The University of Texas at Austin
Chandra Department of Electrical and Computer Engineering
Cockrell School of Engineering

ODEN INSTITUTE FOR COMPUTATIONAL ENGINEERING & SCIENCES

Generalized system identification and joint signal reconstruction with implicit neural representations, with application to MRI

Zach Stoebner & Prof. Jon Tamir
Computational Sensing & Imaging Lab (ECE)

Problem Setup

Imaging systems have spatiotemporal dynamics. The inverse problem can be defined as:

$$\boldsymbol{y}(\boldsymbol{r}, t) = \boldsymbol{F}\boldsymbol{H}_t\boldsymbol{x}(\boldsymbol{r}) + \boldsymbol{n}(\boldsymbol{r}, t)$$

where $[\boldsymbol{r}, t] \in \mathbb{R}^{d+1}$; $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{n} \in \mathbb{C}^{m_1 \times m_2 \times \ldots \times m_d}$; $\boldsymbol{n} \sim N(0, \sigma^2)$. $\boldsymbol{F}$ and $\boldsymbol{H}_t$ are the measurement and system effect operators, respectively.

Conventional methods rely on idealized physical models, but real systems have system-specific effects with unknown models, e.g., magnetic hysteresis, eddy currents, thermal drift.

- Can we identify unknown system effects $\boldsymbol{H}_t$ from a few measurements across time without relying on idealized physical models?
- Can we jointly reconstruct the image $\boldsymbol{x}$?
- Can we accomplish these tasks while only optimizing the parameters θ of a neural network?

Variable Projection

Assume $\boldsymbol{F}$ is full-rank and in a high SNR regime. Approximate $\widehat{\boldsymbol{m}}(\boldsymbol{r}, t) = \boldsymbol{F}^{-1}\boldsymbol{y}(\boldsymbol{r}, t) = \boldsymbol{H}_t\boldsymbol{x}(\boldsymbol{r})$ from k samples across time.

Parameterize $\boldsymbol{H}_\theta(\boldsymbol{r}) = [\boldsymbol{h}_\theta(\boldsymbol{r}, t_i)]_{i=1}^{k}$ where $\boldsymbol{h}_\theta(\boldsymbol{r}, t)$: $\mathbb{R}^{d+1} \to \mathbb{C}^{m_1 \times m_2 \times \ldots \times m_d}$.

Let $\widehat{\boldsymbol{m}}(\boldsymbol{r}) = [\widehat{\boldsymbol{m}}(\boldsymbol{r}, t_i)]_{i=1}^{k}$. We wish to solve:

$$\min_{\boldsymbol{x}, \theta} ||\widehat{\boldsymbol{m}}(\boldsymbol{r}) - \boldsymbol{H}_\theta(\boldsymbol{r})\boldsymbol{x}(\boldsymbol{r})|| \equiv \min_{\theta} ||\widehat{\boldsymbol{m}}(\boldsymbol{r}) - \boldsymbol{H}_\theta(\boldsymbol{r}) \frac{\widehat{\boldsymbol{m}}(\boldsymbol{r})^H \boldsymbol{H}_\theta(\boldsymbol{r})}{\boldsymbol{H}_\theta(\boldsymbol{r})^H \boldsymbol{H}_\theta(\boldsymbol{r})}||$$

where $\widehat{\boldsymbol{x}}_\theta(\boldsymbol{r}) = \frac{\widehat{\boldsymbol{m}}(\boldsymbol{r})^H \boldsymbol{H}_\theta(\boldsymbol{r})}{\boldsymbol{H}_\theta(\boldsymbol{r})^H \boldsymbol{H}_\theta(\boldsymbol{r})}$ is the coordinate-wise variable projection [3] of $\boldsymbol{x}(\boldsymbol{r})$ onto $\mathrm{col}(\boldsymbol{H}_\theta)$.

