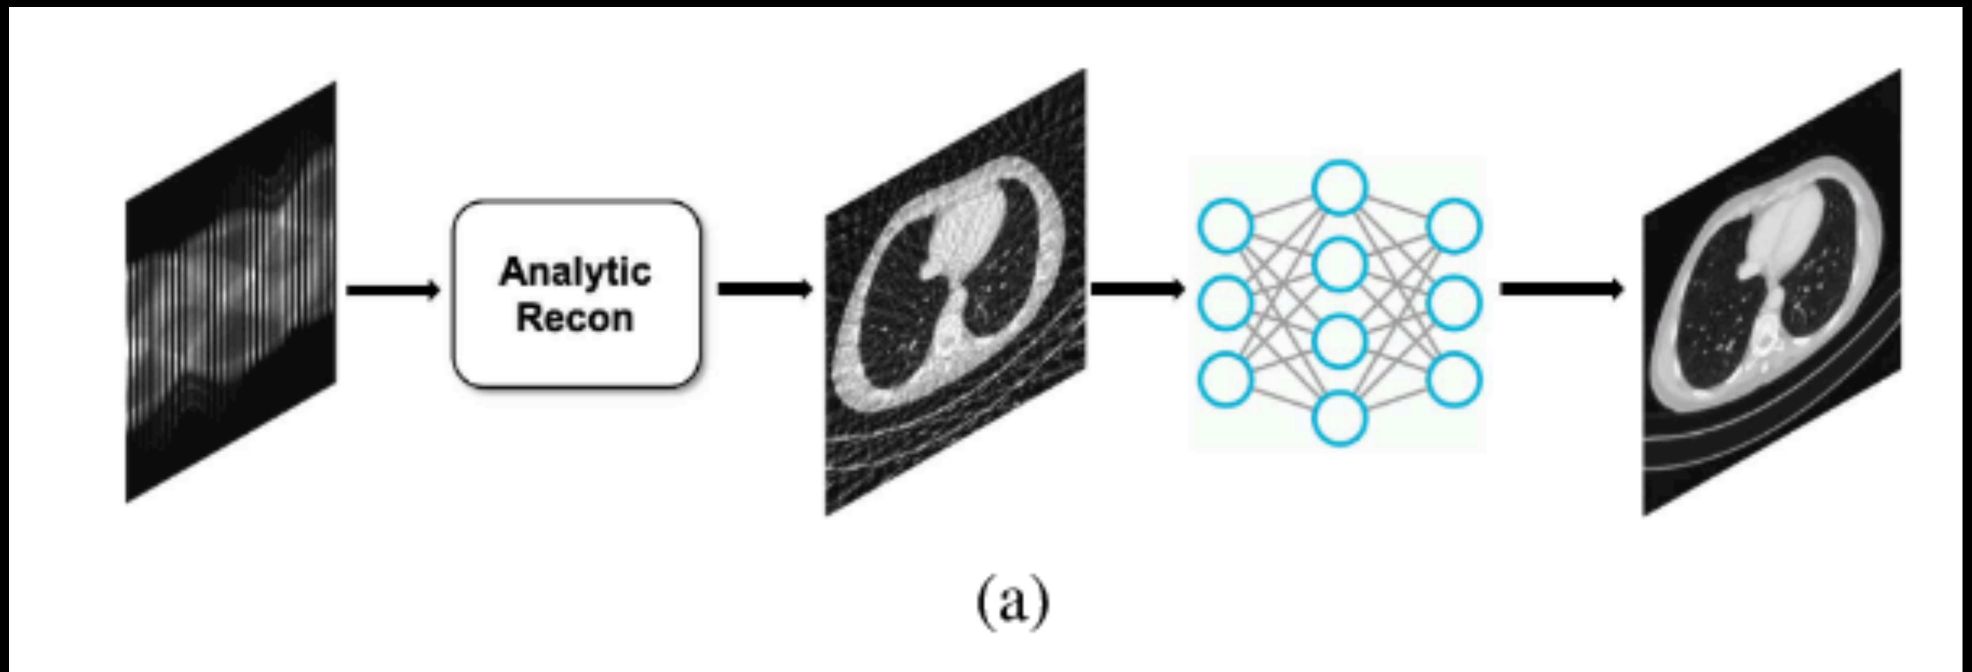

Hybrid-domain learning

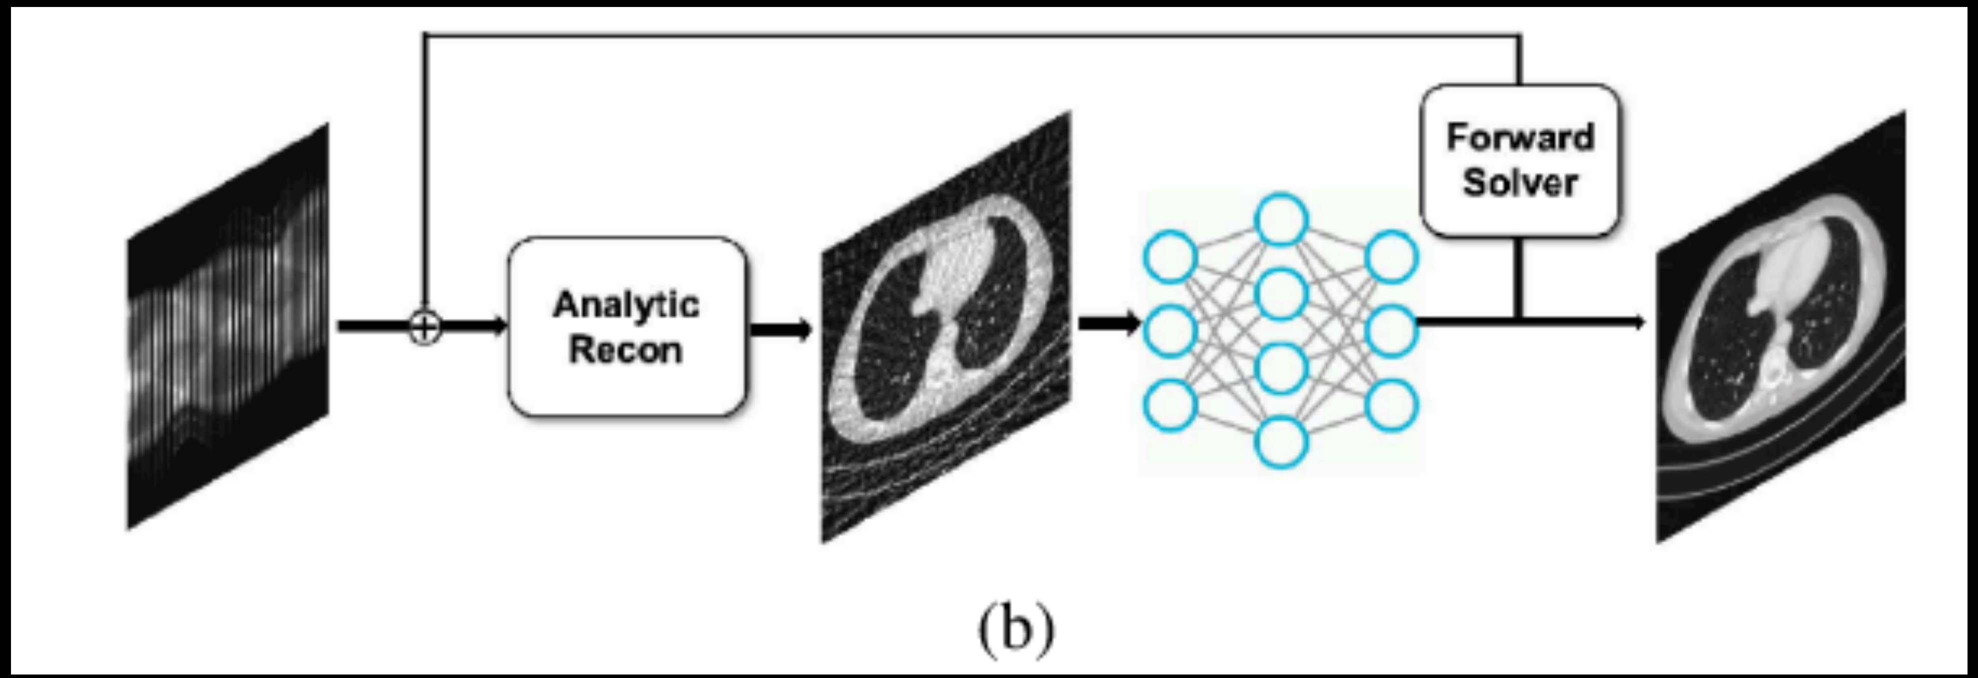

Mapping between sensor domain and image domain

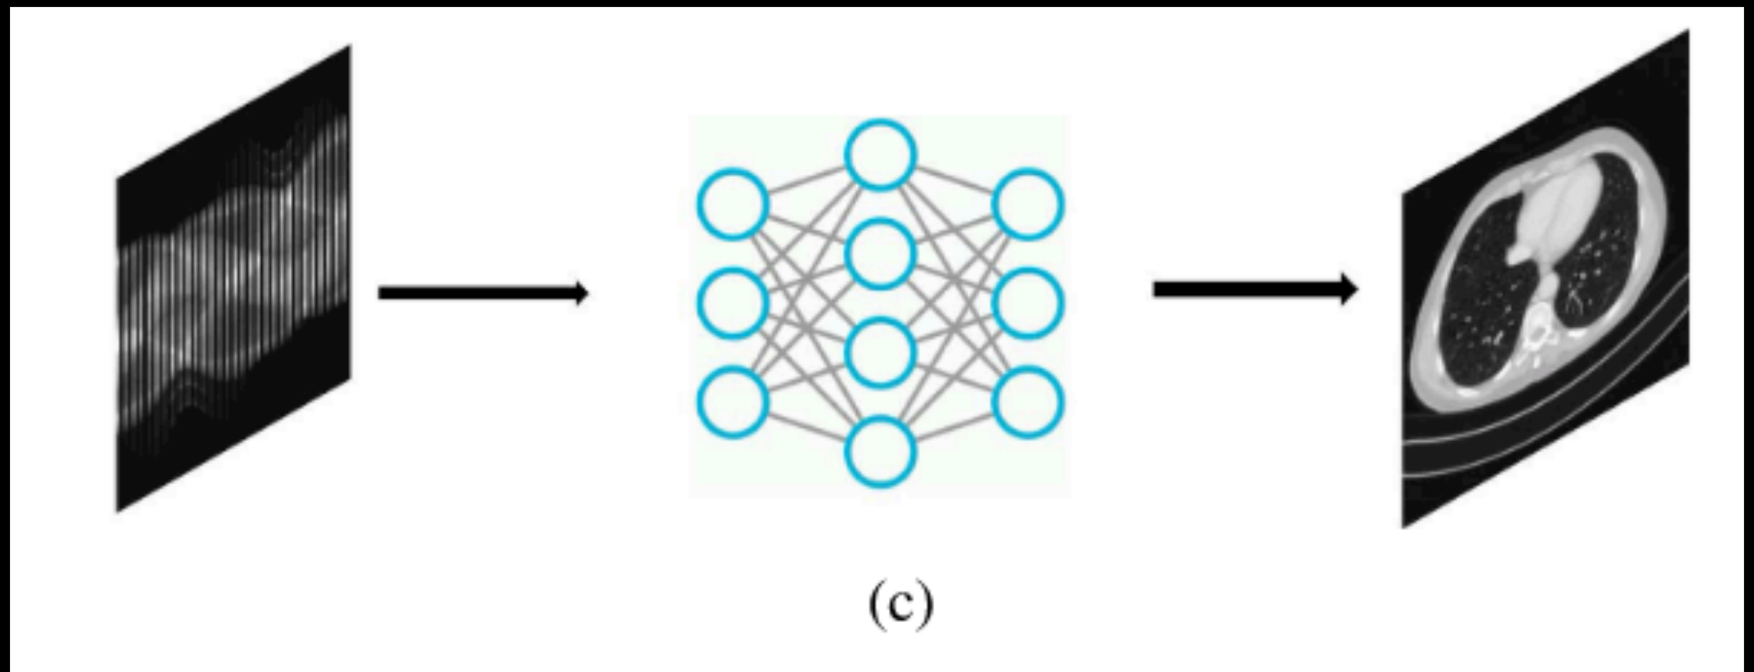

Sensor-domain learning

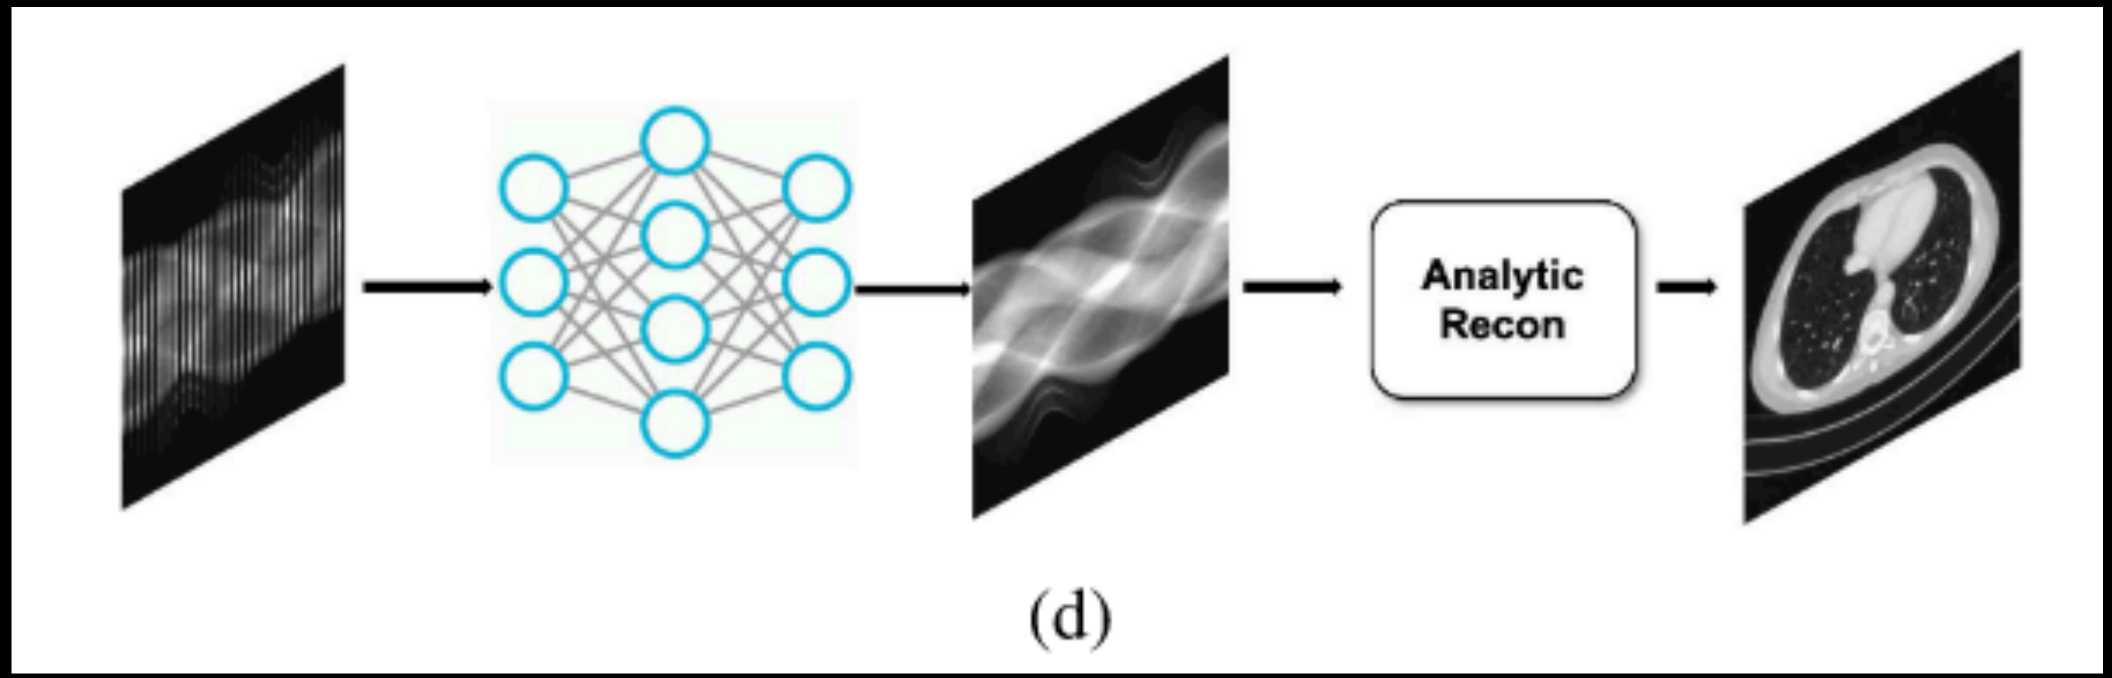

(Figures from: Ravishankar et al., Proceedings of the IEEE 2020)

Transforming compressed sensing models to DL

- MoDL (Model-based Deep Learning architecture for inverse problem)
  - Replace sparsity constraints (in CS formulation) with a deep learning network

Formulate as an optimization problem

$$x_{recon} = argmin_x \parallel UFx - y \parallel_2^2 + \lambda \parallel x - ConvNet(x) \parallel_2^2$$

An unrolled network with two main blocks
(1) A ConvNet to reduce artifacts / improve image quality
(2) A data consistency layer for k-space data consistency

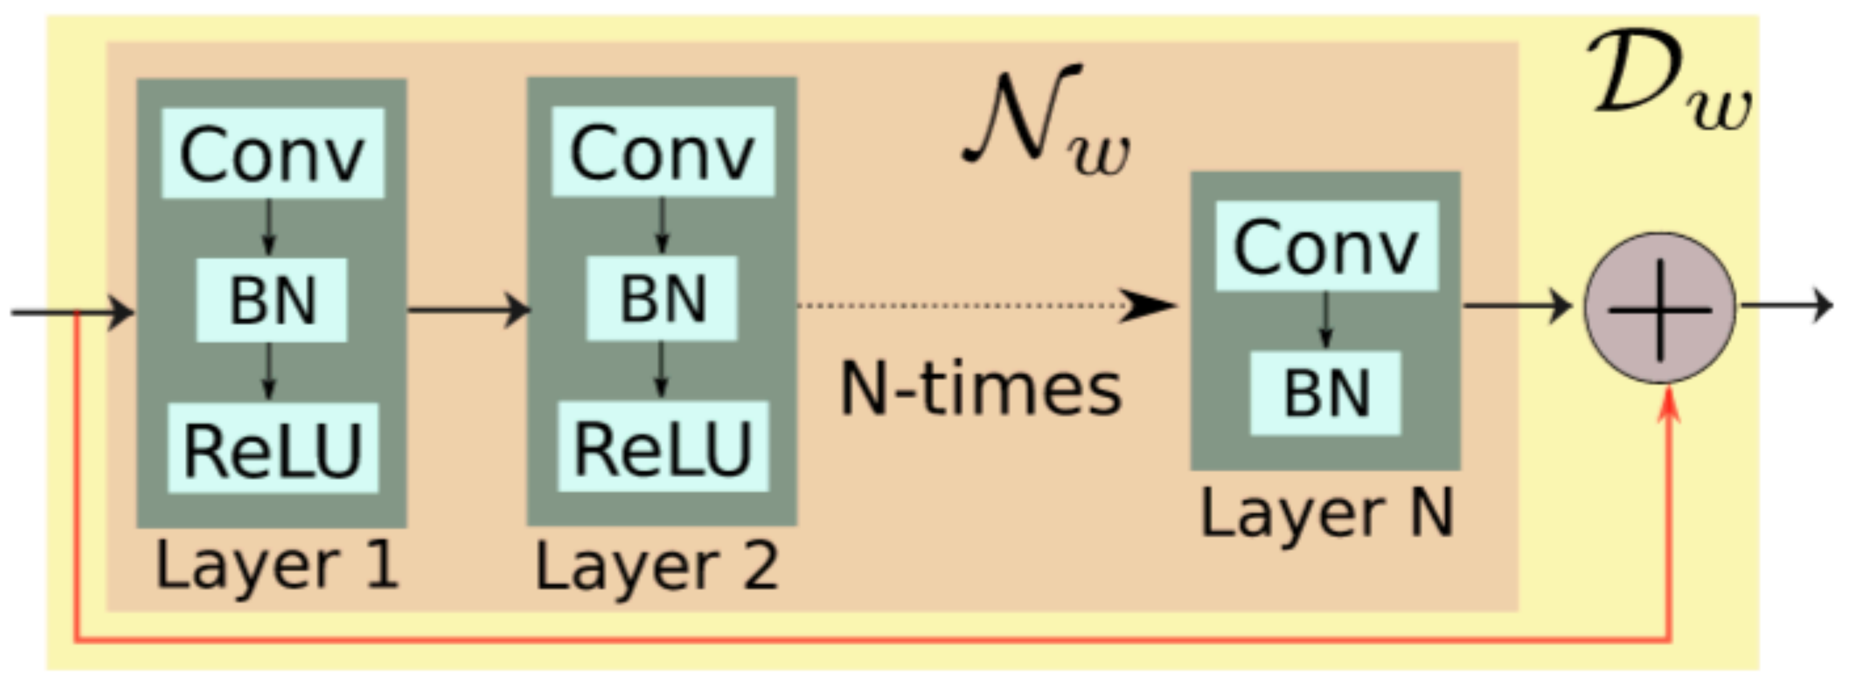

(a) The Residual learning based denoiser

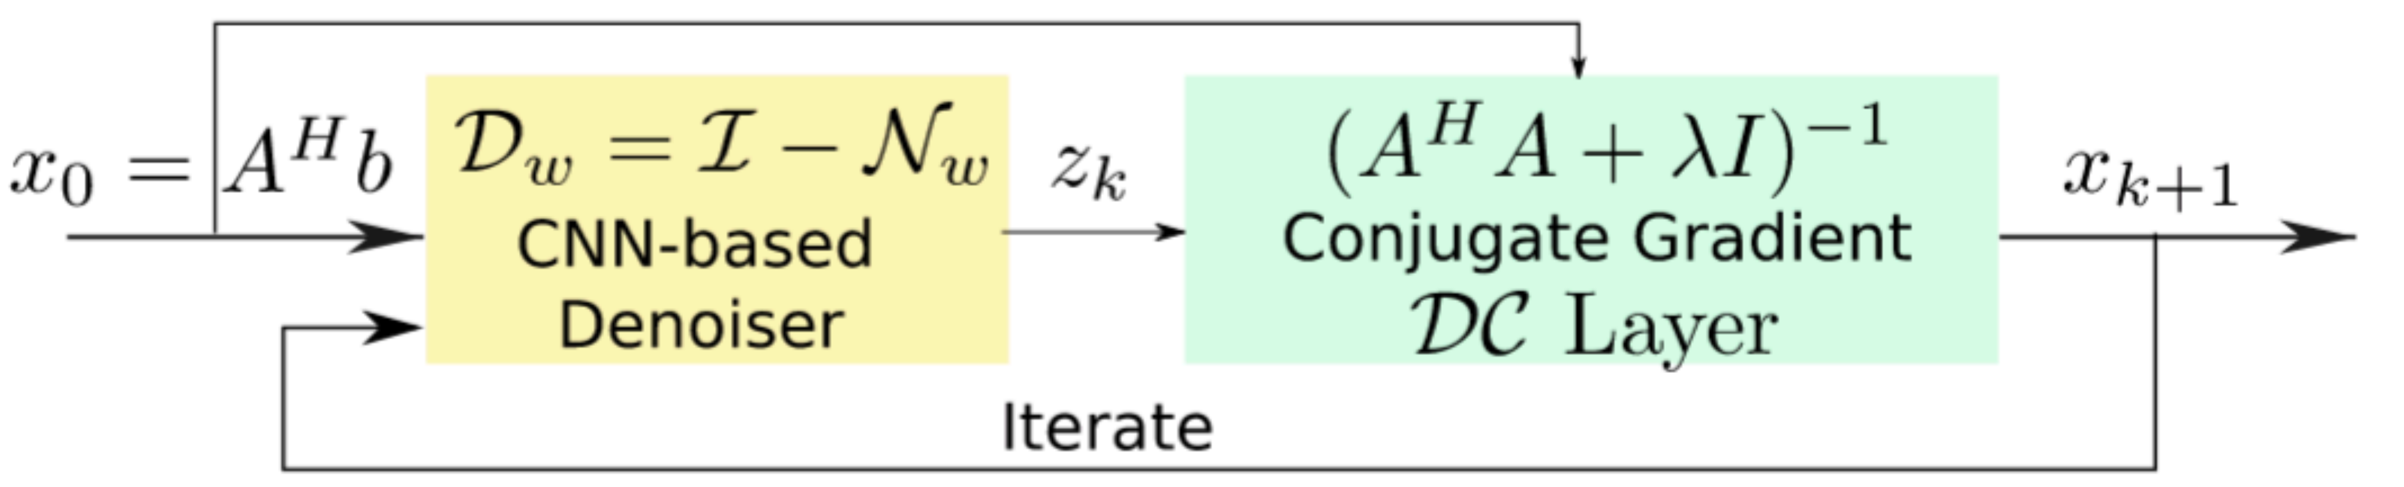

(b) Proposed Model-based Deep Learning (MoDL) architecture

(Figures from: Aggarwal et al., IEEE TMI 2019)

