

facebook
Artificial Intelligence Research

fastMRI

Tullie Murrell
Applied Research Scientist @ FAIR

Outline

- Intro / how do MRIs work?
- The fastMRI dataset
- Traditional MRI Reconstruction
- Deep Learning for MRI Reconstruction
- Future MRI opportunities

A bit about me

- Studied at The University of Adelaide in Australia
- Joined Facebook in 2016
- Worked in Applied Machine Learning (AML) group focusing on real-time inference Computer Vision
- Worked in FB - Building 8 as Computer Vision lead
- Currently working in Facebook AI Research (FAIR) on the FastMRI project

Facebook AI Research (FAIR)

- Advancing the state-of-the-art in artificial intelligence through fundamental and applied research in open collaboration with the community.


FAIR - fastMRI

- Accelerate the acquisition of MRIs
- Aim to achieve up to 10x reduction in scan times
- Collaboration with FAIR and NYU
- Provide open source code
- Change clinical practice

Why accelerate MRI?

- MR scanning is slow: 20 minutes ~ 1 hour
- More people could access MRI
 - Patients too sick to stay still or children
- Potentially replace other, more-harmful imaging types (e.g. X-rays)
- Image moving organs (heart, lung, abdomen)


How do MRIs operate?



- MRI scanners operate in the frequency domain (called k-space)
- K-space is acquired one line at a time
- Scan time is proportional to the number of measurements


Parallel Imaging

- Most modern clinical MRI machines use parallel imaging (PI)
- PI uses multiple receiver coils
- Coils are more sensitive to the region of the object closest to them
 - The sensitivities depend on the MRI machine as well as the object being imaged
- The individual coil images are combined in software



Compressed Sensing

Acquire a subset of the k-space data and reconstruct the image from this subset



Under Sampled K-Space

IFT




Aliased
Image

Reconstruction
Model



Reconstructed Image

Why is this hard?





The fastMRI Dataset





Dataset details

- We released the first version of data + baseline models + leaderboard in Nov.
- Announced at the RSNA radiology conference.
- 1.6K 3D Knee MRIs (57K images) + 1.2M additional pre-processed images
- Public challenge run on dataset starting June 7th



Dataset details

- Multi-coil and Single-coil images
- 5 MRI machine types
- 2 different acquisition types:
 - Fat Suppressed
 - Non-fat suppressed



Publicly available MRI datasets

Dataset	Year	Number of Volumes
NYU Dataset	2018	100
Stanford 2D FSE	2017	89
Stanford 3D FSE	2017	20
Stanford Undersampled Dataset	2017	38
fastMRI (Ours)	2018	1594

Dataset privacy


- All images anonymized (i.e. patient records for each image removed)
- Unable to identify people from knee anatomy (unlike brain images)



Traditional MRI Reconstruction




MRI Reconstruction – kspace vs image



How to undersample images?

- Number of undersampled measurements proportional to scan time acceleration
- Generally low frequency k-space lines are kept
- Sampling patterns:
 - Cartesian
 - Radial
- Random vs equispaced




Traditional Parallel Imaging Algorithms

- When not accelerated, use closed form Root Sum of Squares (RSS)
- Two main approaches on undersampled data:
 - SENSE – reconstruct in image space
 - GRAPPA – reconstruct in K-space


Parallel Imaging Algorithms - SENSE

- Equispaced k-space sampling causes wrap-around effect in images (aliasing)
- SENSE solves simultaneous equations based on where image pixels wrap around and estimated coil sensitivities
- Sensitivities are estimated using approximate reconstruction from low frequency k-space lines



Parallel Imaging Algorithms - GRAPPA

- Idea: Unobserved k-space values can be expressed as a convolution of the observed values from all coils
- The convolution weights need to be estimated independently for each image




Compressed Sensing Algorithms

- Traditional compressed sensing - adds sparsity constraint to underdetermined system created from undersampled measurements
- This can be uniquely solved with L1 regularization (LASSO)

$$\hat{x} = \underbrace{\|Ax - y\|_2}_{\text{Data consistency term}} + \underbrace{\|\phi(x)\|_1}_{\text{Sparsity term}}$$

Compressed Sensing – sparsity constraints

Total variation



Wavelet




Challenge

- Traditional parallel imaging provides 2x acceleration
- Sparsity compressed sensing provides another 2x acceleration
- How do we extend this to higher accelerations?
 - Combine traditional methods with Deep Learning!



Deep Learning for MRI Reconstruction


MRI Reconstruction – kspace vs image



Model architectures


- Common architectures (taken from Computer Vision literature):
 - U-Net
 - Full resolution residual network (FRRN)
- SOTA: Learning a Variational Network for Reconstruction of Accelerated MRI Data (Hammernik et al.)

U-Net



Traditionally built for image segmentation but works as a good baseline for image reconstruction.


Full resolution residual network (FRRN)




Slight improvement over U-Net by learning high resolution short-cuts.

How to exploit uniqueness of reconstruction problem?

Variational Block



Variational Network (Hammernik et al.)



Variational Network (Hammernik et al.)

- Data consistency network layers (soft consistency)
- Relies on sensitivity estimation (e.g. from SENSE)
- Can be interpreted as rolling out each optimization step of traditional compressed sensing

Model size

- Key problem – accuracy keeps improving with larger models
- Solutions:
 - Half precision training
 - Remove batch-normalization layers
 - Architecture search

Losses

- L2 (Mean squared error)
- L1
- SSIM – Structural Similarity Index
- PSNR – Peak Signal to Noise Ratio
- Perceptual loss

Structured Similarity Index (SSIM)


- Intuition: evaluates luminance, contrast and structure by comparing reconstruction and ground-truth mean and variance.

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)}$$

How to make sure reconstructions are diagnosable?

- Current:
 - Weekly radiologist review
 - Clinical studies
 - **Not reproducible or scalable**
- Future:
 - Strong perceptual losses


Perceptual loss



- Can we incorporate classification information into a learned loss?

Active sampling pattern

- Not all sampling patterns are created equal
- Question: Can we predict the best sampling pattern?
 - Iteratively identify the best sample to acquire next.
- Leads to $> 50\%$ reduction in MSE compared to random subsampling at 4x acceleration



Training environment

- Pytorch
- Distributed training
- Hyperparameter search
- Reproducibility
- Visualization - Tensorboard






Future MRI opportunities



Future fastMRI opportunities

- Different body types
- Dynamic MRI
- Accelerating other image modalities
- Volumetric reconstruction



Beyond MRI reconstruction

- Segmentation
- Classification
 - Abnormality detection
 - Stanford, MRNet dataset
 - Classification directly from under-sampled data



Thank you!