Define the cost function:

$$f(\boldsymbol{h}_\theta; \boldsymbol{r}, \widehat{\boldsymbol{m}}) = \frac{1}{2}\|\widehat{\boldsymbol{m}}(\boldsymbol{r}) - \boldsymbol{H}_\theta(\boldsymbol{r})\widehat{\boldsymbol{x}}_\theta(\boldsymbol{r})\|^2 \text{ where } \widehat{\boldsymbol{x}}_\theta \in \mathcal{X}, \boldsymbol{h}_\theta \in \mathcal{H}$$

MRI Signal & Objective

The spatiotemporal MRI signal equation [4] is:

$$\boldsymbol{y}(\boldsymbol{k}, t) = \int \boldsymbol{x}(\boldsymbol{r})\, e^{\frac{-t}{T_2^*(\boldsymbol{r})}} e^{-j2\pi\boldsymbol{\omega}(\boldsymbol{r})\cdot t} e^{-j2\pi\boldsymbol{k}(t)\cdot\boldsymbol{r}} d\boldsymbol{r} + \boldsymbol{n}(\boldsymbol{r}, t)$$

where $\boldsymbol{\omega}$ is off-resonance and $\boldsymbol{k}$ is the k-space trajectory. With the assumptions, we have:

$$\widehat{\boldsymbol{m}}(\boldsymbol{r}, t) = \boldsymbol{x}(\boldsymbol{r}) e^{\frac{-t}{T_2^*(\boldsymbol{r})}} e^{-j2\pi\boldsymbol{\omega}(\boldsymbol{r})\cdot t}$$

This implies that $\widehat{\boldsymbol{m}}(\boldsymbol{r}, 0) = \boldsymbol{x}(\boldsymbol{r})$. However, these ideal models may not be exactly reflected in the real measurements, implying $\widehat{\boldsymbol{m}}(\boldsymbol{r}, 0) = \boldsymbol{H}_0\boldsymbol{x}(\boldsymbol{r})$ where $\boldsymbol{H}_0$ is the initial condition.

Define the objective:

$$f(\boldsymbol{h}_\theta; \boldsymbol{r}, \widehat{\boldsymbol{m}}) = \frac{1}{2}\Big(\underbrace{\|\widehat{\boldsymbol{m}}(\boldsymbol{r}) - \boldsymbol{H}_\theta(\boldsymbol{r})\widehat{\boldsymbol{x}}_\theta(\boldsymbol{r})\|^2}_{\text{Data consistency}} + \underbrace{\lambda_1\|\boldsymbol{h}_\theta(\boldsymbol{r}, 0)\|^2}_{\text{Time intercept}} + \underbrace{\lambda_2\|\min\{|\angle\boldsymbol{h}_\theta(\boldsymbol{r}, t) \bmod 2\pi - \pi|, 2\pi - |\angle\boldsymbol{h}_\theta(\boldsymbol{r}, t) \bmod 2\pi - \pi|\}\|^2}_{\text{Phase deviations from } \pi} + \underbrace{\lambda_3\|D_{\boldsymbol{r}}^2\{\boldsymbol{h}_\theta(\boldsymbol{r}, t)\}\|^2}_{\text{2nd-order finite differences over space}}\Big)$$

Implicit Neural Representations

An implicit neural representation is $\boldsymbol{h}_\theta(\boldsymbol{r}, t) = \boldsymbol{f}_\theta(\boldsymbol{\gamma}(\boldsymbol{r}, t))$. [1] Define the lifting operator of random Fourier features:

$$\boldsymbol{\gamma}(\boldsymbol{r}, t) = [\cos(2\pi\boldsymbol{B}[\boldsymbol{r}, t]^T), \sin(2\pi\boldsymbol{B}[\boldsymbol{r}, t]^T)]$$

where $\boldsymbol{B} \in \mathbb{R}^{l \times (d+1)}$ is sampled from $N(0, s^2)$. The embedding size l and the scale s are tunable. [2]