MoDL

Overall MoDL architecture

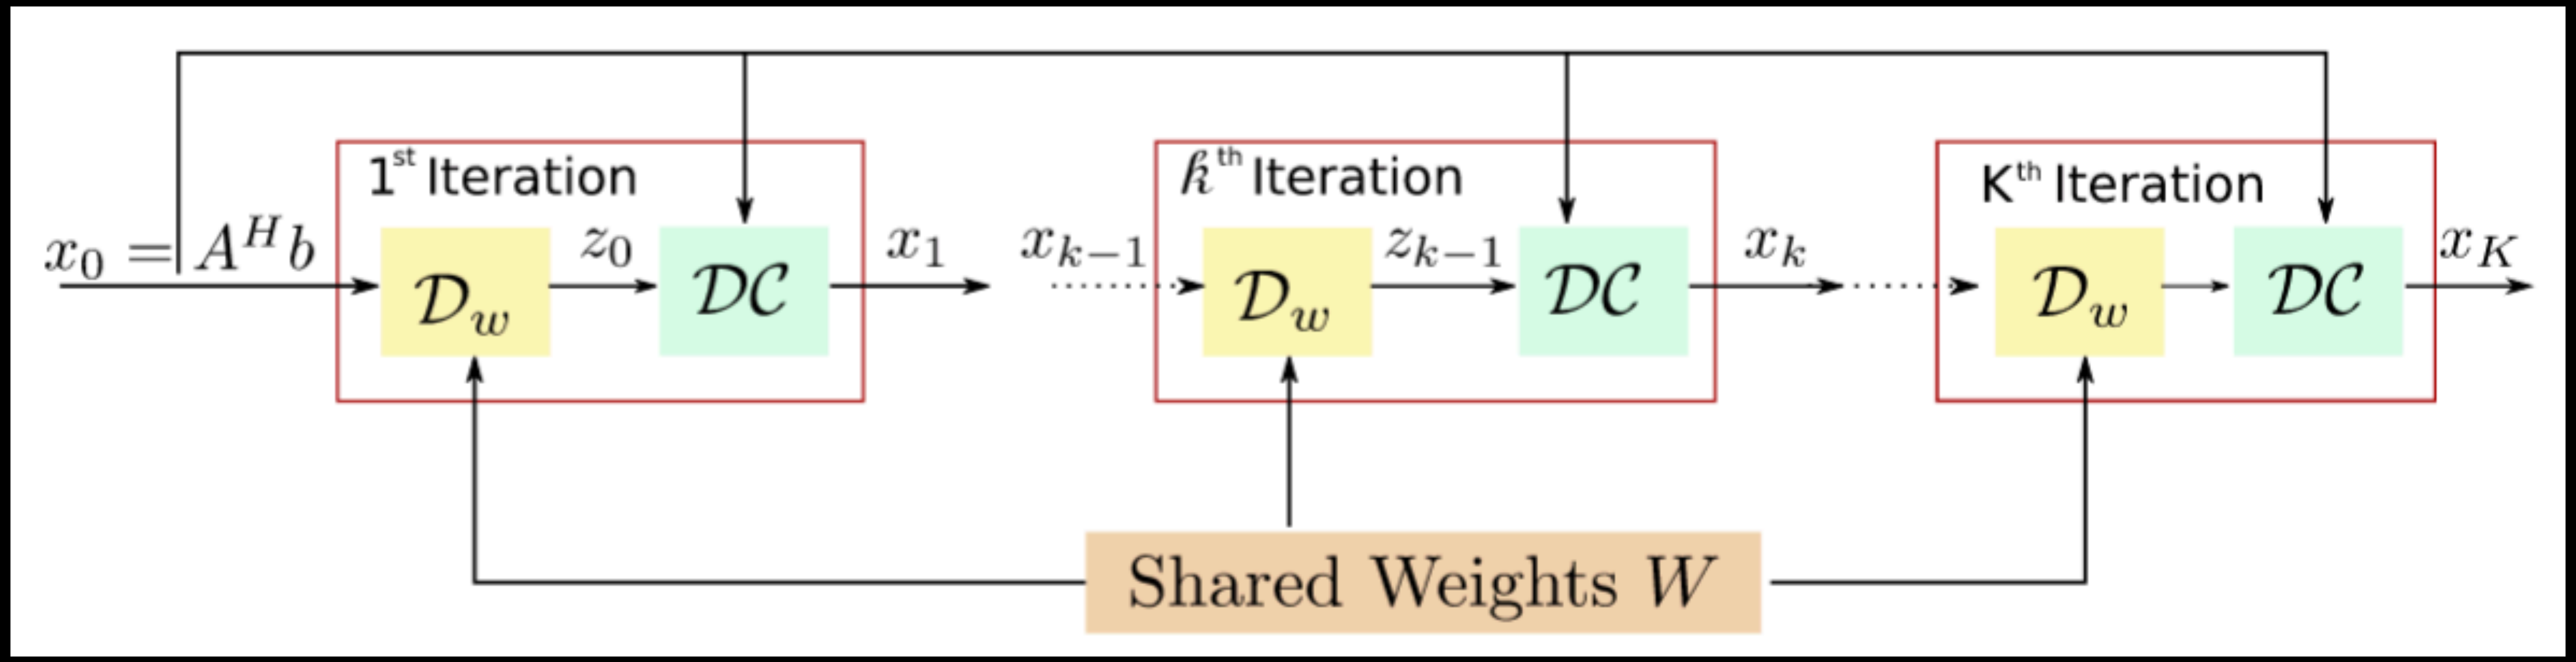

k-space sampling pattern

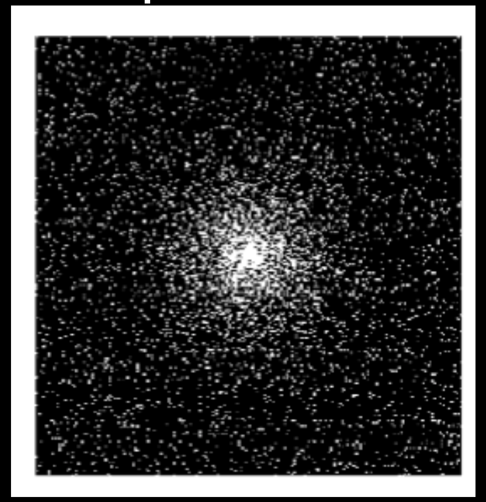

Zero-padding

Compressed sensing

MoDL

Image results

Error

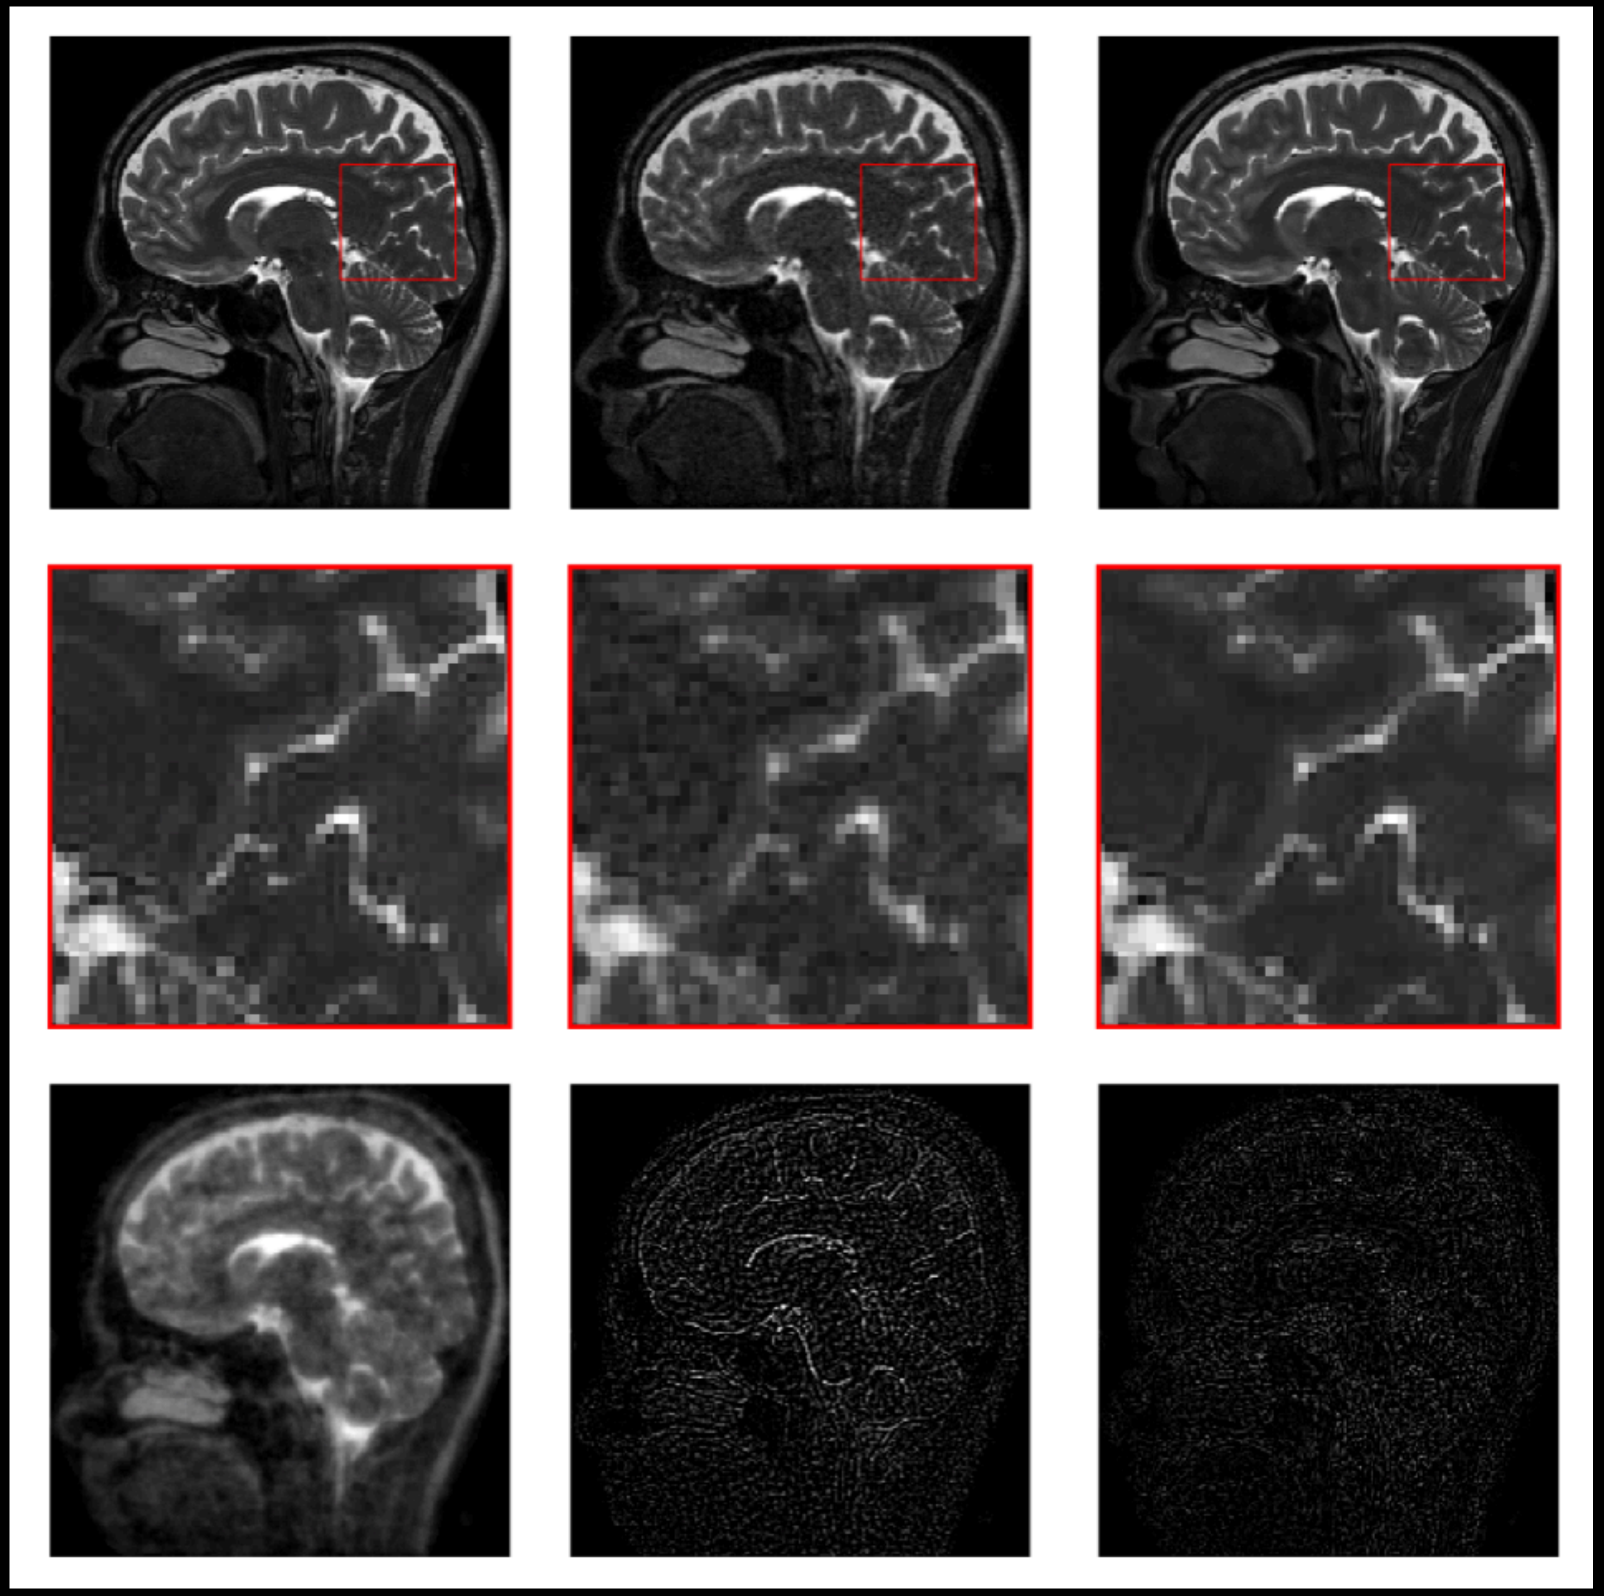

(Figures from: Aggarwal et al., IEEE TMI 2019)

Unrolled networks

- Unrolled networks are one of the most popular frameworks for DL MRI reconstruction as it integrates the strengths of traditional iterative optimization methods with the learning power of deep neural networks.
- This approach may offer better interpretability and theoretical guarantees of convergence.

Training in dual domains

- KIKI-net[1]: Use cross-domain ConvNets for image reconstruction
  - One sub-network for k-space completion
  - One sub-network for image restoration

