


Unified Deep Learning Framework for MRI Image Contrast Translation

Rachel Yang
Scarsdale High School

May 29, 2026

New England Statistics Symposium (NESS) 2026


Introduction of MRI and Image Contrasts



- MRI: magnetic resonance imaging
- Important diagnostic imaging tool
- **Multiple image contrasts**
 - T1, T2, FLAIR, diffusion, perfusion, etc.
- Each contrast makes certain structures more or less visible

Why Do We Need Multiple MRI Contrasts?

- Each contrast makes certain structures visible or prominent
- They provide complementary information




- Radiologists rely on multiple MRI contrasts for accurate diagnosis

T1 vs T2 Contrasts

T1-Weighted MRI


- **Best for:** anatomy, fat tissue
- **Fat :** bright
- **Fluid :** dark



T1 Image

T2-Weighted MRI

- **Best for:** inflammation, fluid, lesions
- **Fat :** less bright
- **Fluid :** bright (fluid-sensitive)




T2 Image

Both contrasts together give doctors a more complete picture — neither alone is sufficient.

How are MRI Contrasts Acquired?

- To get multiple MRI contrasts, a series of MRI sequences are performed to acquire contrast-specific data sequentially



Total scan time: $3 + 4 + 5 = 12$ minutes (for 3 contrasts)

The Challenges to Acquire Multiple MRI Contrasts

Scan Time

Multiple contrasts = 30 to 60+ min per session



Patient Burden

Exhausting for ill or anxious patients



Cost & Access

Unavailable in under-resourced settings




What if we could generate the missing MRI contrast from the acquired data?


Inspiration from Language Translation




Language
Translate



MRI Contrast
Translate



MRI Contrast Translation: The Solution




Why this matters:

- ✓ Shorter scan sessions
- ✓ Lower cost per patient
- ✓ Helps when a contrast is corrupted, missing, or inaccessible
- ✓ Enables retrospective analysis on incomplete datasets
- ✓ Only need to acquire T1 (or T2)

Using a Deep Learning Approach

Deep Learning



- Deep learning is powerful for medical image processing
- Deep learning learns directly from real MRI data
- Once trained, synthesis is instant


Data Preprocessing

Registration


- Aligns T1 and T2 images so pixels correspond to the same location
- Essential for accurate training

Skull Stripping

- Removes non-brain tissue (skull, skin)
- Helps the model focus on what matters



Before



After

Source: ResearchGate, Fatima et al.

Dataset

Dataset	PMC Paired Multi-Contrast MRI (public, AAAI 2024)
Patients	40 total — 30 training, 10 testing
Image type	Paired T1 and T2 scans per patient
Slice extraction	5 central 2D slices per 3D volume (most information-dense)
Input size	256 × 256 pixels
Normalization	Pixel values scaled to $[-1, 1]$ range
Train/val split	80% training, 20% validation
Batch size	16 slices per training step

Source: Bagheri & Uludag, "PMC: paired multi-contrast MRI dataset," ML4CMH @ AAAI 2024.

Training Setup

Optimizer

Adam

Learning Rate

1×10^{-4} (adaptive)

Max Epochs

100

Early Stopping

Patience = 20 epochs

LR Schedule

Reduced on validation loss plateau


Why these choices?

Adam adapts the learning rate automatically — well-suited for image synthesis tasks.

Early stopping prevents overfitting: training halts when validation loss stops improving.

Research Objective: Optimizing Loss Functions

- *Which combination of loss functions produces the highest quality MRI contrast translation?*
- We tested 5 combinations, from a simple L1 baseline to a full composite loss
- Evaluated on T1-T2 and T2-T1 translation



Loss Functions: Teaching the Model What 'Good' Looks Like

A loss function measures how wrong the prediction is — the model learns by minimizing it.

L1 Loss

Pixel-by-pixel average error. Simple but can produce blurry images.

+ Perceptual Loss

Compares deep features (not just pixels). Preserves texture and realistic appearance.

+ Gradient Loss

Penalizes blurry edges by comparing image gradients. Sharpens boundaries.

+ Frequency Loss

Works in frequency domain (FFT). Captures fine-grained detail that spatial losses miss.