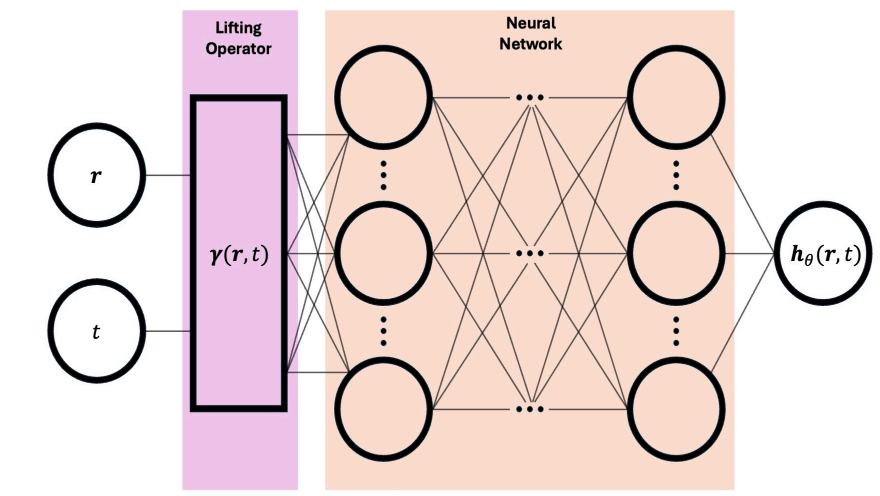

1D Simulation & 2D GRE MRI

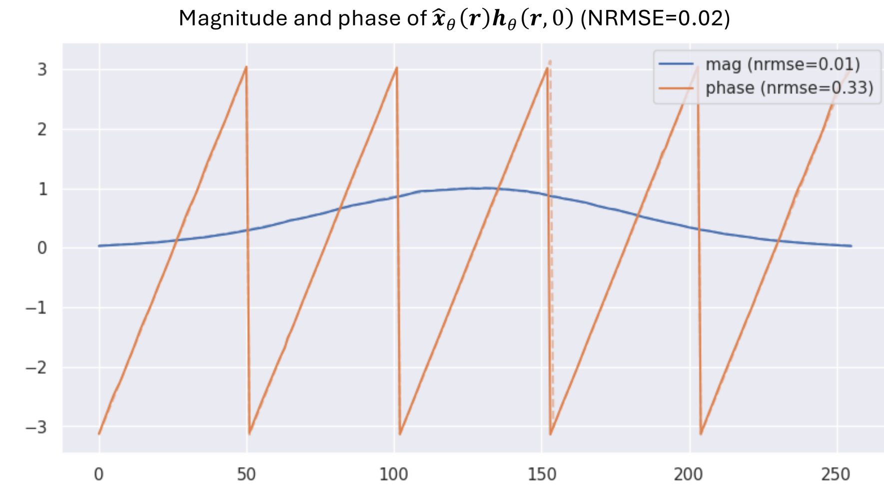

Magnitude and phase of $\widehat{\boldsymbol{x}}_\theta(\boldsymbol{r})\boldsymbol{h}_\theta(\boldsymbol{r}, 0)$ (NRMSE=0.02)

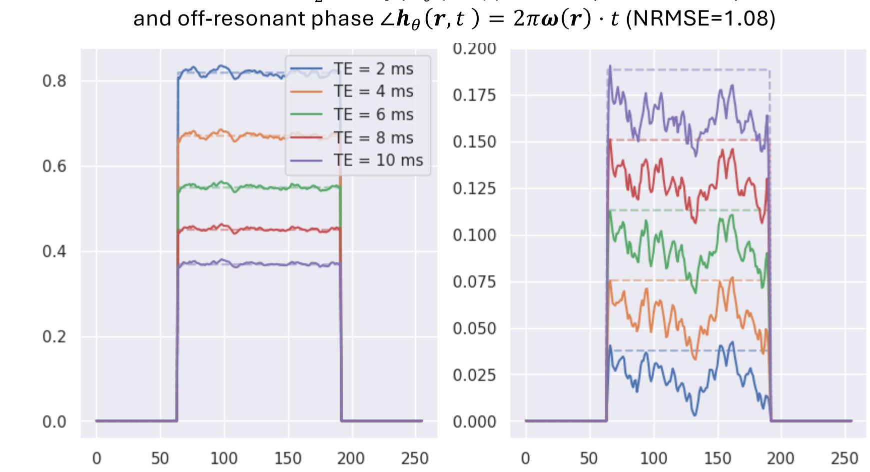

Estimates of T_2^* decay $|\boldsymbol{h}_\theta(\boldsymbol{r}, t)| = e^{\frac{-t}{T_2^*(\boldsymbol{r})}}$ (NRMSE=0.67) and off-resonant phase $\angle\boldsymbol{h}_\theta(\boldsymbol{r}, t) = 2\pi\boldsymbol{\omega}(\boldsymbol{r}) \cdot t$ (NRMSE=1.08)