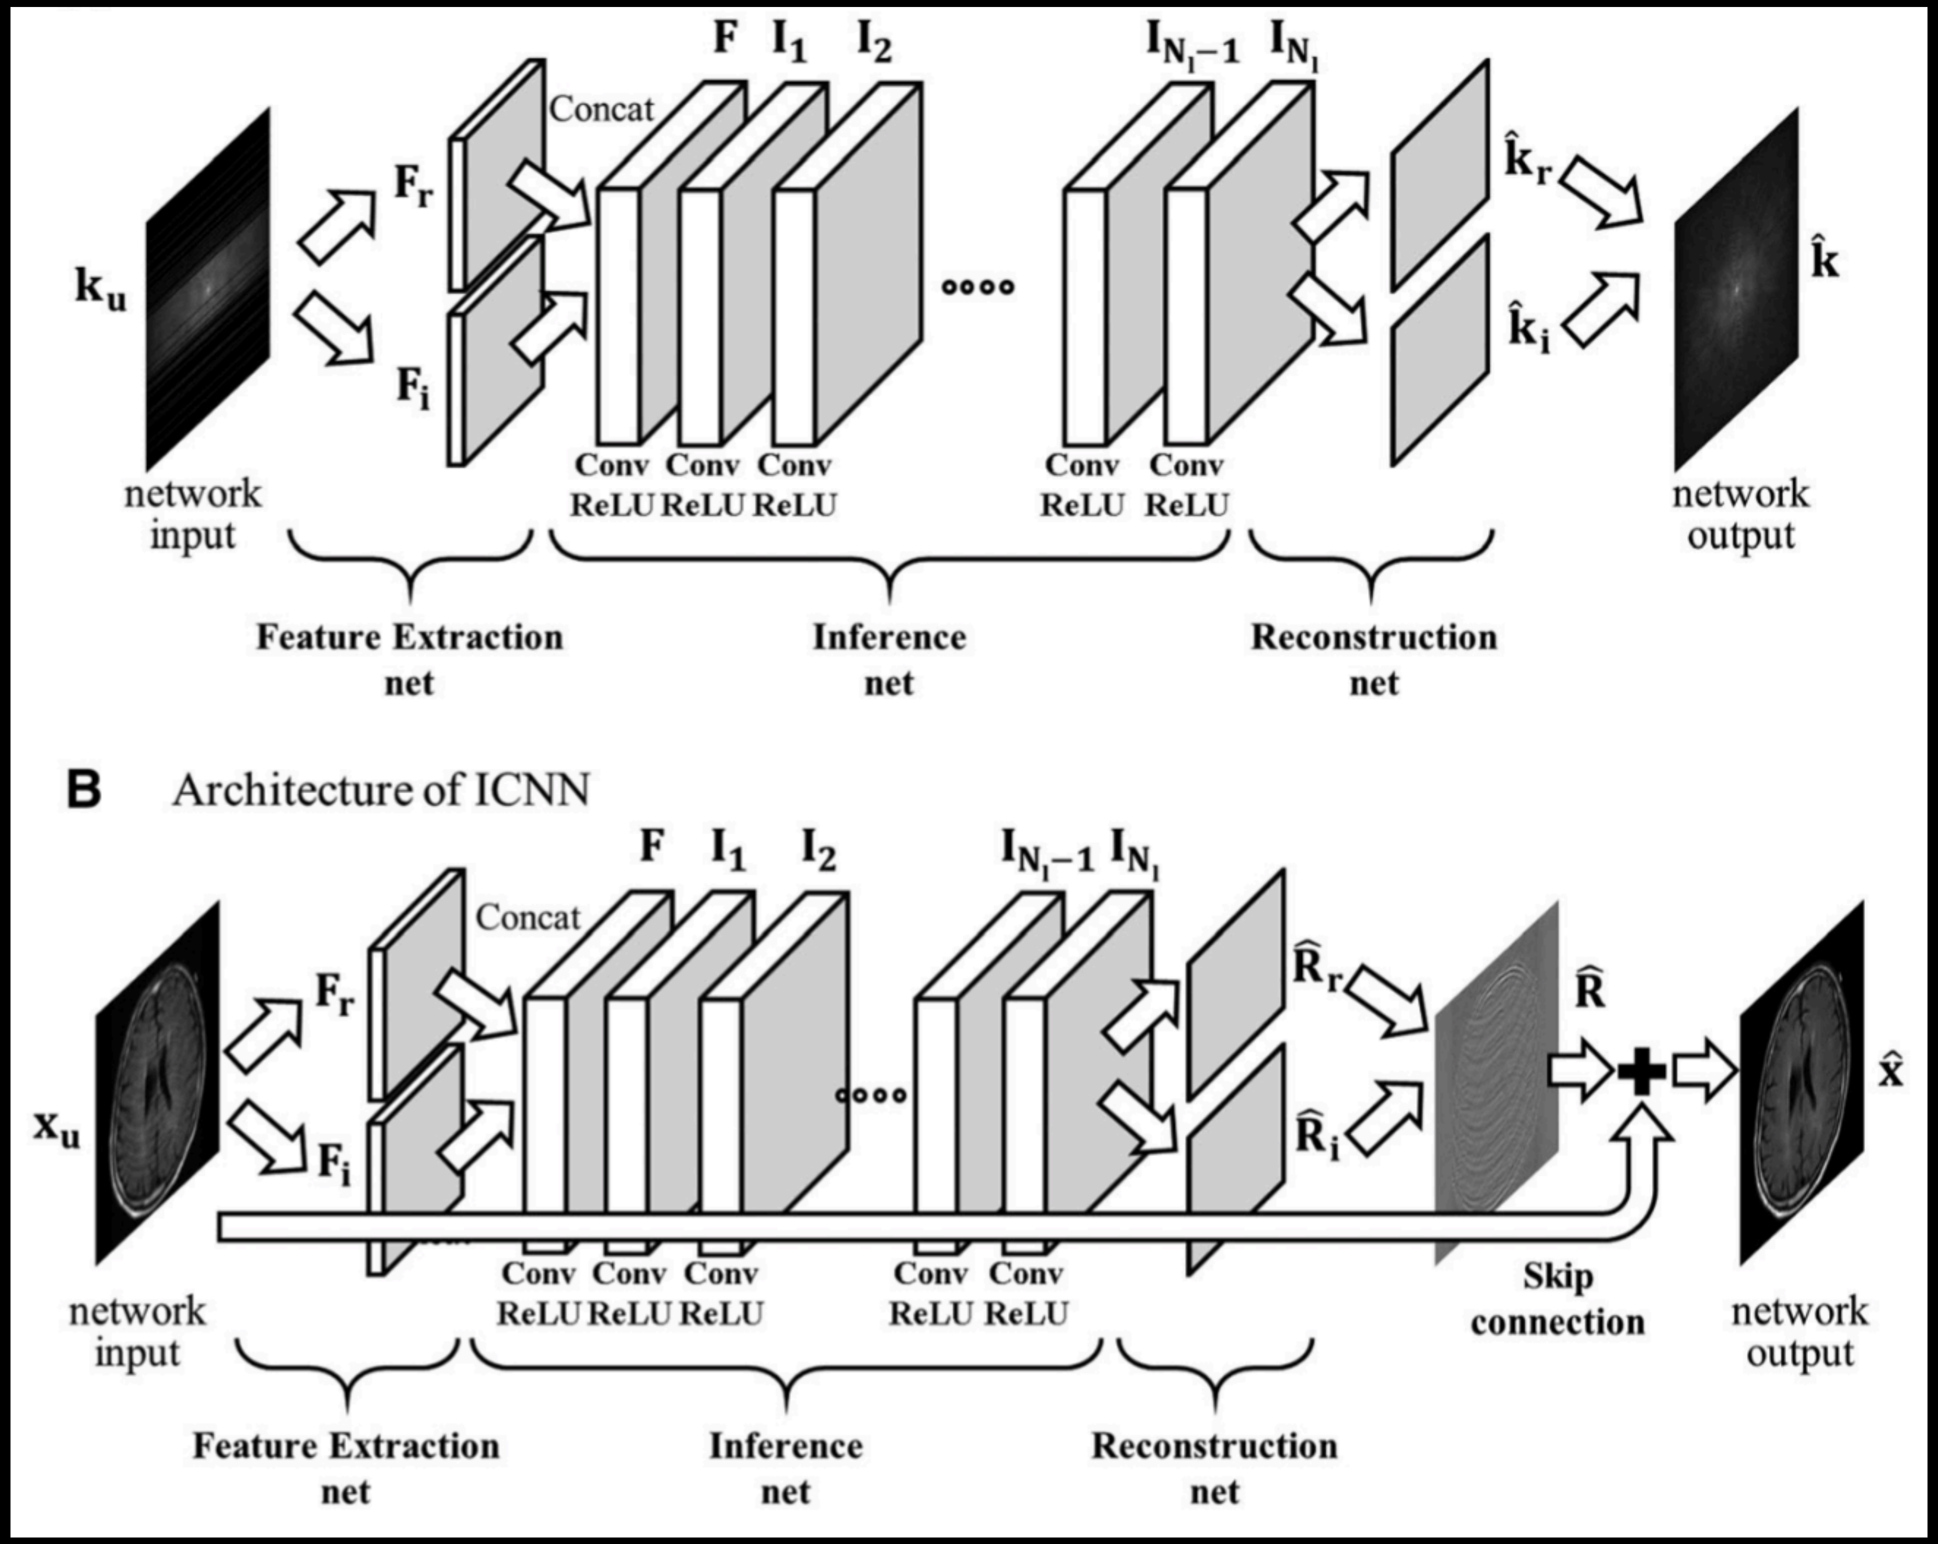

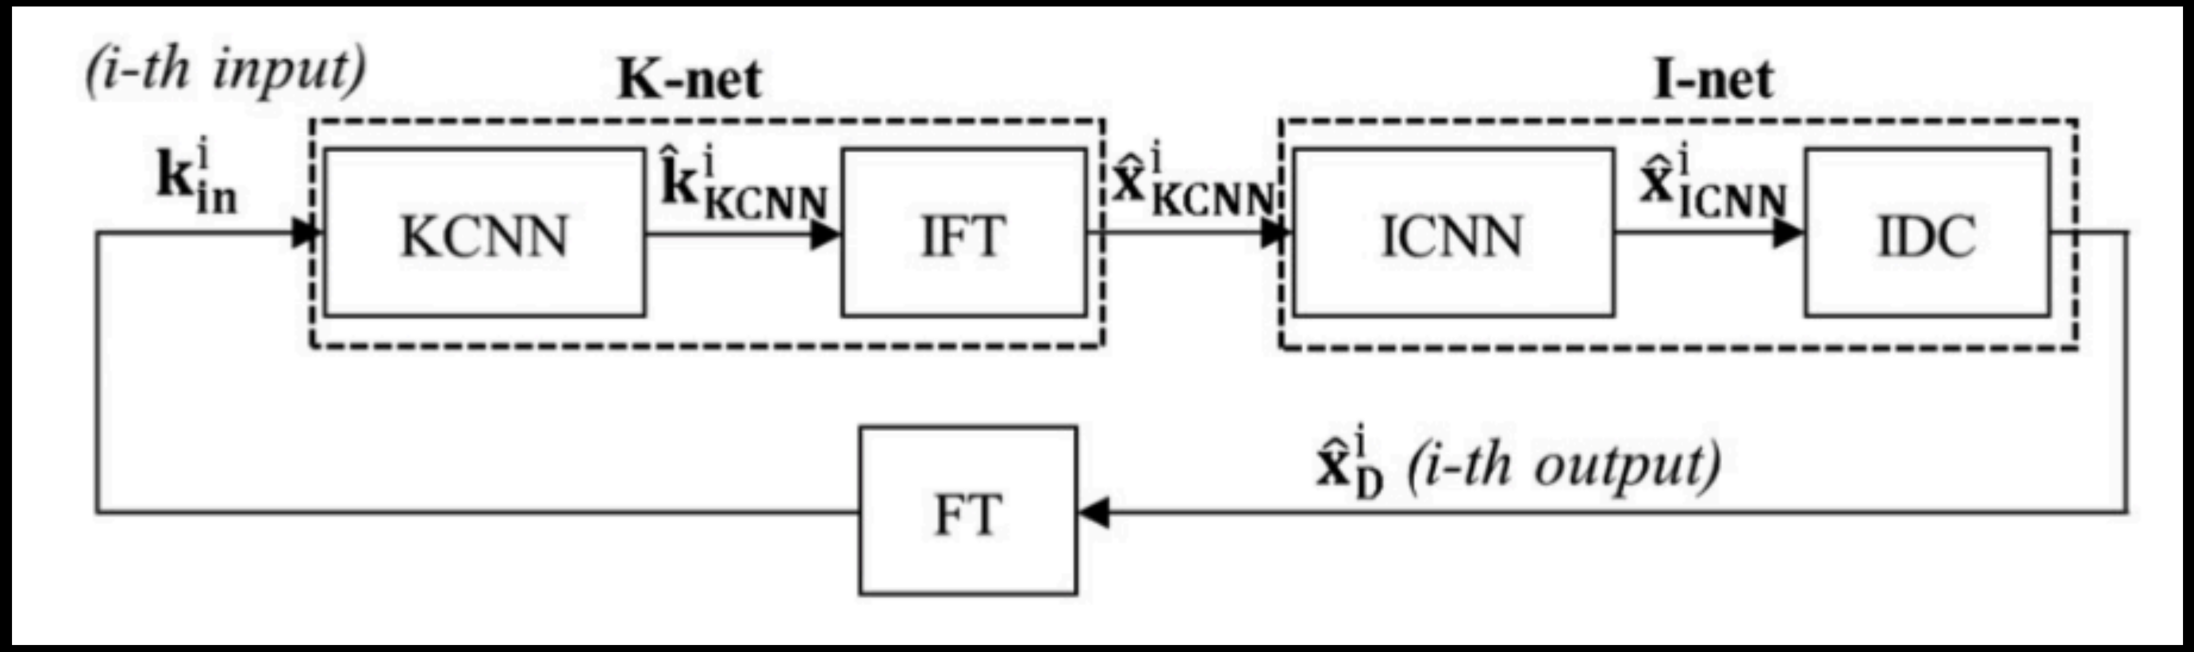

[1] Eo et al., MRM, 2018

Training in dual domains

Results from single-domain CNN vs. cross-domain CNN
(undersampled factor R=4)

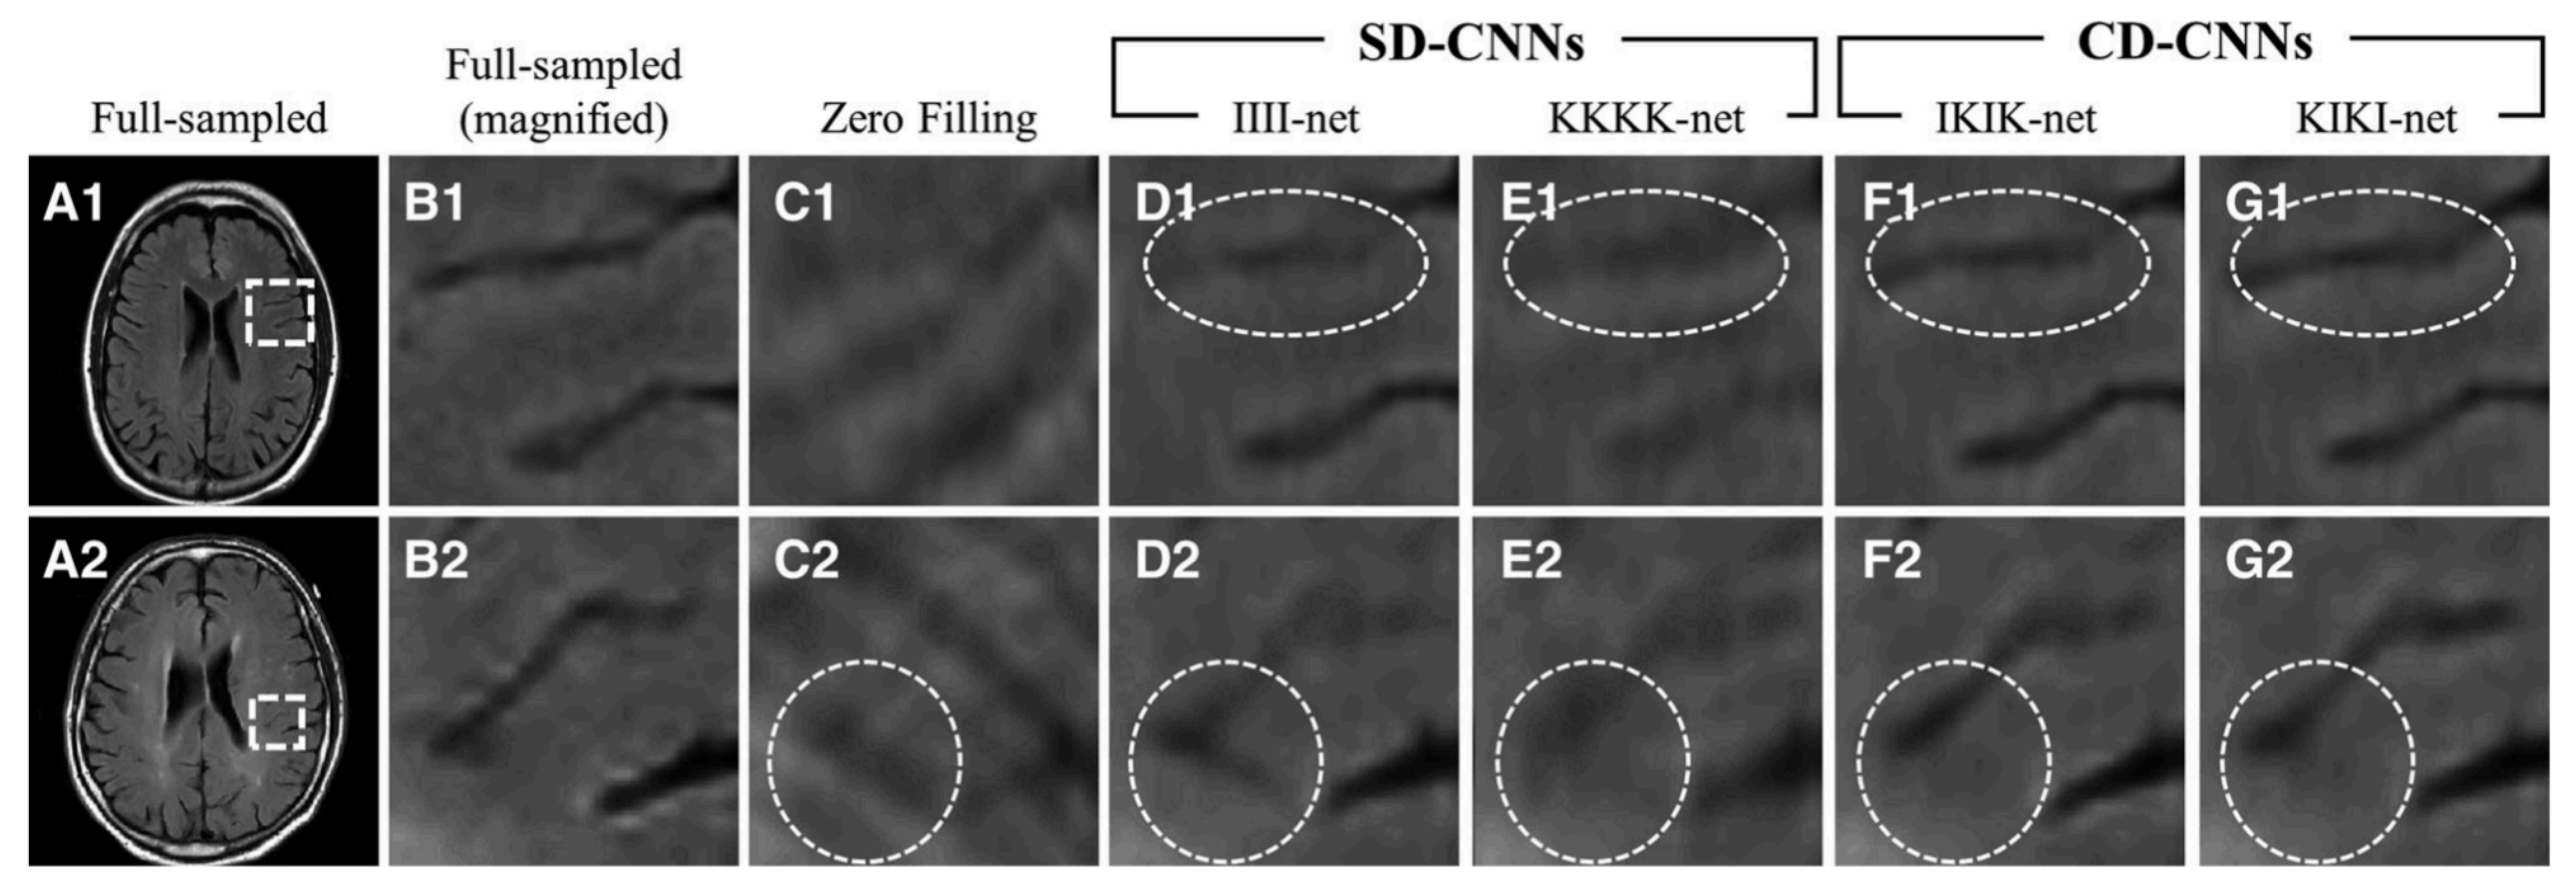

(Figure from: Eo et al., MRM 2018)

Utilizing the information from multi-contrast images

- MRI dataset sometimes involves images with multiple contrasts. Using this information shared across different contrast may reconstruction accuracy.

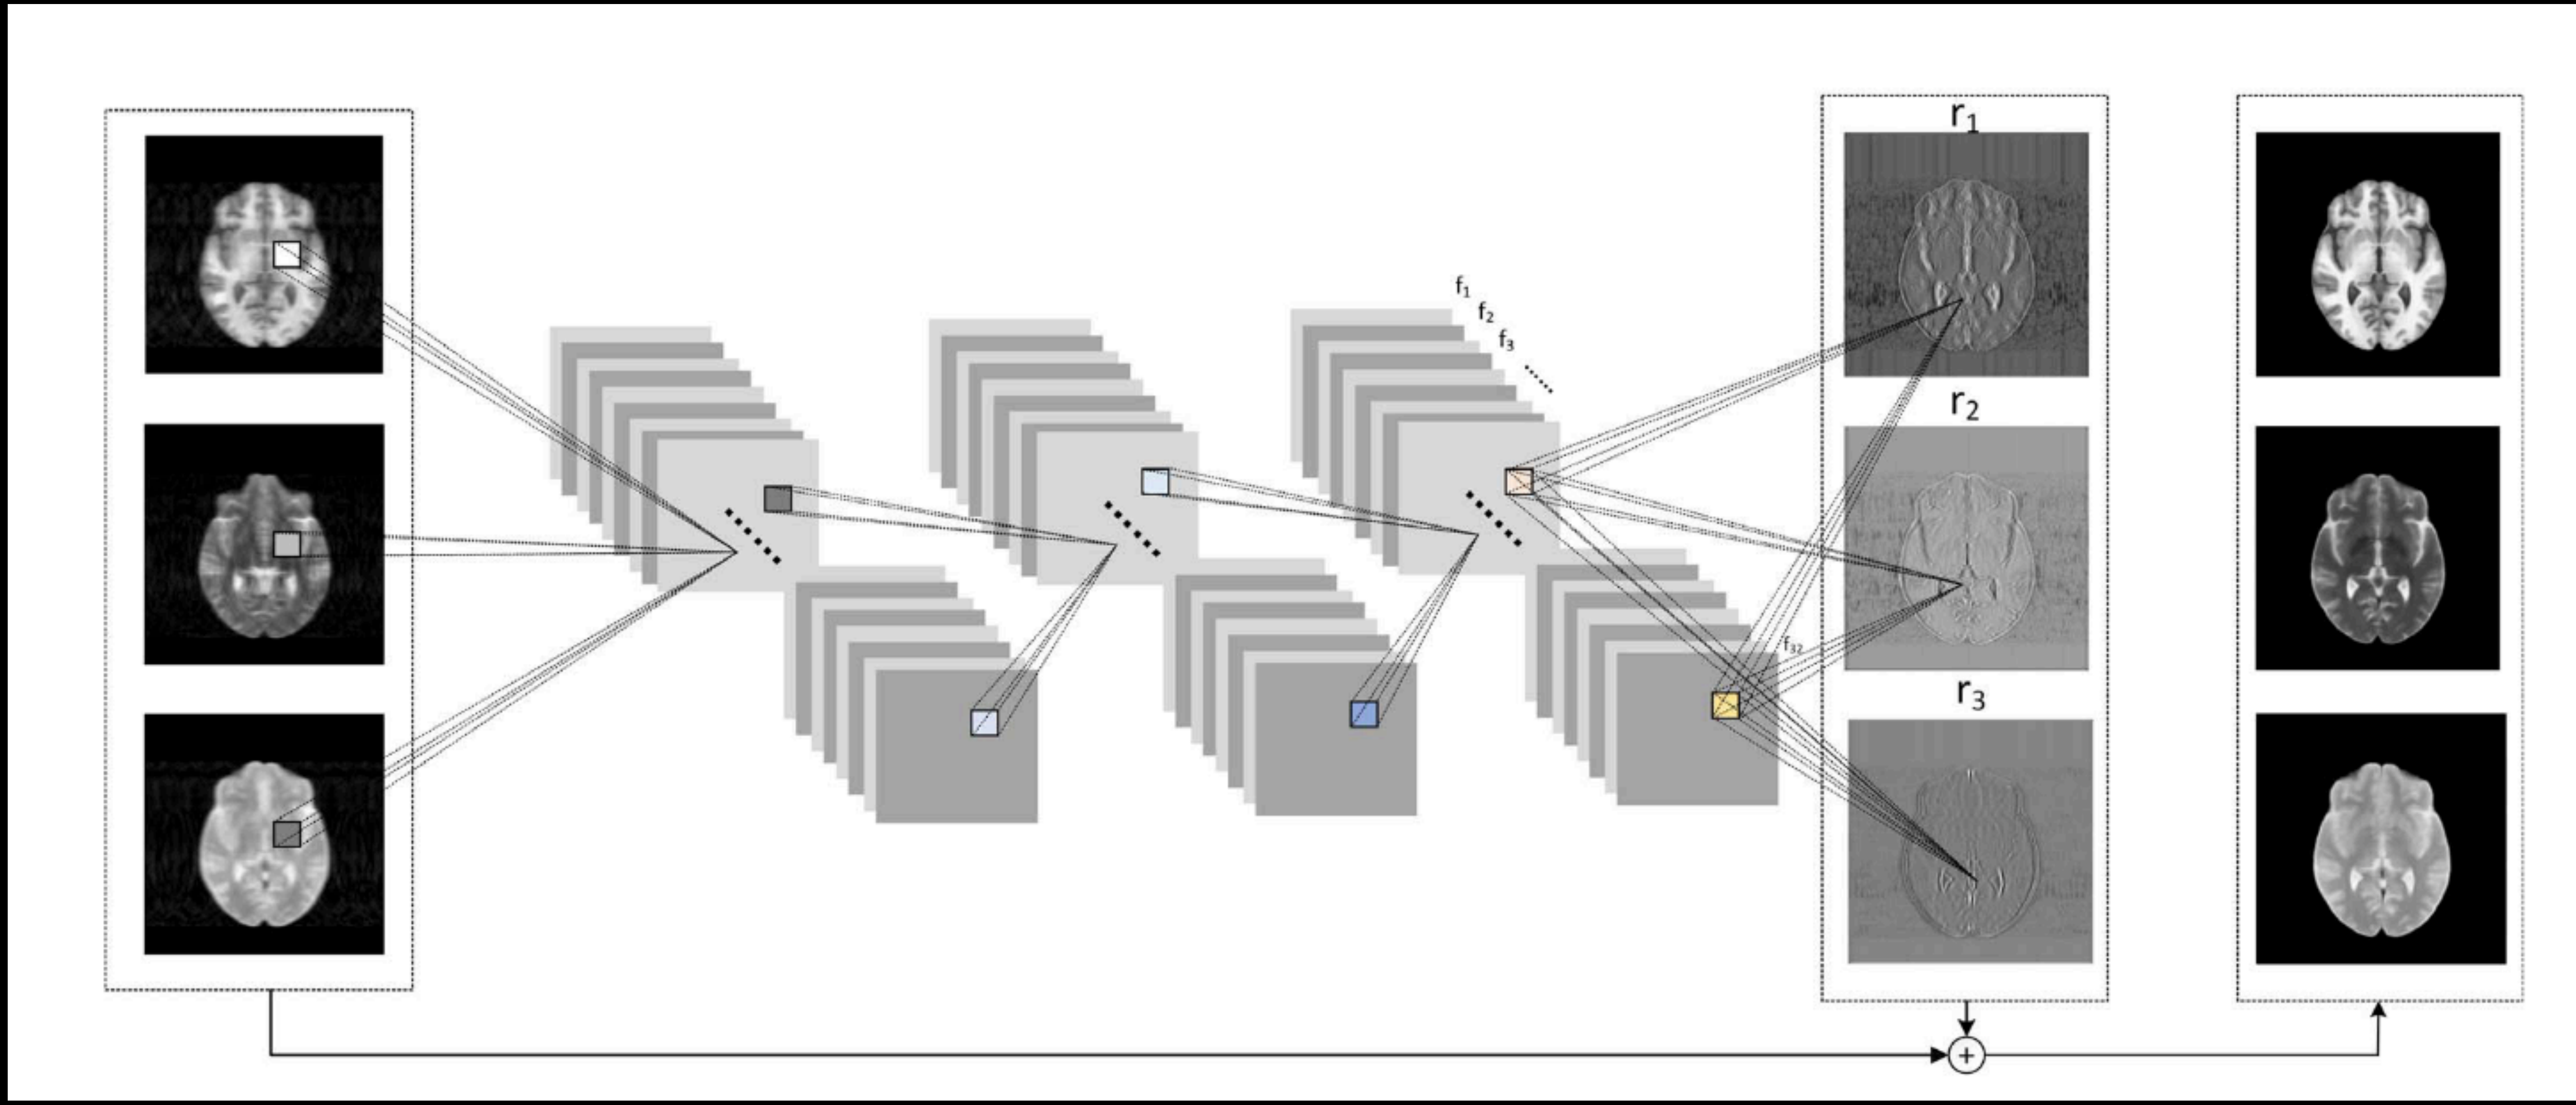

(Figures from: Sun et al., IEEE TIP 2019)

Clinical workflow considerations

- Eventually, if you want to make your DL applications useful and compatible with the clinical workflow, you may need to consider…

- (1) Integration with existing systems: PACS compatibility, DICOM compliance, seamless integration
- (2) Reconstruction speed: low latency, hardware efficiency
- (3) Robustness and generalizability: to be applied in different sequences, different scan setups

Part 4: Challenges of deep learning MRI reconstruction

Challenges of DL MRI reconstruction

- (1) Hallucinations
  - Unrealistic image features may appear on the reconstructed images
- (2) Data scarcity
  - Healthcare data are more sensitive and public datasets are less available than other computer vision tasks
  - Fully-sampled clean data may not be available due to physics limitations
- (3) Generalizability
  - Some DL models will be rather data dependent and may not generalize well to all sequences or body parts
- (4) Interpretability
  - The “failure mode” for DL recon is sometimes not clear, compared to conventional model-based approaches
- Let’s see how these problems can be (partially) mitigated

Can we reduce the occurrence of hallucinations?