Experimental Design: 5 Loss Combinations

We systematically tested which combination of loss functions produces the best MRI synthesis:

1	L1 only (1.0)	<i>Baseline</i>
2	L1 + Perceptual (0.4)	
3	L1 + Gradient (0.6)	
4	L1 + Frequency (0.3)	
5	L1 + Perceptual + Gradient + Frequency	✓ <i>Final model</i>

Weights were tuned by testing 0.2, 0.4, 0.6, 0.8 for each component — best values selected per metric.

How Do We Measure Image Quality?

SSIM

SSIM *Structural Similarity Index Measure*

Measures structural similarity between predicted and real image. Range: 0–1. Higher is better.

PSNR *Peak Signal-to-Noise Ratio*

Image fidelity in decibels. Higher = less noise/error. Common in image processing research.

MAE *Mean Absolute Error*


Average pixel-level prediction error. Lower is better. Intuitive and easy to interpret.

MSE *Mean Squared Error*

Like MAE but squares the errors — penalizes large mistakes more heavily. Lower is better.

T1 to T2 Visual Results


Does adding more loss terms visually improve the synthetic T2 image?



Key observation: The composite loss produces the sharpest reconstruction with the best-preserved fine structural details.

T2 to T1 Visual Results

The model also works in reverse — generating synthetic T1 images from T2 input.



- Differences between models are more subtle from T2 to T1
- Performance was already very high across all loss variants.

Quantitative Results: All Five Models

Averaged over 10 test patients. ↑ = higher is better, ↓ = lower is better.

Loss Function	T1→T2 SSIM ↑	PSNR ↑	MAE ↓	T2→T1 SSIM ↑	PSNR ↑	MAE ↓
L1 1.0	0.9200	24.85	0.0036	0.9828	24.63	0.0037
L1 + Perceptual 0.4	0.9319	25.39	0.0032	0.9847	25.06	0.0033
L1 + Gradient 0.6	0.9357	25.50	0.0031	0.9867	25.68	0.0030
L1 + Frequency 0.3	0.9480	26.31	0.0025	0.9917	27.54	0.0018
Composite (final)	0.9485	26.34	0.0025	0.9913	27.35	0.0019

Key Finding: Why Did We Choose the Composite Loss?

The Nuance:

- L1 + Frequency slightly outperformed Composite on T2→T1 (SSIM 0.9917 vs 0.9913)
- Difference was marginal —performance has saturated from T2 to T1.

Our reasoning:

1. **Balance across both tasks:** Composite loss improved T1→T2 more significantly, making it the stronger overall model.
2. **Consistency across metrics:** composite loss performed near- best across all three metrics in both directions.
3. **Diminishing returns:** The T2→T1 gap (0.0004 SSIM) is below the threshold of clinical significance.

What Could This Mean Clinically?

Our results demonstrate strong synthetic image quality — here's what that could enable:



Reduced Scan Time

Acquire one contrast, synthesize the other — potentially cutting MRI session length significantly.



Pediatric or Claustrophobic Patients

Some patients can't tolerate long scans, so synthesis of images can eliminate the time constraint for them.



Retrospective Datasets

Old clinical datasets often have missing sequences. Synthesis enables research on otherwise incomplete data.

Limitations & Future Work

Current Limitations

- Small dataset (only 40 patients)
- 2D slices only — 3D network may perform better
- Tested on one public dataset — may not generalize to all scanners/hospitals
- No clinical validation by radiologists yet

Future Directions

- Train on larger datasets
- Extend to other MRI contrasts (FLAIR, DWI)
- Explore 3D U-Net for full volumetric translation
- Work toward radiologist evaluation of synthetic images
- Clinical deployment — integrating the model into actual hospital workflows

Conclusion

01

A U-Net-based model can effectively perform bidirectional cross-contrast MRI translation ($T1 \leftrightarrow T2$).

02

Combining multiple loss functions (L1 + perceptual + gradient + frequency) outperforms any single loss function.

03

The composite loss enhances edge definition and preserves fine structural details compared to the L1 baseline.

04

Careful loss function design is critical for achieving high-quality medical image synthesis.

Thank You

Questions & Discussion

Acknowledgments

Mentor: Can Wu, Memorial Sloan Kettering Cancer Center

Dataset: Bagheri & Uludag, PMC Dataset (AAAI 2024)

Architecture: Ronneberger et al., U-Net (MICCAI 2015)

Review Reference: Acs & Zhuang (Discover Imaging 2025)

Rachel Yang | Scarsdale High School | 2026