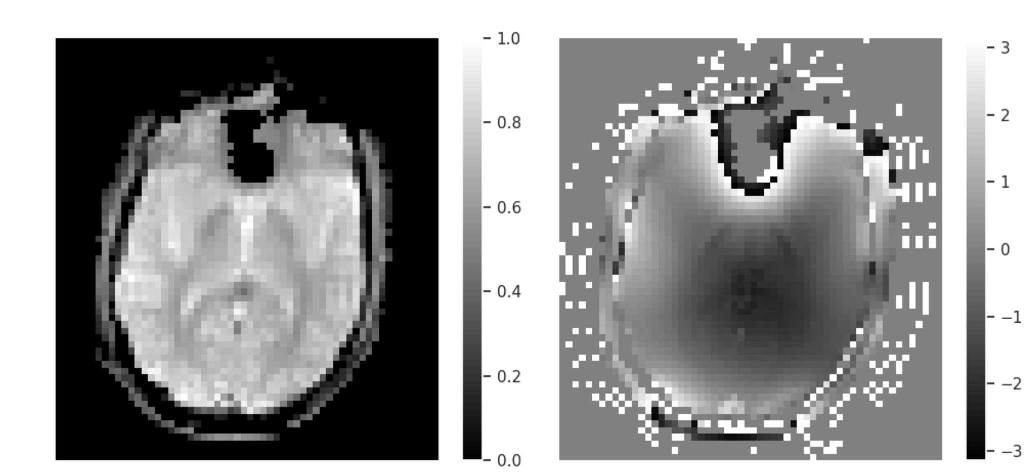
Magnitude and phase of ground truth $\boldsymbol{x}^*(\boldsymbol{r})$

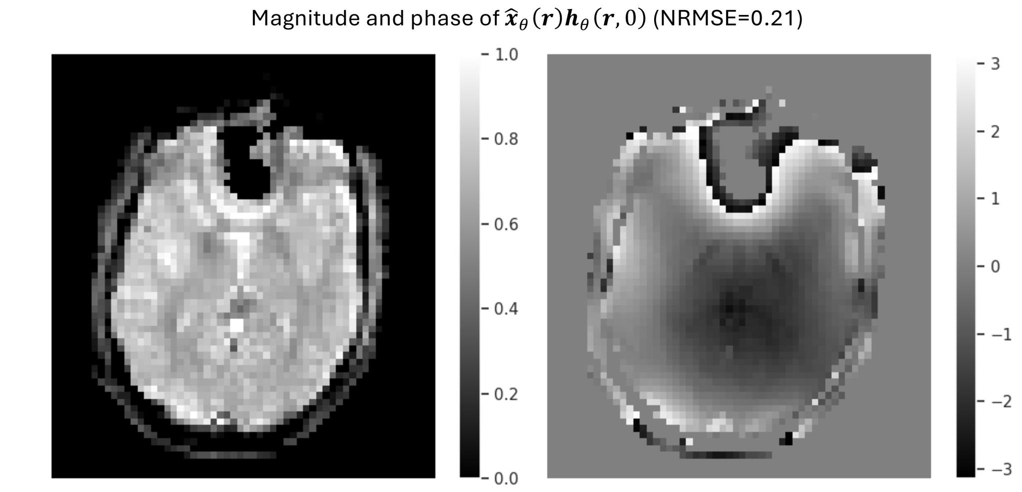
Magnitude and phase of $\widehat{\boldsymbol{x}}_\theta(\boldsymbol{r})\boldsymbol{h}_\theta(\boldsymbol{r}, 0)$ (NRMSE=0.21)

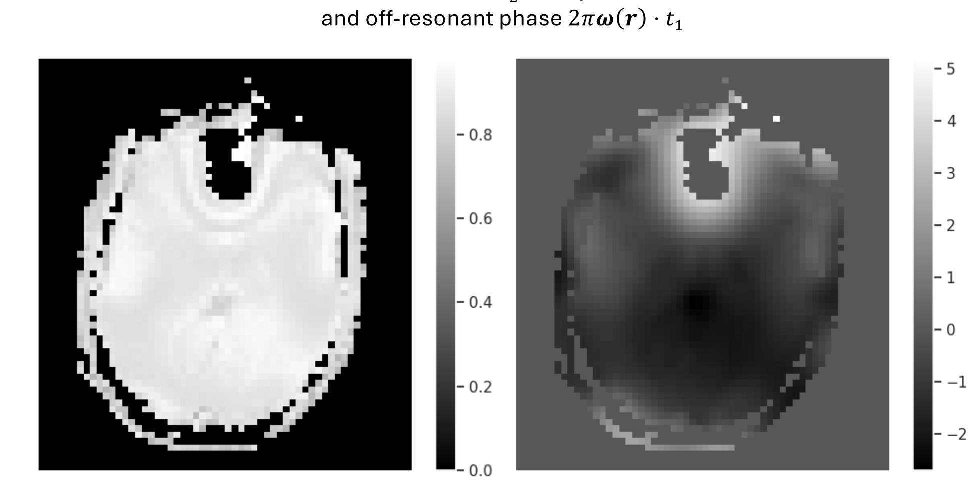
Ground truth T_2^* decay $e^{\frac{-t_1}{T_2^*(\boldsymbol{r})}}$ and off-resonant phase $2\pi\boldsymbol{\omega}(\boldsymbol{r}) \cdot t_1$