- Hallucinations may be inevitable.
- However, this effect may be reduced through training with a large datasets and better training strategies.
- Furthermore, we can perform “perturbation analysis” to study how a trained network model can distort the images.

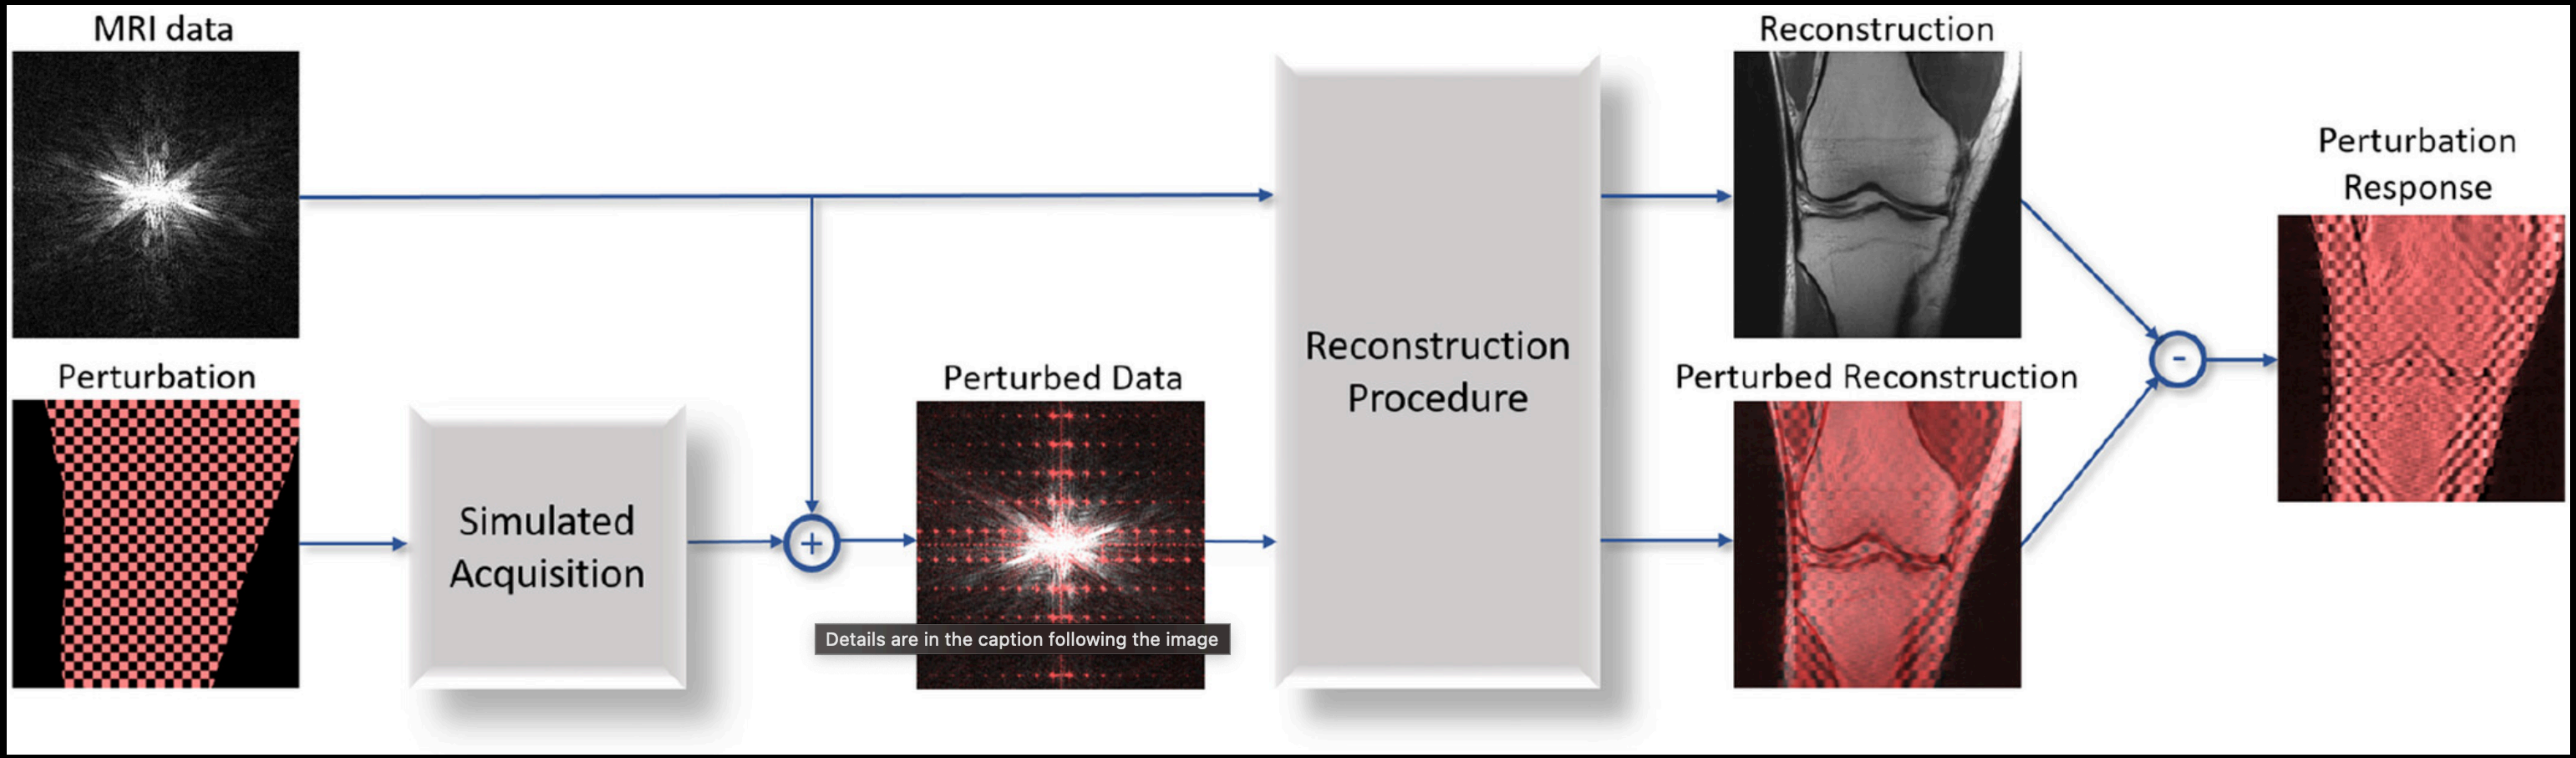

(Figures from: Chang et al., MRM 2021)

Perturbation analysis

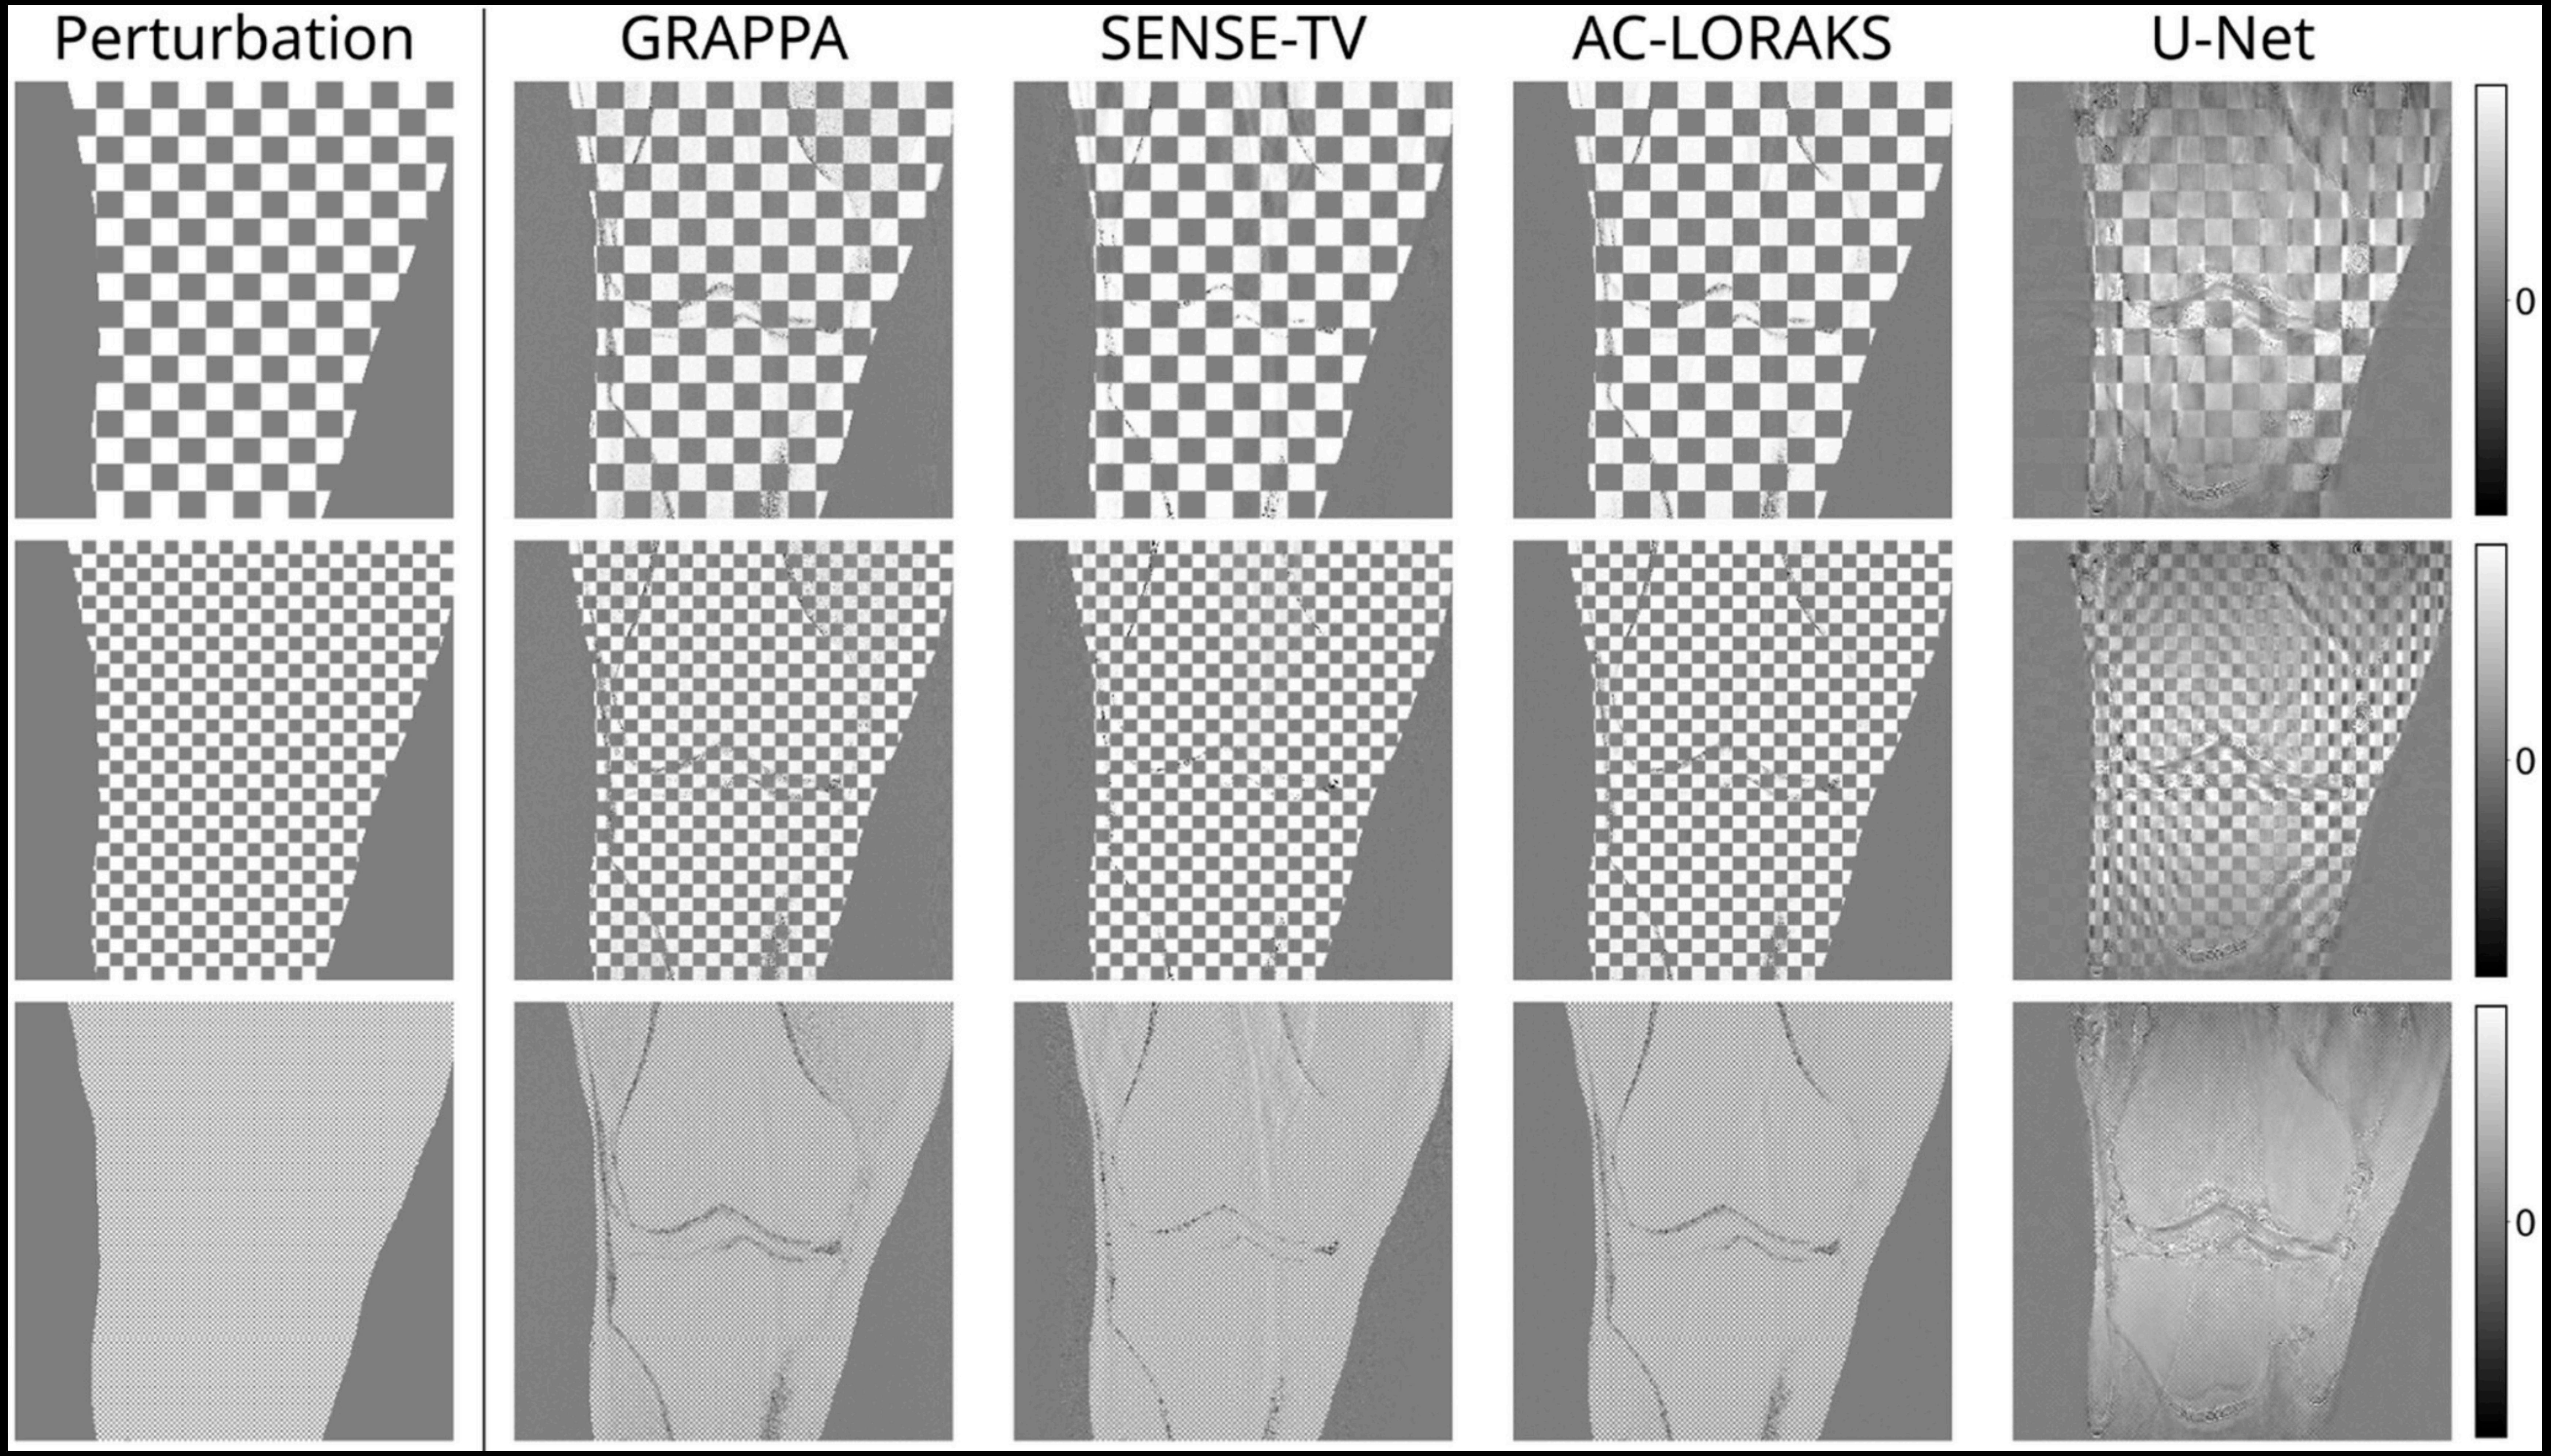

(Figures from: Chang et al., MRM 2021)

Self-supervised training with limited data

- Self-supervised physics-guided reconstruction[1]
  - Deep learning reconstruction without fully-sampled reference dataset
  - Acquired k-space was split into 2 disjoint sets for self-supervision during training.

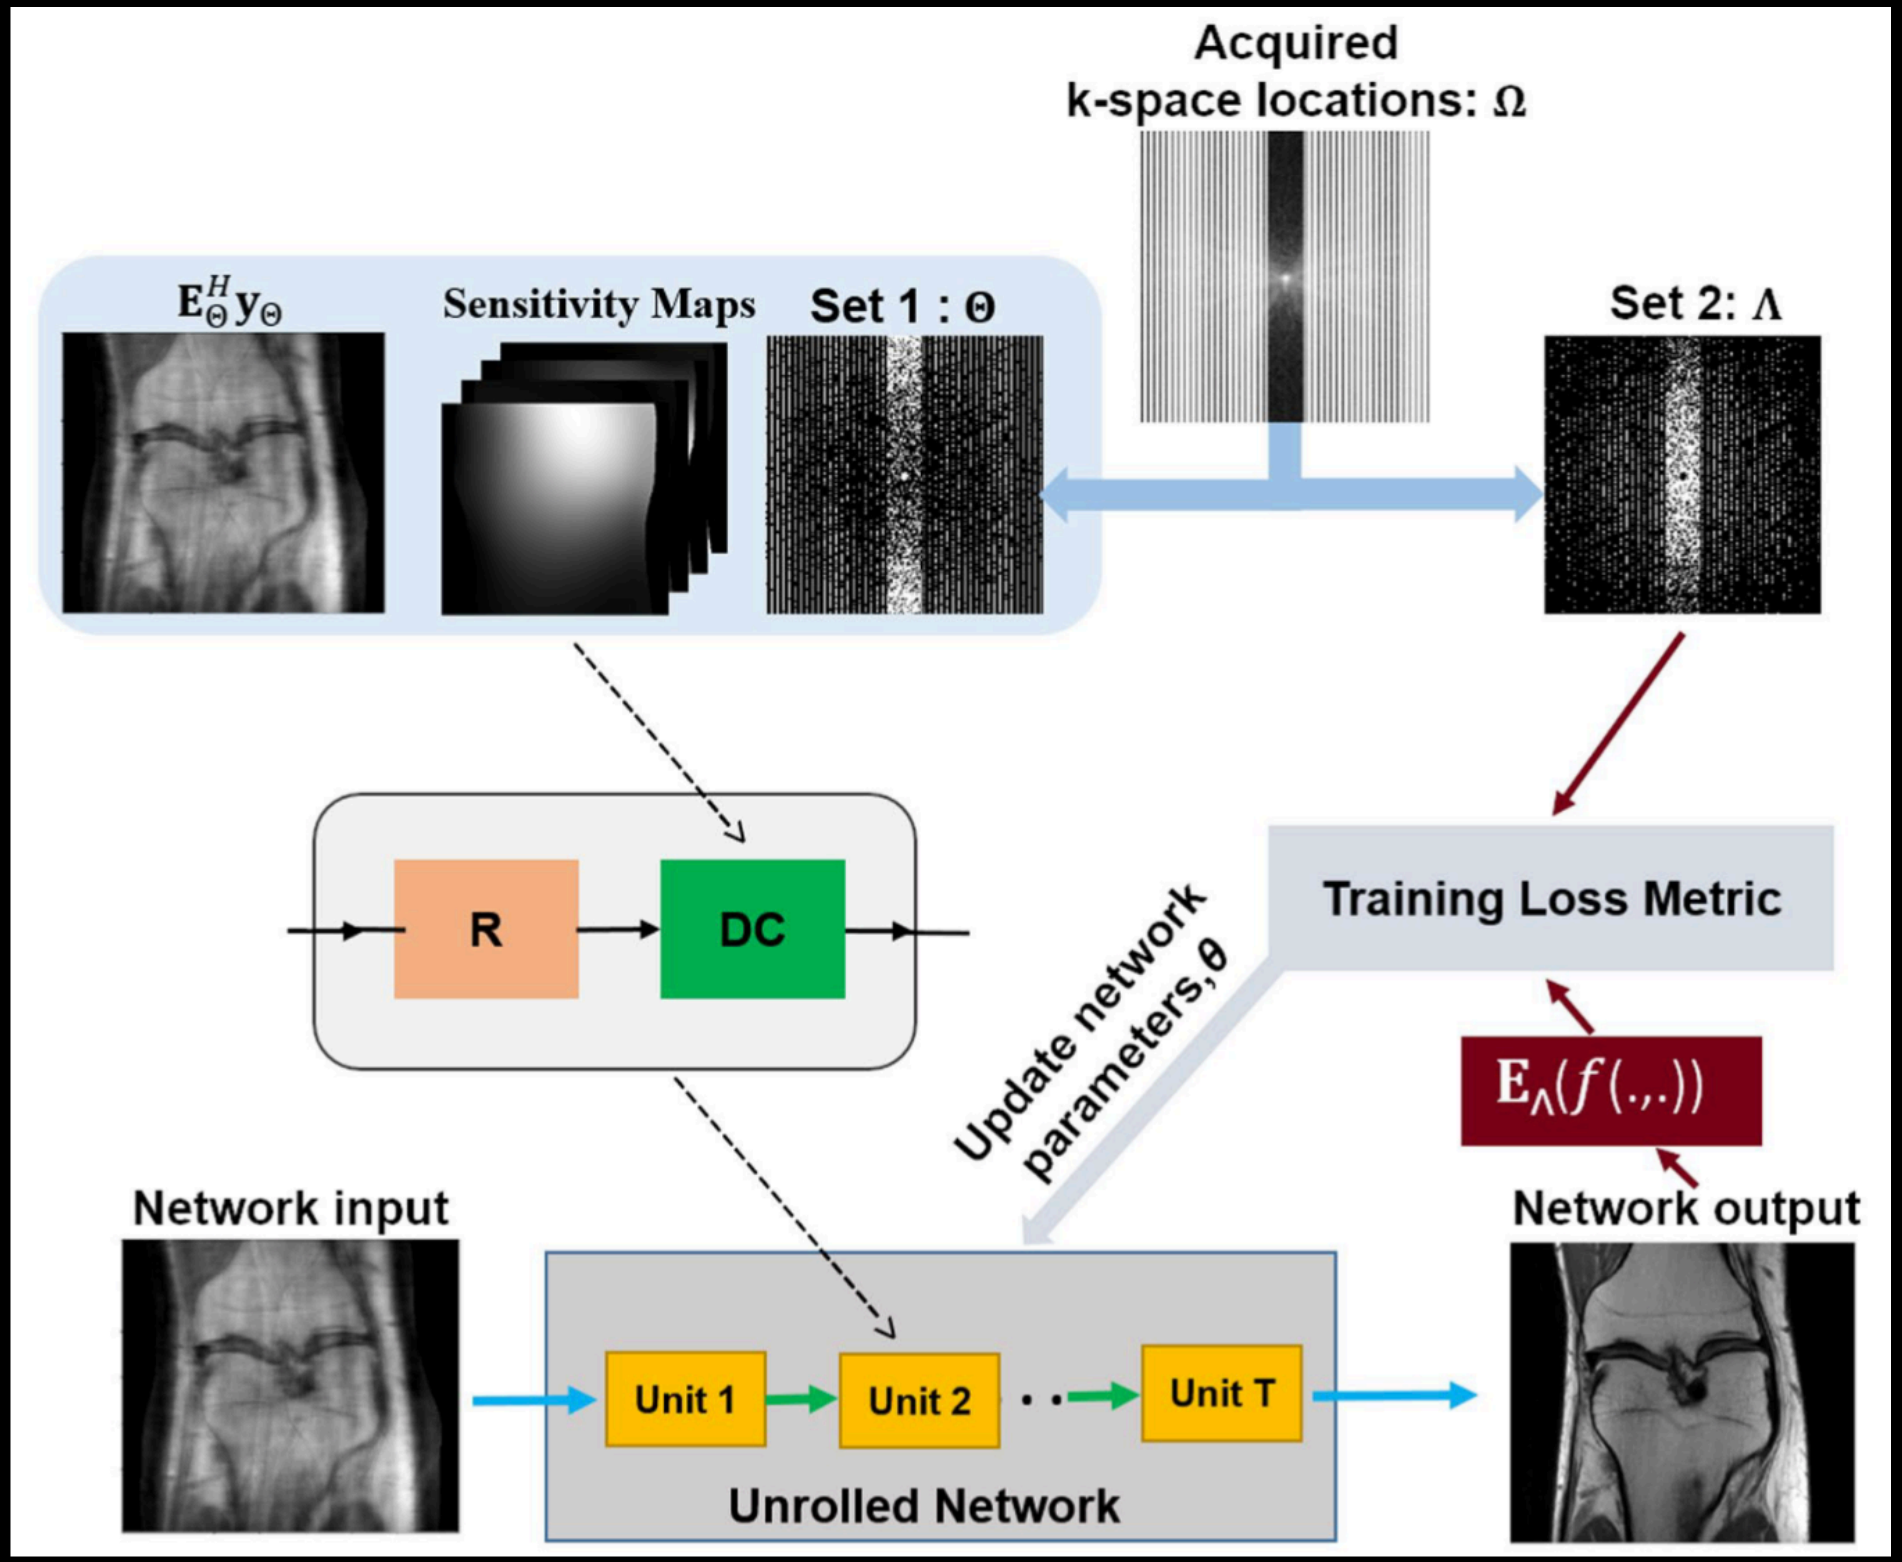

[1] Yaman et al., MRM, 2020

Self-supervised training with limited data

- Image from self-supervised learning show similar performance compared to the supervised method.

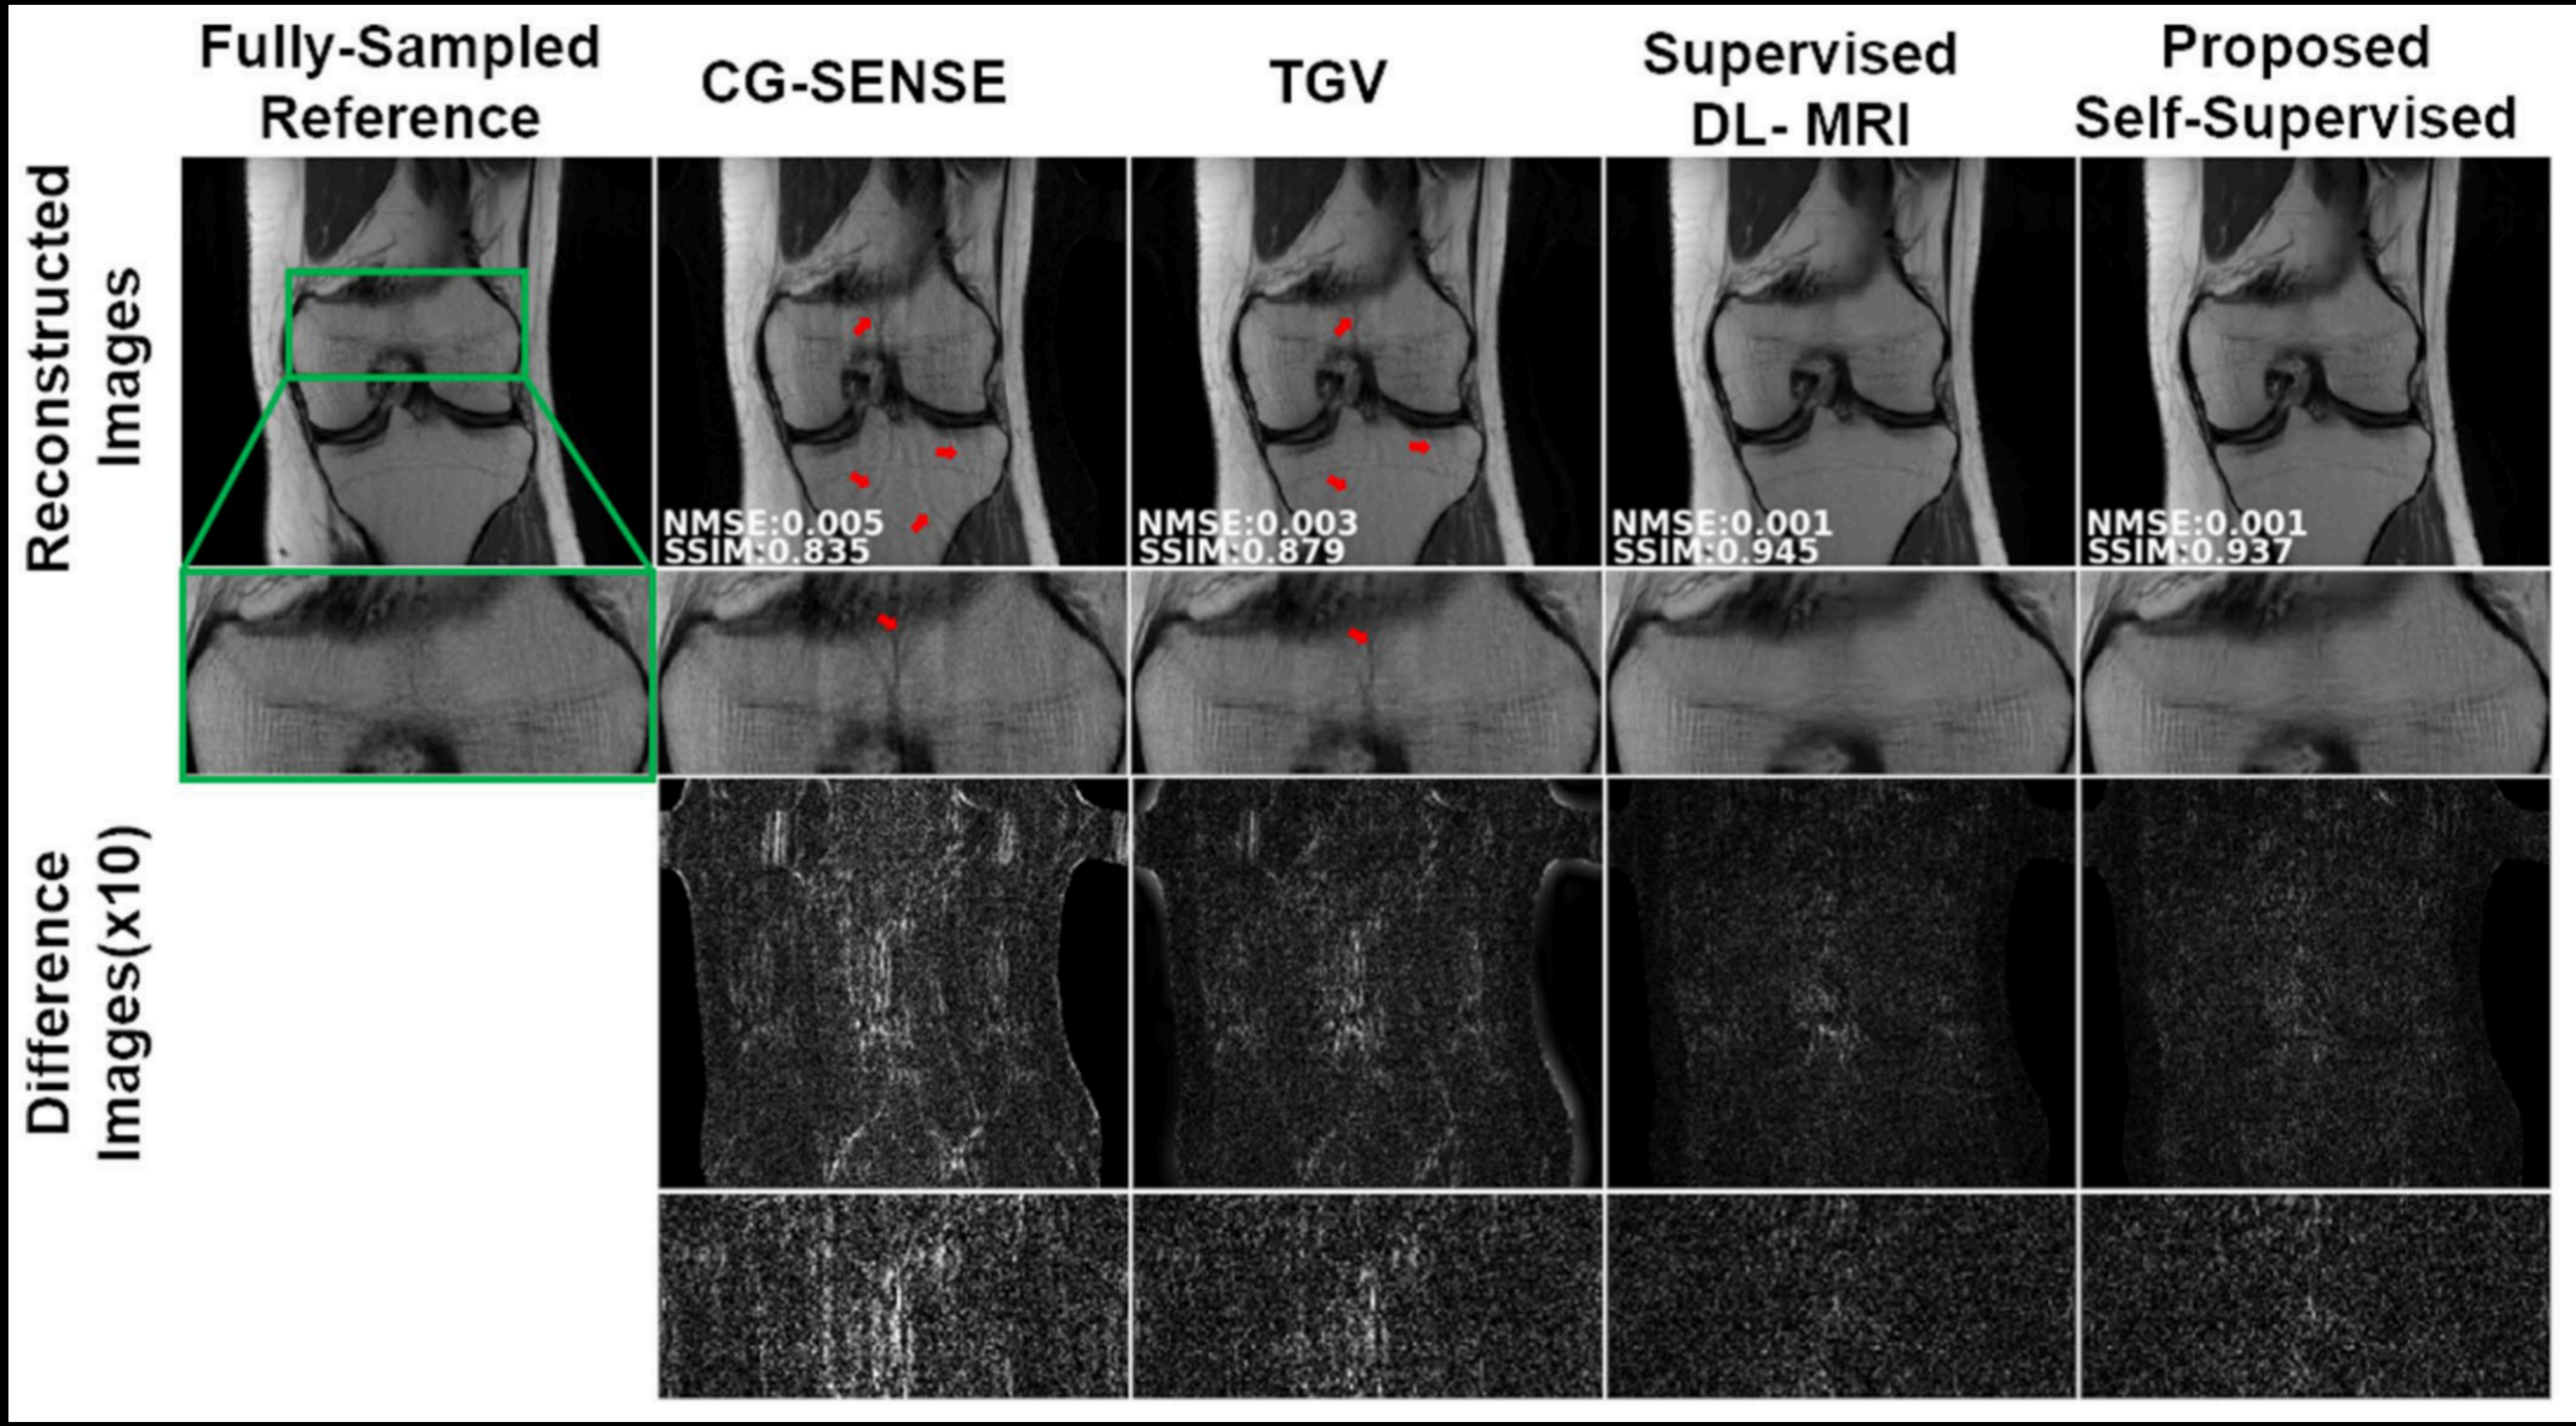

(Figure from: Yaman et al., MRM 2020)

Publicly available MRI k-space datasets

- fastMRI (https://github.com/facebookresearch/fastMRI)
  - Knee, brain and prostate MRI
- SKM-TEA (https://github.com/StanfordMIMI/skm-tea)
  - Quantitative knee MRI with tissue segmentation
- M4Raw (https://github.com/mylyu/M4Raw)
  - Multi-contrast multi-repetition 0.3 T brain MRI
- CMRxRecon (https://github.com/CmrxRecon/CMRxRecon-SciData)
  - Cardiac Cine MRI and cardiac quantitative MRI