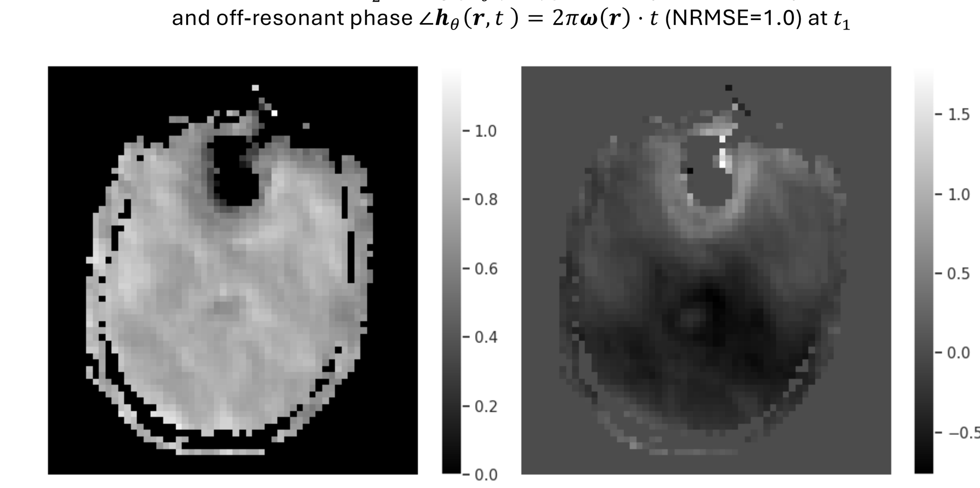
Estimates of T_2^* decay $|\boldsymbol{h}_\theta(\boldsymbol{r}, t)| = e^{\frac{-t}{T_2^*(\boldsymbol{r})}}$ (NRMSE=0.42) and off-resonant phase $\angle\boldsymbol{h}_\theta(\boldsymbol{r}, t) = 2\pi\boldsymbol{\omega}(\boldsymbol{r}) \cdot t$ (NRMSE=1.0) at t_1

Insights & Next Steps

- If the physical system has non-zero initial conditions, then we expect that $\boldsymbol{x}^*(\boldsymbol{r}) = \widehat{\boldsymbol{m}}(\boldsymbol{r}, 0) \approx \widehat{\boldsymbol{x}}_\theta(\boldsymbol{r})\boldsymbol{h}_\theta(\boldsymbol{r}, 0)$.
- For later sample times $t > 30\ ms$, $NRMSE(\widehat{\boldsymbol{x}}_\theta(\boldsymbol{r})\boldsymbol{h}_\theta(\boldsymbol{r}, t), \widehat{\boldsymbol{m}}(\boldsymbol{r}, t))$ is low but $\boldsymbol{h}_\theta(\boldsymbol{r}, t)$ and $\widehat{\boldsymbol{x}}_\theta(\boldsymbol{r})$ are incorrect due to phase-wrapping, nonlinear long-range dependency → phase unwrap before fitting, penalize large phase deviations, explore different architectures for spatiotemporal sequences.

References

[1] Sitzmann, Vincent, et al. "Implicit neural representations with periodic activation functions." *NeurIPS* (2020).
[2] Tancik, Matthew, et al. "Fourier features let networks learn high frequency functions in low dimensional domains." *NeurIPS* (2020).
[3] Golub, Gene, and Victor Pereyra. "Separable nonlinear least squares: the variable projection method and its applications." *Inverse problems* (2003).
[4] Funai, Amanda K., et al. "Regularized field map estimation in MRI." *IEEE-TMI* (2008).

Acknowledgements

This work is supported in part by IFML

inrlib