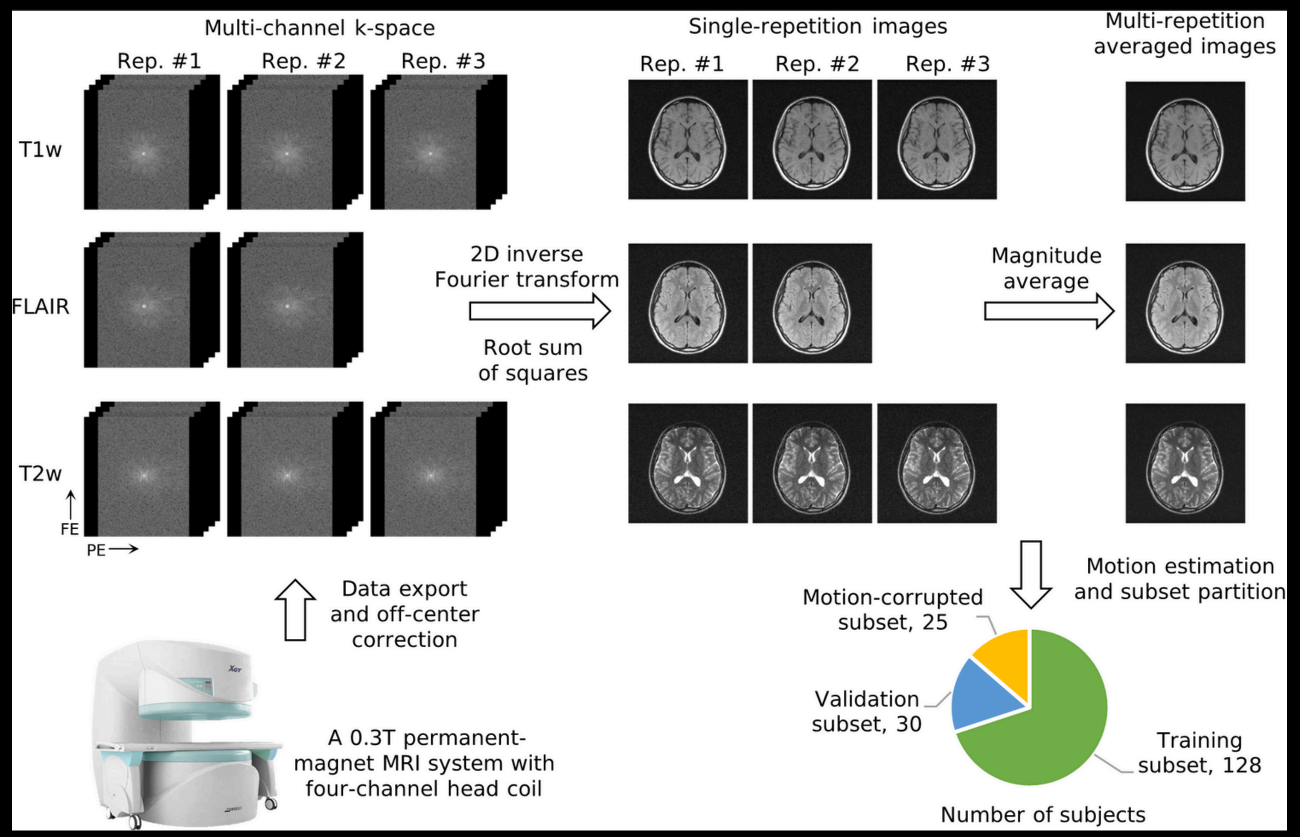

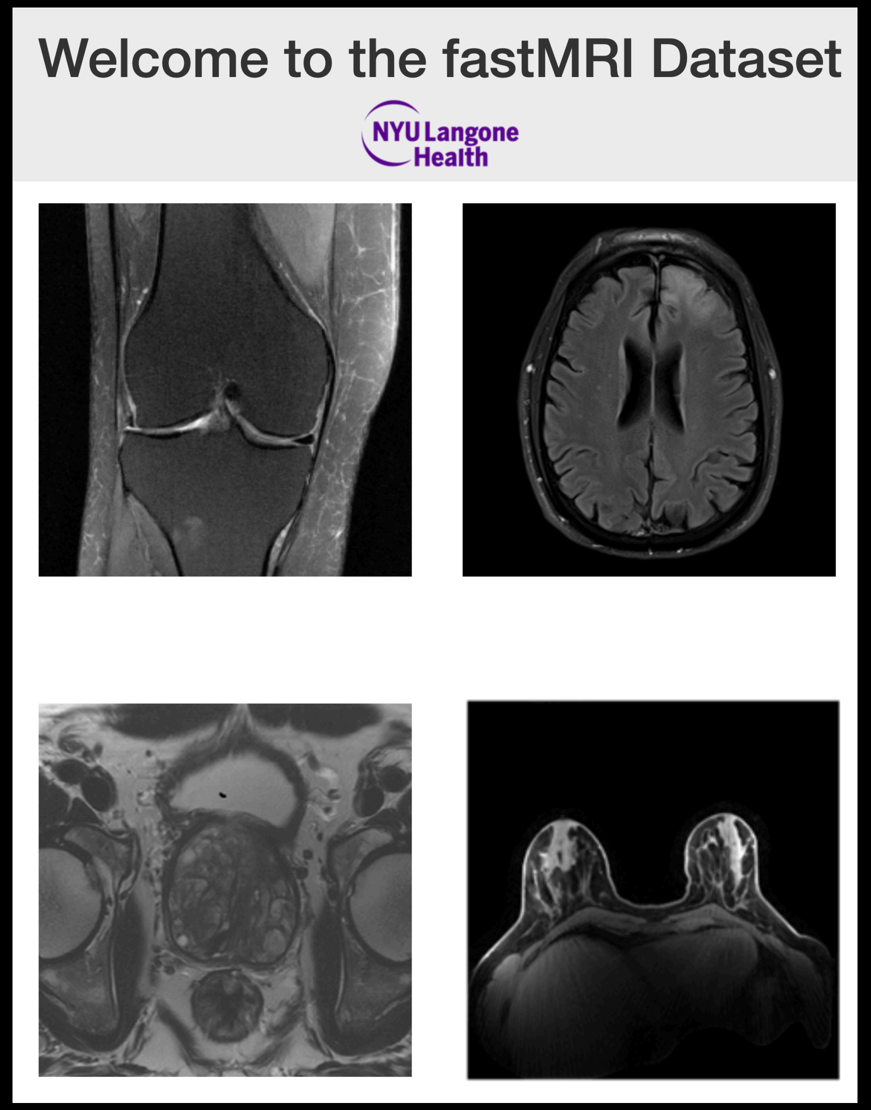

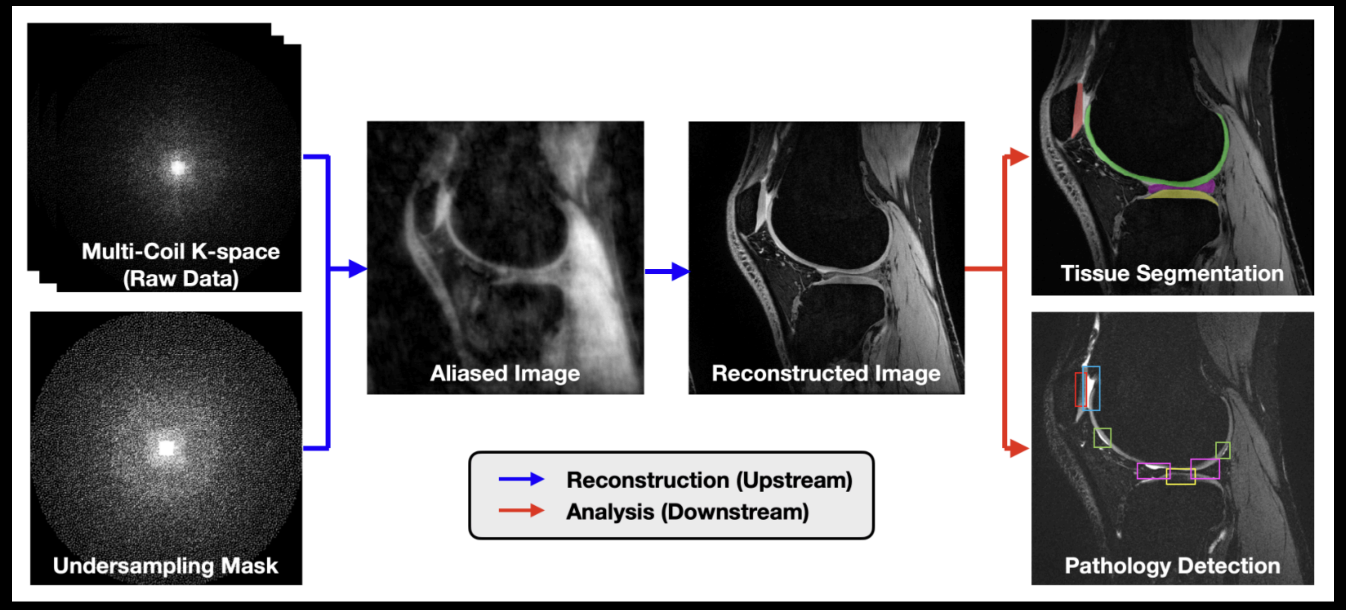

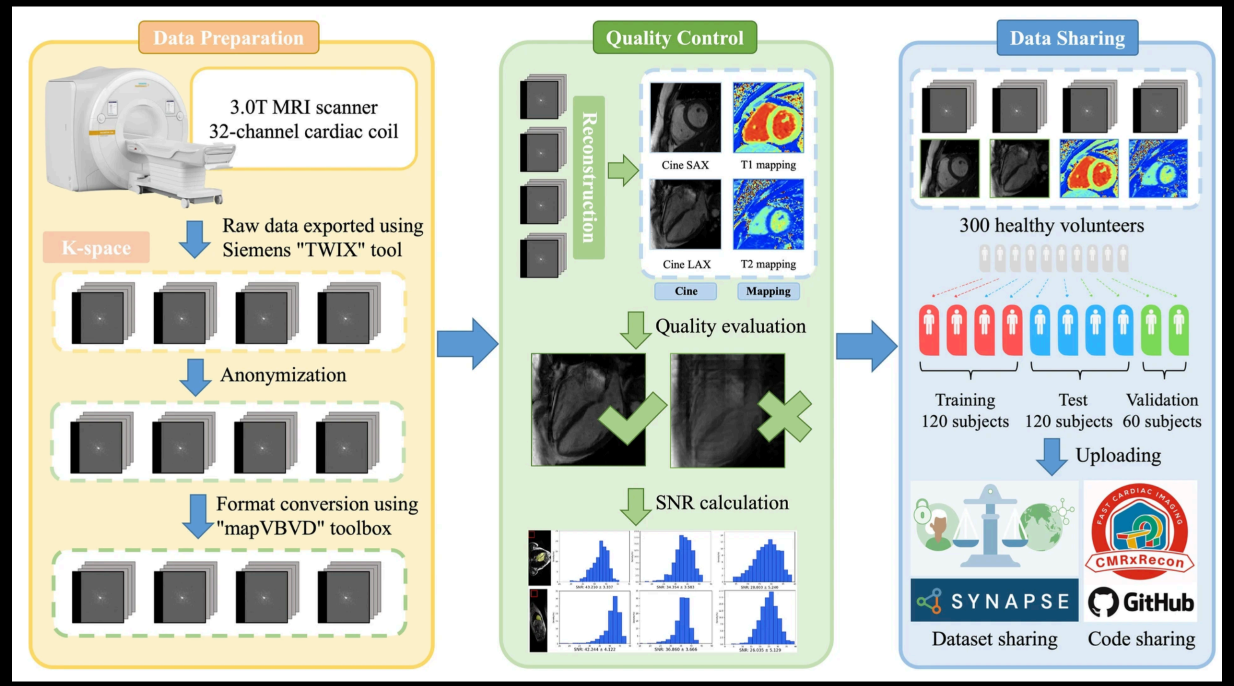

Failure mode of DL recon is not clear

- This below image is reconstructed using parallel imaging…

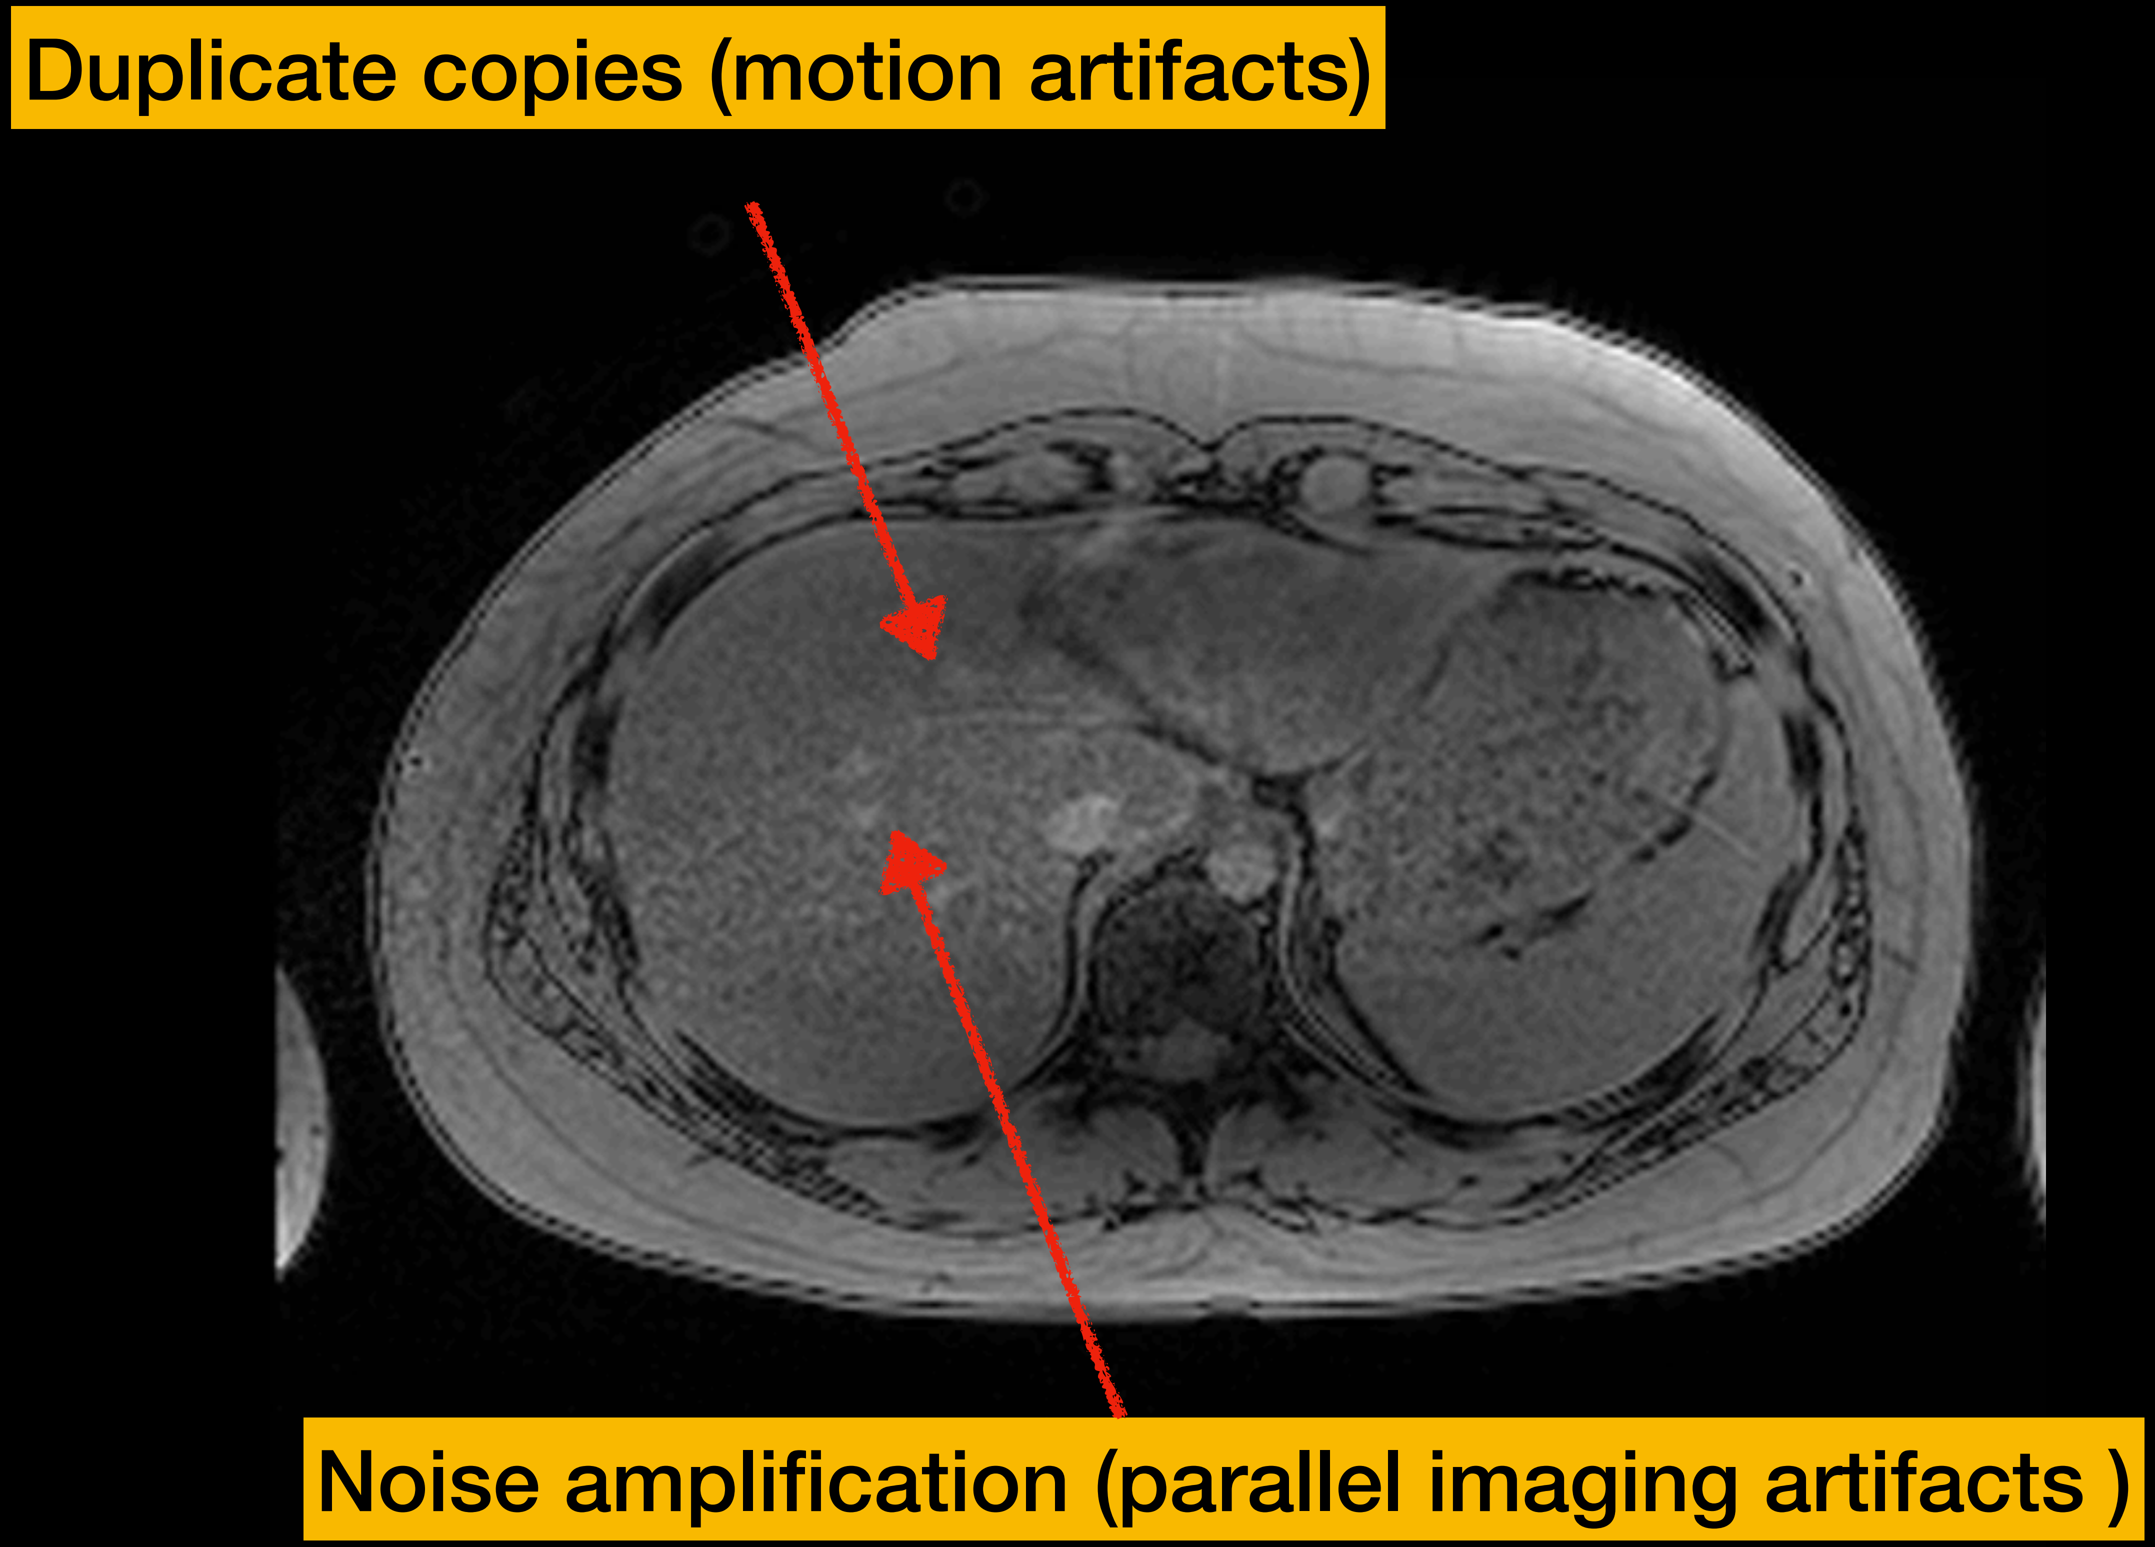

Failure mode of DL recon is not clear

- "Failure mode" of parallel imaging techniques, such as noise amplification, is pretty well-known. Radiologists may "read through" those artifacts.

- However, when and how DL recon can fail is still not clearly known…

Uncertainty quantification in DL MRI reconstruction

- UP-Net (Uncertainty-aware Physics-driven deep learning network)
  - Uncertainty information incorporated into deep learning-based artifact suppression and parameter mapping

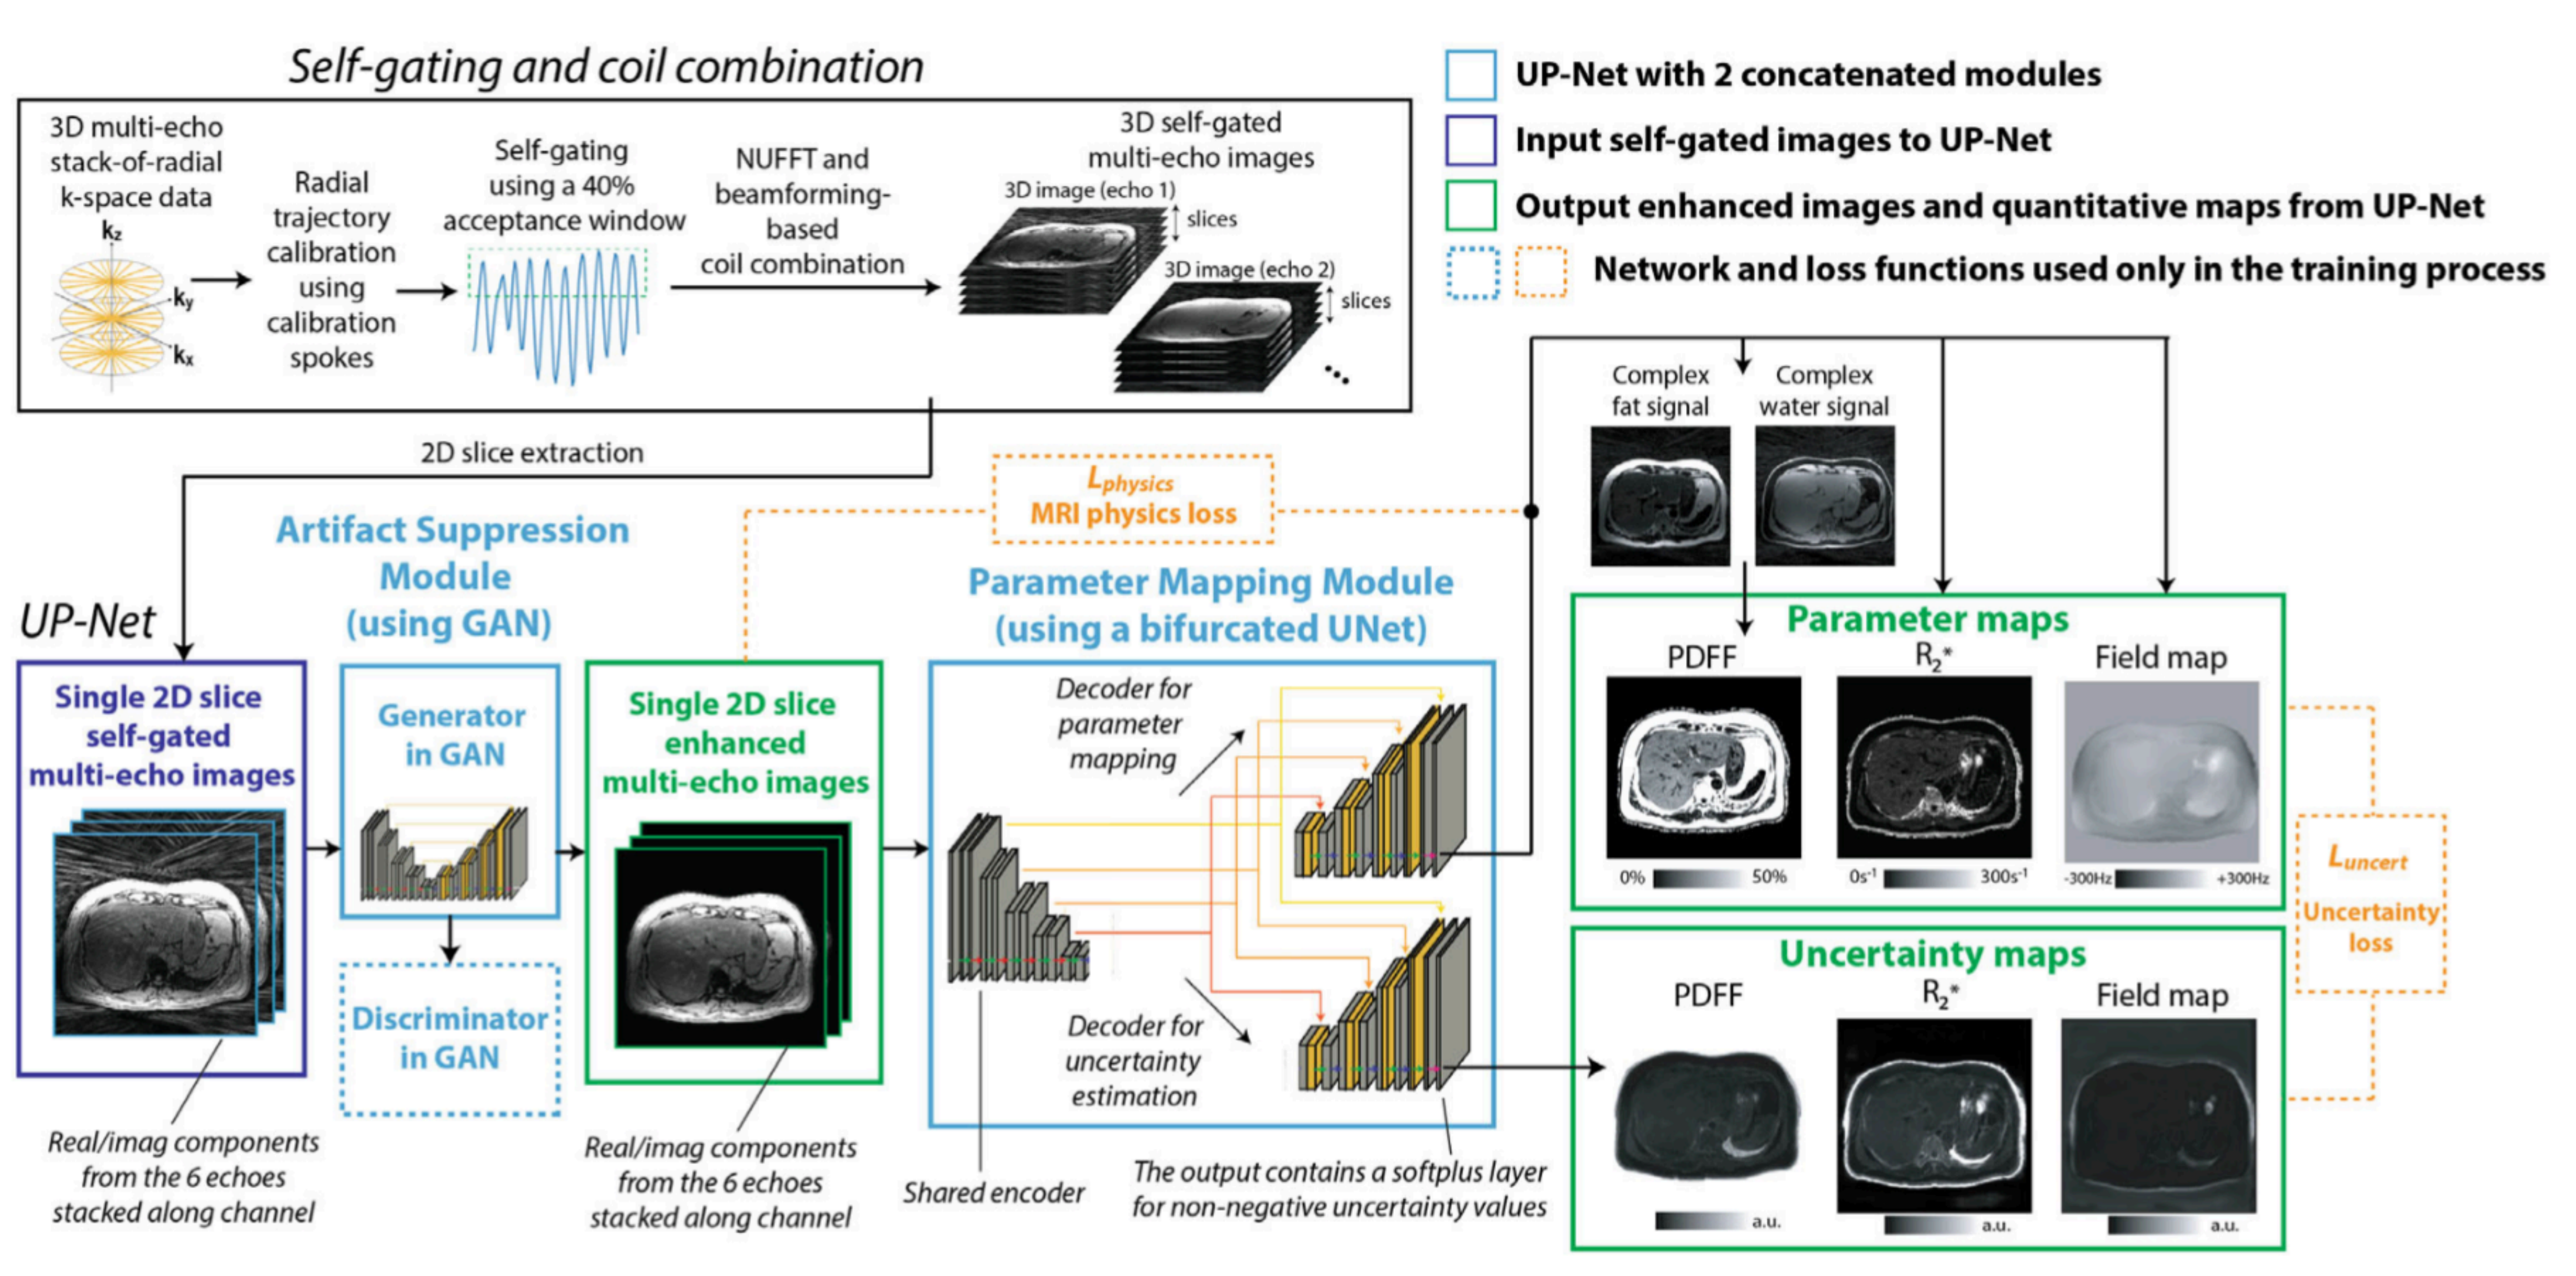

$$L_{\text{uncert}} = \frac{\|\hat{p} - p\|_1}{\hat{u}} + \log(\hat{u}).$$

[1] Shih et al., MRM, 2023

UP-Net

- Additional uncertainty map provided by the deep learning network can be used to estimate errors in the deep learning results

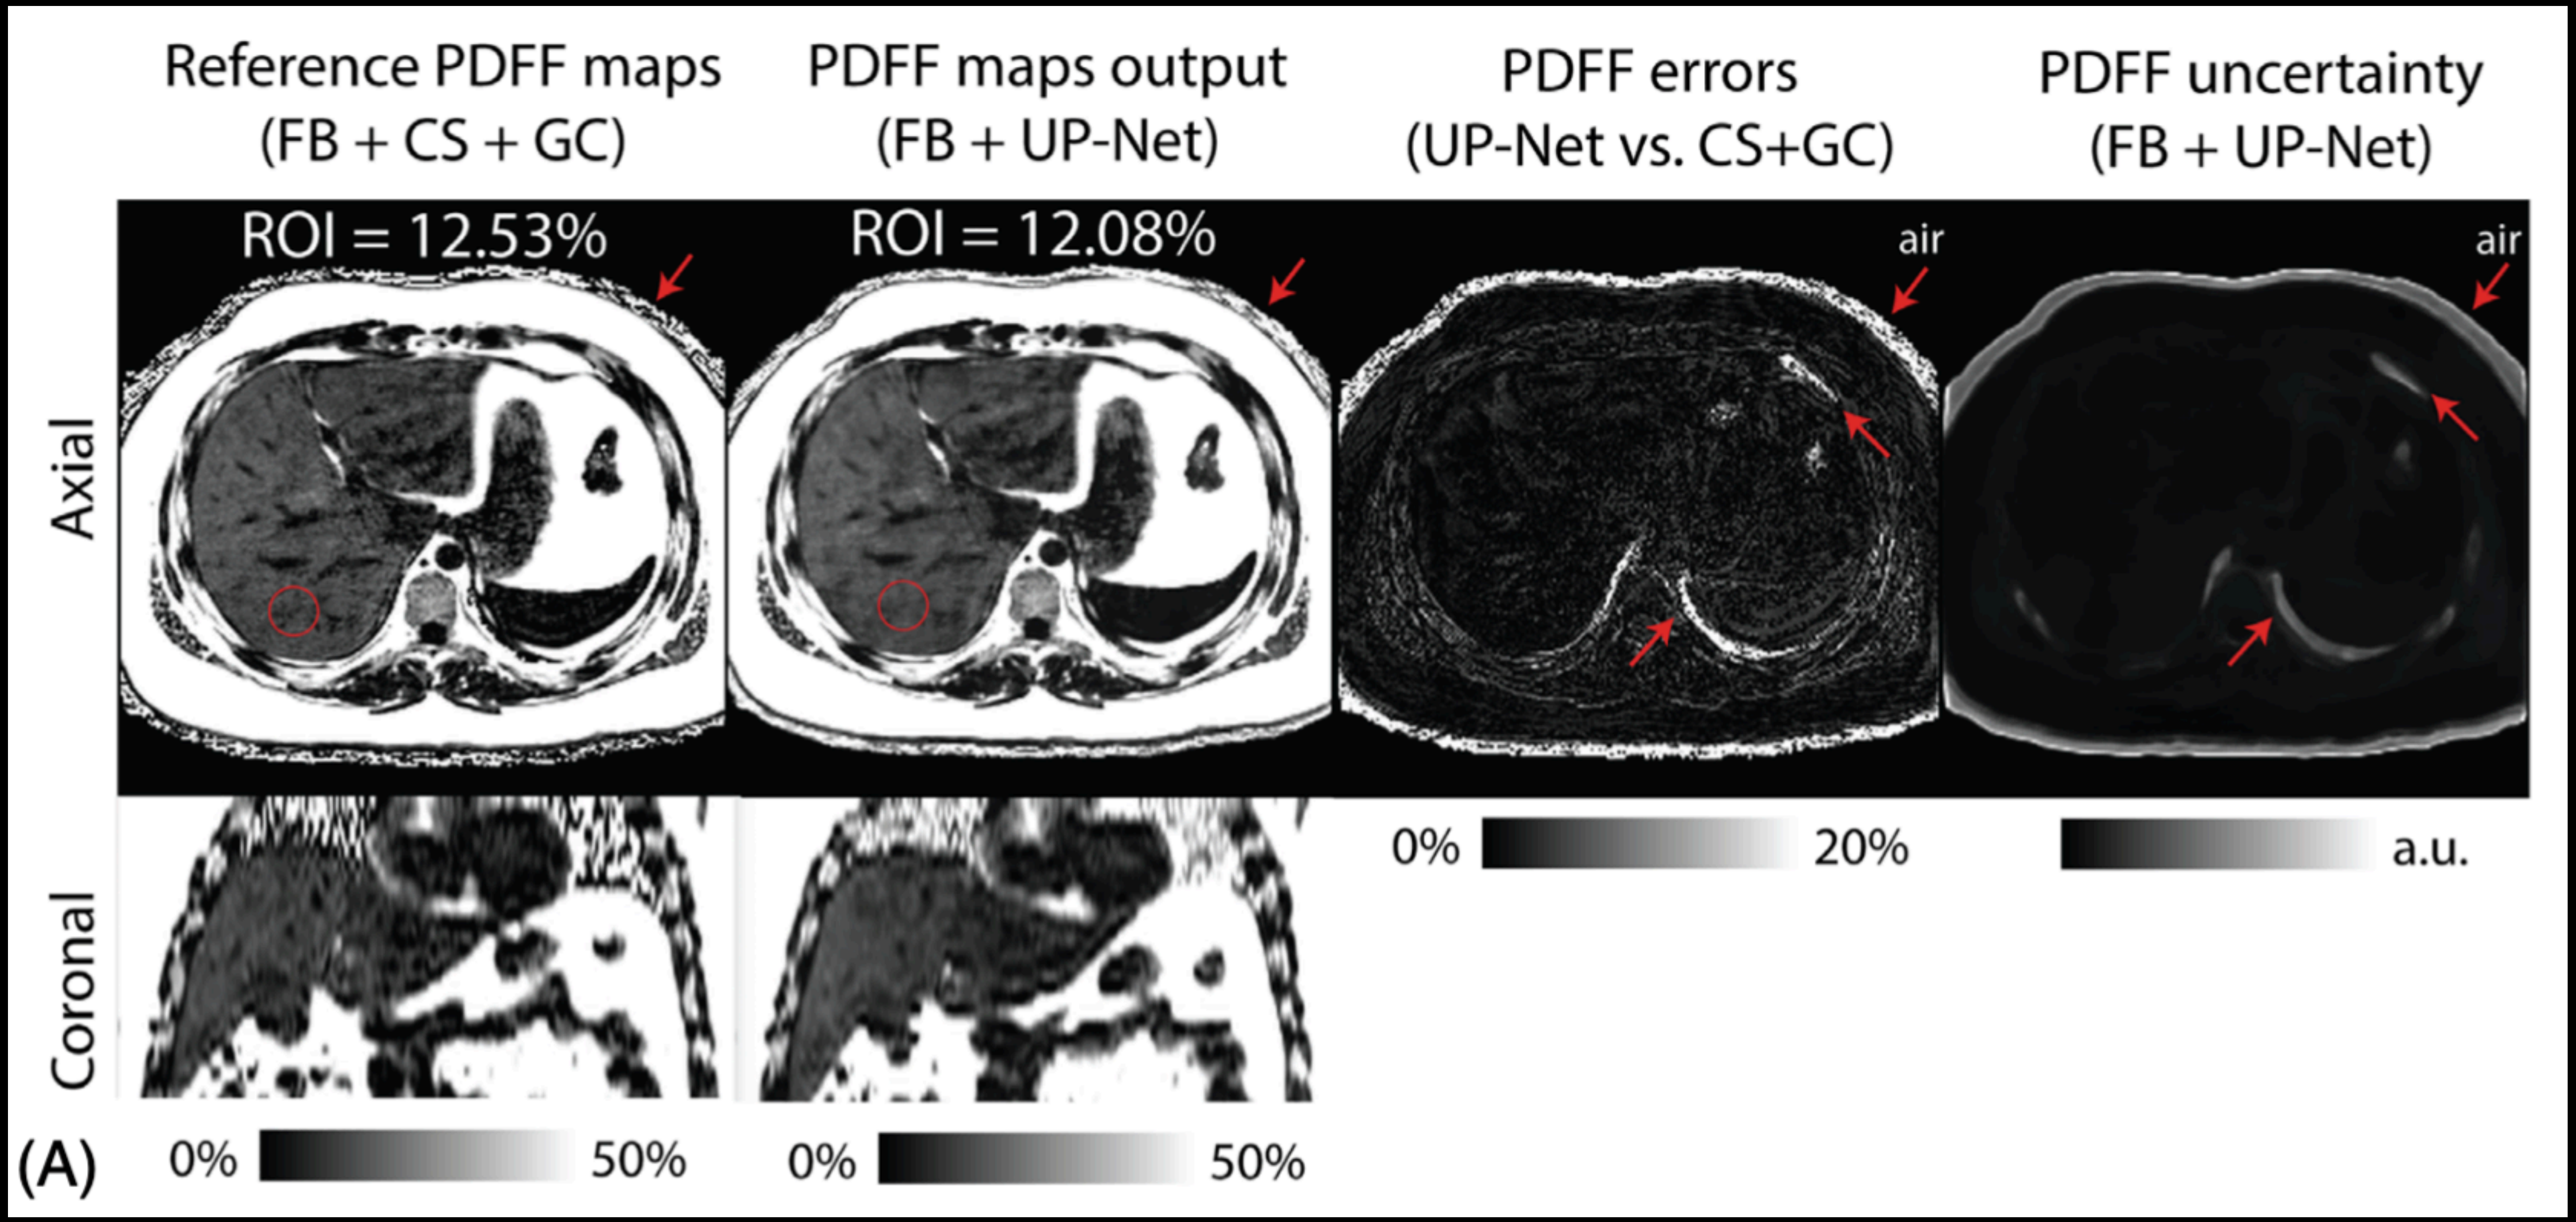

(Figure from: Shih et al., MRM 2023)

Part 5: Discussion

Discussion

- Major MRI vendors have started to provide deep learning reconstruction products that can reduce scan time or reduce image noise.
- The products may still be limited to certain sequences or body parts.

Siemens - Deep Resolve

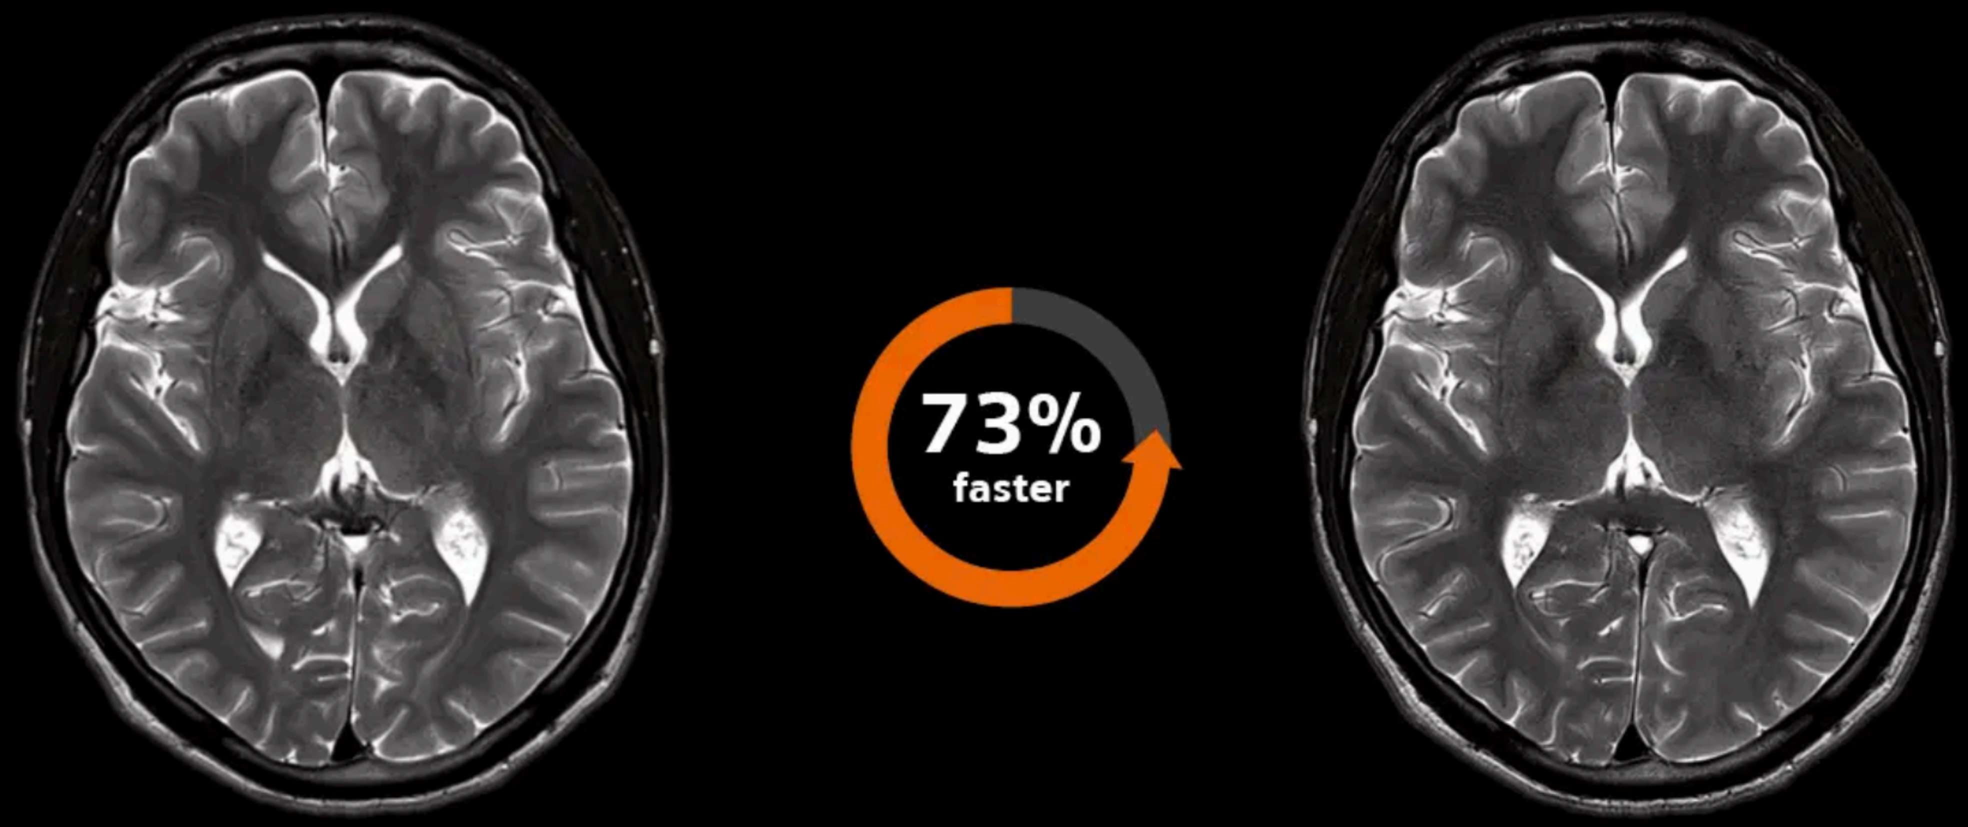

Conventional
MAGNETOM Vida
PAT 1, TA 2:12 min
28 slices, 0.4 x 0.4 x 4.0 mm³

Deep Resolve
MAGNETOM Vida
PAT 4, TA 0:36 min
28 slices, 0.2 x 0.2 x 4.0 mm³

GE - AIR Recon DL

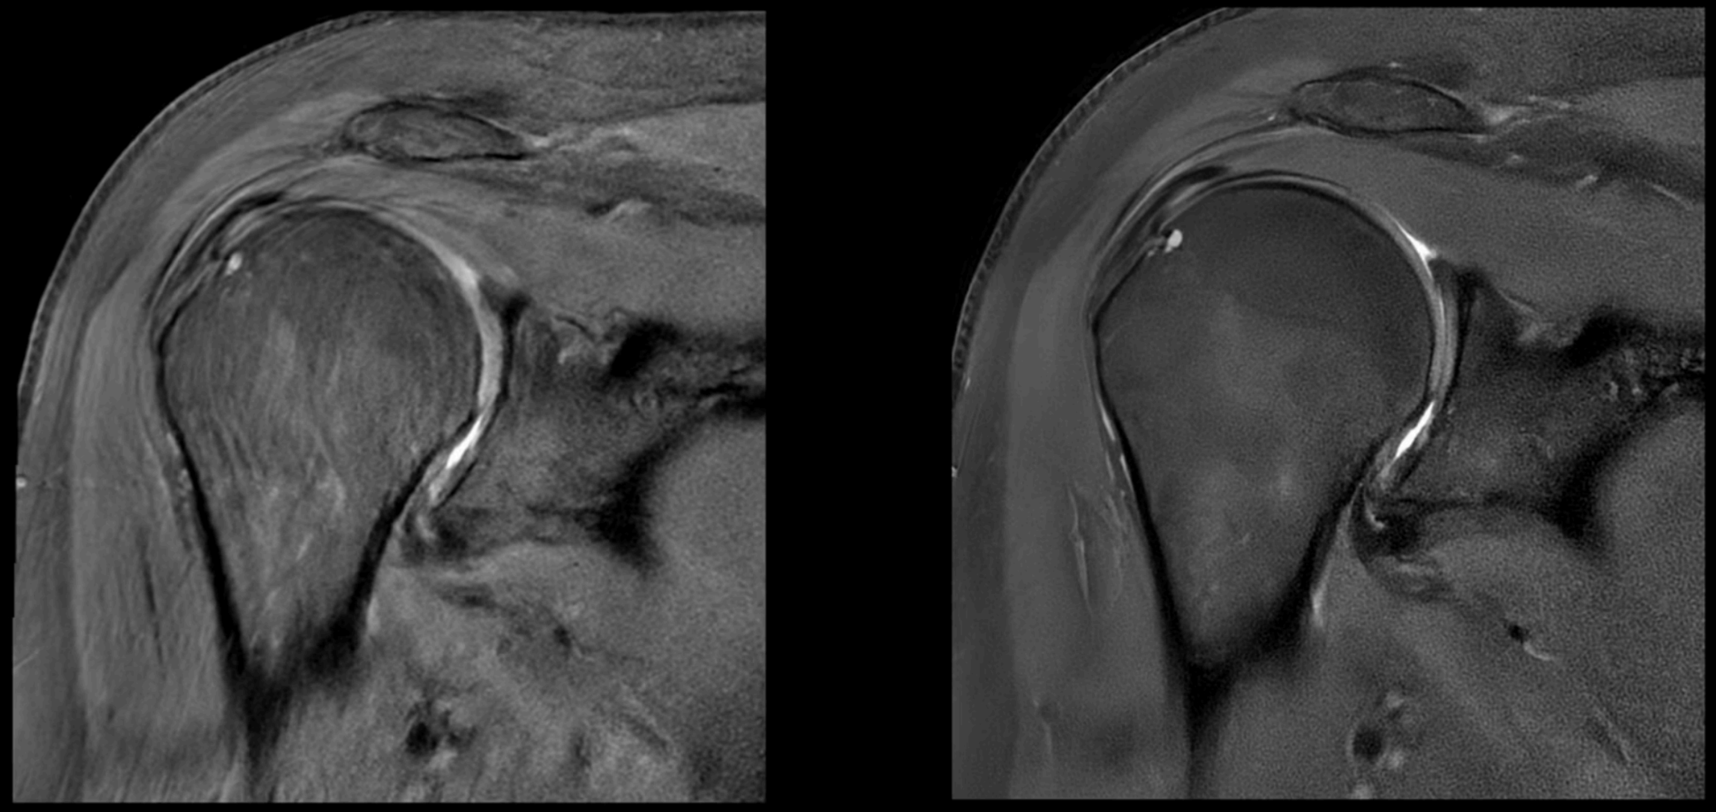

Left: Conventional
Coronal PD FatSat FSE,
0.3 x 0.4 x 3 mm, 2:13 min.

Right: AIR™ Recon DL
Coronal PD FatSat PROPELLER
0.3 x 0.3 x 3 mm, 2:57 min.

From: https://www.siemens-healthineers.com/magnetic-resonance-imaging/technologies-and-innovations/deep-resolve

From: https://www.gehealthcare.com/products/magnetic-resonance-imaging/air-recon-dl

Publicly available MRI k-space datasets

- fastMRI (https://github.com/facebookresearch/fastMRI)
    - Knee, brain and prostate MRI
- SKM-TEA (https://github.com/StanfordMIMI/skm-tea)
    - Quantitative knee MRI with tissue segmentation
- M4Raw (https://github.com/mylyu/M4Raw)
    - Multi-contrast multi-repetition 0.3 T brain MRI
- CMRxRecon (https://github.com/CmrxRecon/CMRxRecon-SciData)
    - Cardiac Cine MRI and cardiac quantitative MRI

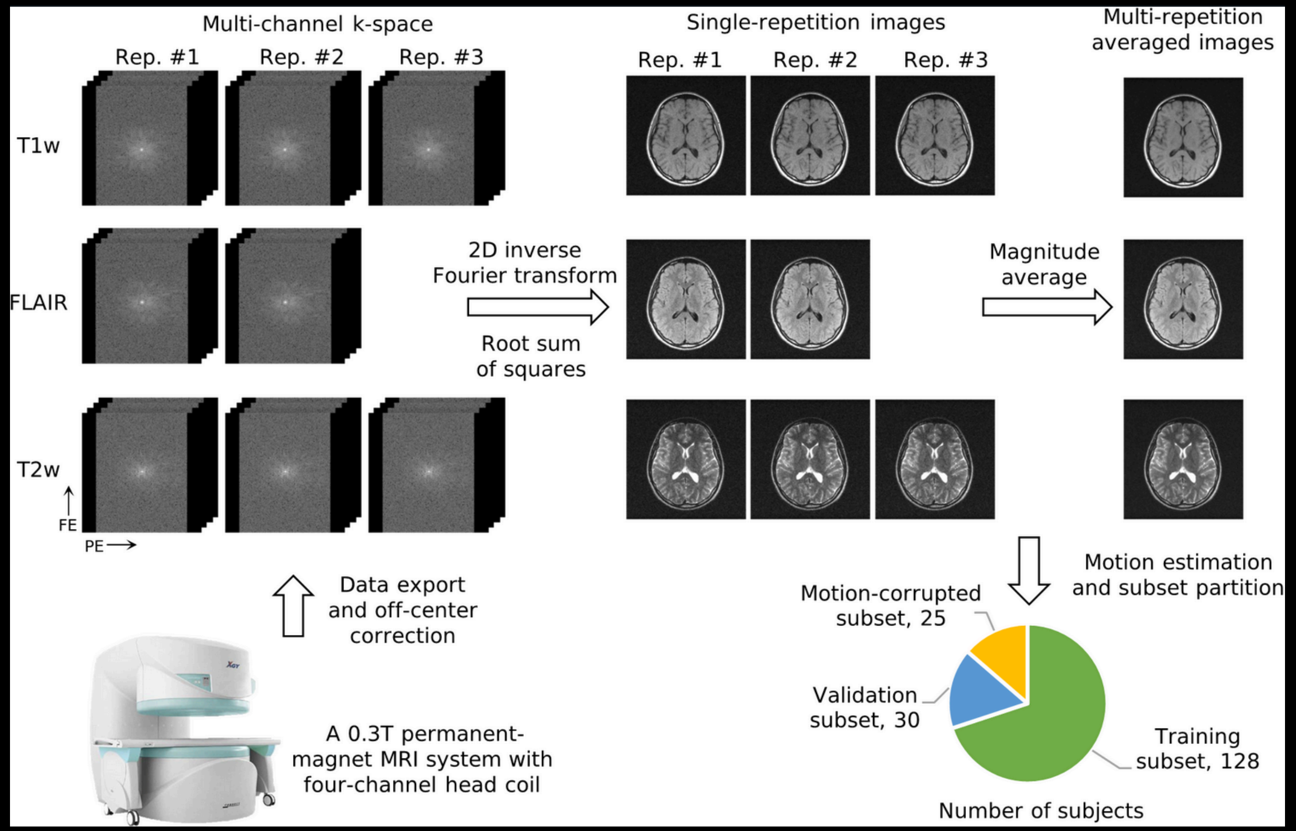

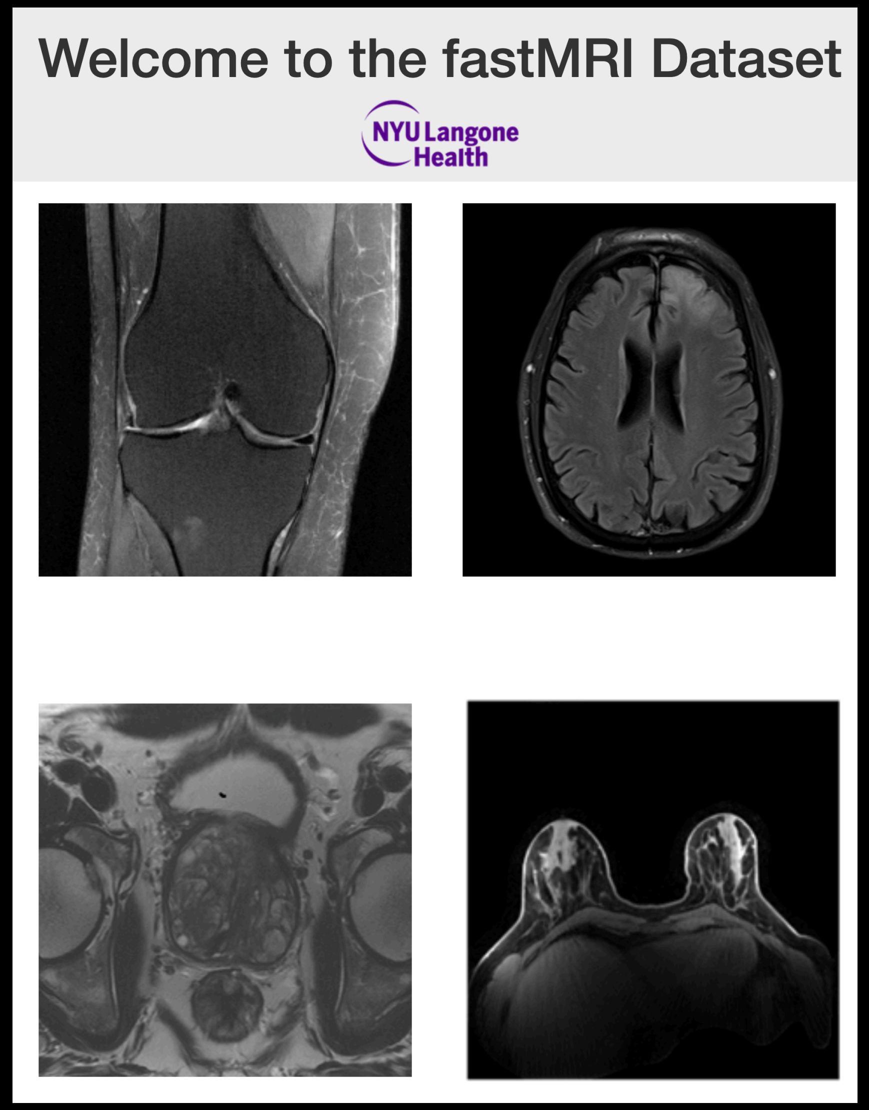

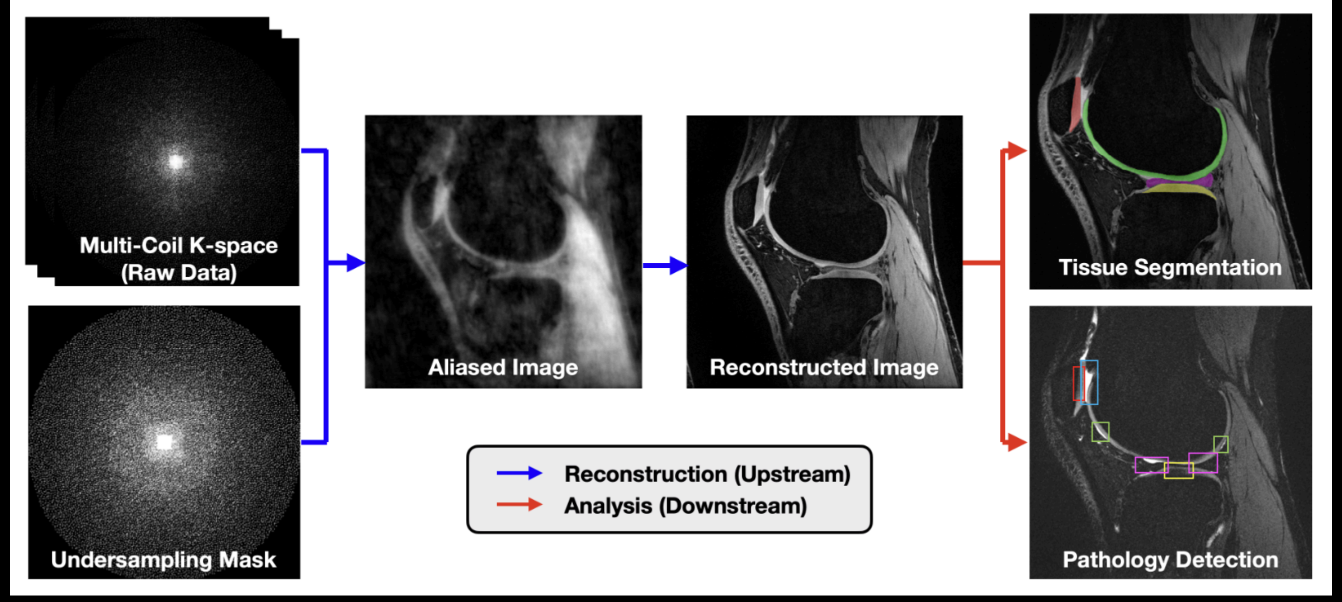

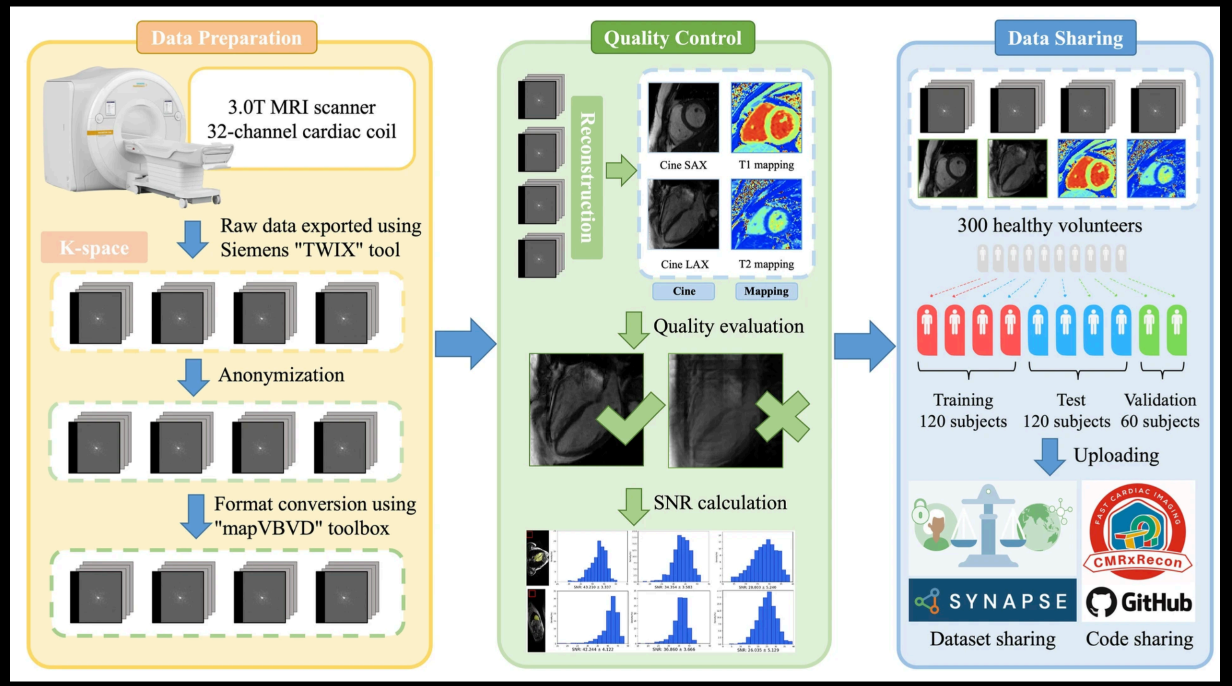

Discussion

- There are many opportunities, but there are also many open questions.
- What are the limitations for deep learning-based MRI reconstruction?
  - Let’s ask ChatGPT…

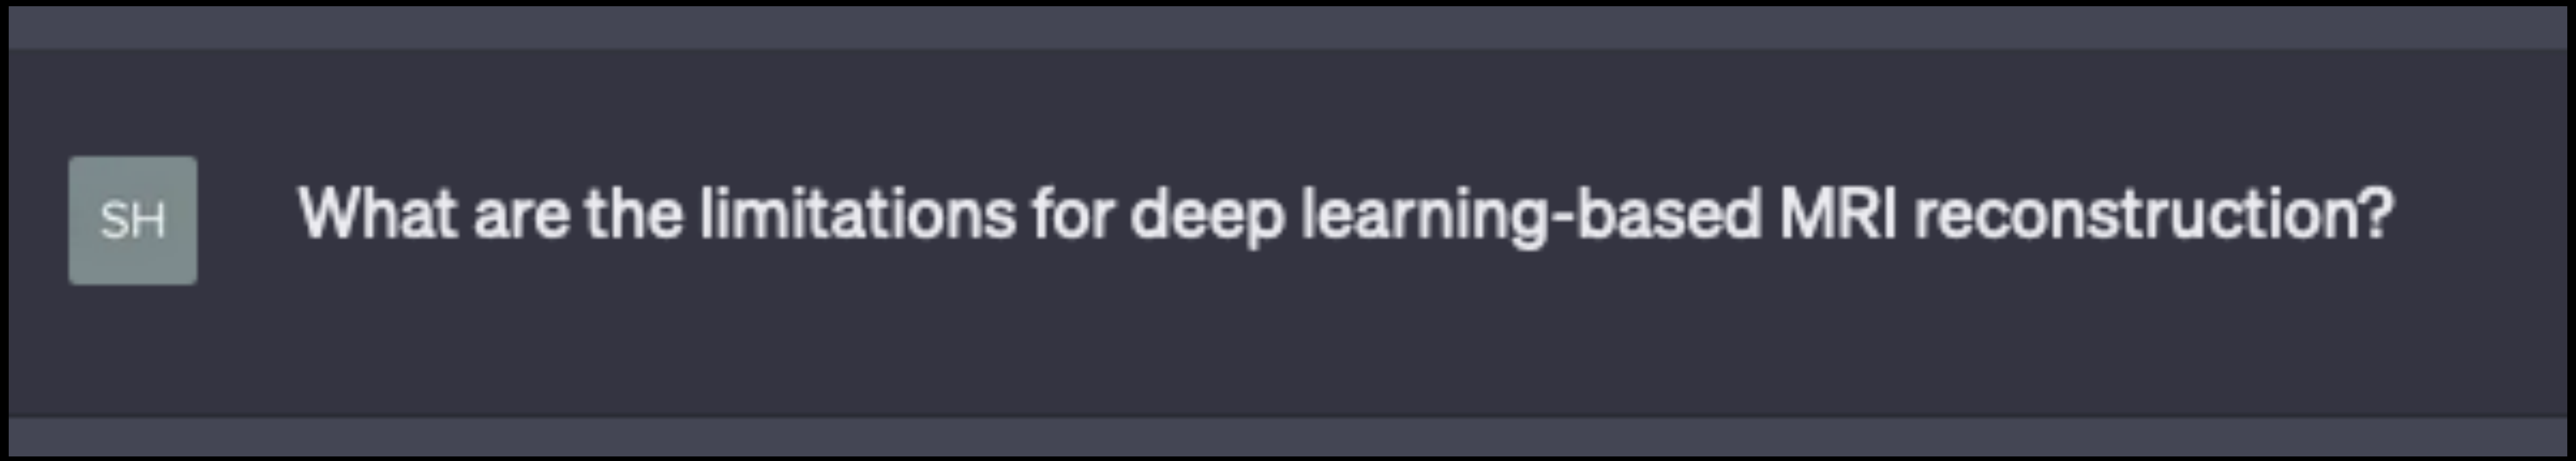

Discussion

- Limitations of deep learning-based MRI reconstruction
  - Insufficient training data
    - Even though there are public large datasets, obtaining diverse and representative dataset is still challenging.
  - Lack of interpretability / “Failure mode” not clear
    - The black-box nature of deep learning can be problematic for clinical acceptance and trust.
    - Uncertainty quantification or theory to explain deep learning are being investigated
  - Generalization to different acquisition parameters
    - Potential solution would be including large datasets with all different acquisition parameters or including sequence parameters as inputs
  - Computational complexity
    - The hardware keeps advancing and it can still be expensive

A few *personal* suggestions…

- Focus on the problem you want to solve (*to improve image quality? to allow for higher undersampling factors? to train without ground truth images?…*).

- Have a good understanding on the deep learning tools you have. Choose or develop methods or architectures that can solve the problem.

- Understand your data and be aware of the MRI signal model and acquisition process. There can be constraints or there can be some prior information to utilize.

- Don’t get lost in numbers! Don’t forget the clinical problem.

Thanks!

- To provide feedback for the lectures:

Questions?

Contact: Shu-Fu Shih

Email: sshih@mednet.ucla.edu