# PHYSICS-BASED DECODING IMPROVES MAGNETIC RESONANCE FINGERPRINTING

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Magnetic Resonance Fingerprinting (MRF) is a promising paradigm to perform fast quantitative Magnetic Resonance Imaging (QMRI). However, existing MRF methods suffer from slow imaging speeds and poor generalization performance on radio frequency pulse sequences generated in various scenarios. To address these issues, we propose a novel model-based MRF method that learns better representations by integrating a fast and differentiable MRI physics model as a form of causal decoding. The proposed approach adopts a supervised auto-encoder framework consisting of an encoder and a decoder, where the encoder predicts the target tissue properties (anti-causal mapping) and the decoder reconstructs the inputs (causal mapping). Specifically, the encoder embeds high-dimensional MRF time sequences into a low-dimensional tissue property space, while the decoder exploits an MRI physics model to reconstruct the input signals using the estimated tissue properties and associated MRI settings. The implicit causal regularization induced by the physics-based decoder improves the generalization performance and uniform stability by a considerable margin. Our experiments validate the effectiveness of the proposed physics-based decoding by achieving the state-of-the-art performance on tissue property estimation.

## 1 INTRODUCTION

Quantitative Magnetic Resonance Imaging (QMRI) is used to identify tissue's intrinsic properties, such as the spin-lattice magnetic relaxation time (T1), the spin–spin magnetic relaxation time (T2), and other physical properties. Compared to conventional weighted (qualitative) MRI that focuses on tissue's contrast of brightness and darkness, QMRI reveals tissue's intrinsic properties with quantitative values and associated physical interpretations. Since different tissues are characterized by their distinct properties values, QMRI shows great potential to reduce subjectivity, with advantages in many areas including diagnosis, tissue characterization, investigation of disease pathologies, etc. [2].

Magnetic Resonance Fingerprinting (MRF) provides an alternative QMRI framework to achieve multi-property quantification simultaneously [38]. Given a pseudo-random radio frequency (RF) pulse sequence, a distinct magnetic response – a.k.a. fingerprint, signature, or signal evolution – from each specific tissue is observed and then used to predict the target tissue properties. Therefore, multi-property quantification boils down to an inverse problem (*i.e.*, an anti-causal task) that aims to infer underlying tissue properties (causal factor) from the magnetic responses (effect factor).

Various approaches have been developed to solve the MRF problem, using model-based techniques, e.g. dictionary matching (DM), compressive sensing, as well as learning-based / data-driven techniques [5–7, 9, 14, 15, 17, 24, 34, 38, 41, 42, 44, 48–50]. In spite of good performance in particular situations, they rarely take the MRI dynamics into consideration. This can cause reduced robustness and generalizability to potential data shifts occurred in practical scenarios with serious negative consequences. For example, the T1 and T2 value range and distribution are patient-specific and subject to pathological tissue types, development phase and other factors, which may cause label shift. In addition, as specific RF settings can often be applied to different situations, hospitals, and MRI instruments, MRF models are naturally expected to be able to handle such varied cases and be generalized to different RF settings. Motivated by these issues, we aim to develop a new MRF approach that combines the benefits of both model-based and learning-based techniques to achieve superior robustness and generalizability.

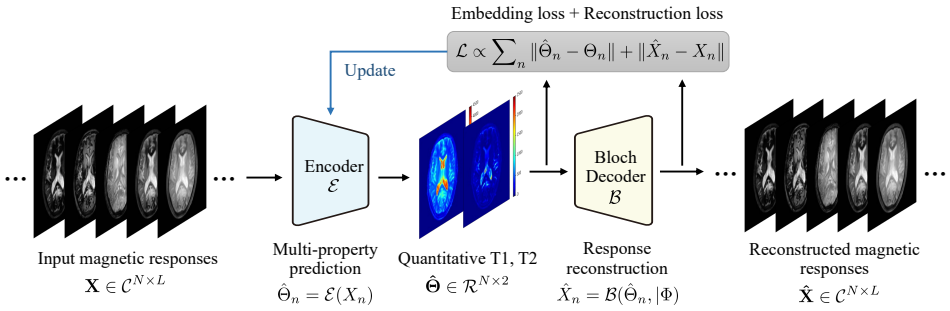


Figure 1: Diagram of the proposed BlochNet, a physics-based autoencoder for MRF. BlochNet adopts a supervised auto-encoder framework where the encoder solves an inverse problem that predicts tissue properties from input magnetic responses – an anti-causal task, while the decoder leverages an Bloch equation based MRI physics model to address a causal task that reconstructs the input responses from the estimated tissue properties. Such design helps the encoder capture anti-causal mapping effectively with the aid of causal feedback from the Bloch decoder.

In this work, we propose a physics-based auto-encoder model, called *BlochNet*, to learn generalizable representations for MRF, as shown in Fig. 1. BlochNet adopts a supervised auto-encoder framework where the encoder solves the inverse problem that predicts tissue properties from input magnetic responses – an anti-causal task, while the decoder leverages a Bloch equation based MRI physics model to address a causal task that reconstructs the input responses from the estimated tissue properties. The Bloch equations, a.k.a. equations of motion of nuclear magnetization, are a set of ordinary differential equations that model the magnetization dynamics and calculate the nuclear magnetization as a function of relaxation times T1 and T2 [3]. They enable the formulation of the magnetic response of a tissue with specific intrinsic properties T1, T2 under varied magnetic field (we provide more details in Section 3.0.3 and Appendix). The physics-based decoder acts as a causal regularization (without the need to perform causal discovery like [30]), to perform supervised reconstruction to the input magnetic responses from estimated tissue properties as well as the associated pseudo-random excitation pulse sequence, such as repetition time (TR), time of echo (TE), and radio frequency flip angle (FA) over time. The rationale underlying the design is that domain knowledge such as physics principles can help to reduce the solution space of an inverse problem by applying additional constraints. This contributes to finding a better solution for an (ill-posed) inverse problem and capturing an obscure anti-causal mapping [27, 32, 37, 40].

However, it is not feasible to naively apply the MRF physics model directly, as solving the Bloch equations is highly computationally intensive. Thus, we significantly improve the implementation efficiency of the Bloch equation based MRI physics model. To the best of our knowledge, we are the first to apply exact Bloch equation in a training procedure by explicitly using it as differential decode. Our major contributions include:

- We propose BlochNet, a supervised autoencoder framework for MRF, where the Bloch equation based MRI physics model serves as the decoder and guides the encoder to learn underlying robust representations by providing causal feedback of searching possible tissue properties.
- Compared to earlier methods, BlochNet shows consistently better generalization performance across synthetic, phantom and real MRF data, and across different types of RF pulse sequences. This demonstrates well the benefits of physics-based decoding in MRF in practical scenarios.
- We provide a fast and end-to-end solvable MRI physics model that can be used directly as a differentiable module in neural networks (*e.g.*, it acts as a decoder in BlochNet).

## 2 Related Works and Preliminaries

### 2.1 Causal Learning and Physics-Informed Learning

In causal learning, the model learns to predict the output that is generated by the causal factors, while anti-causal learning infers the causes from the given output [27]. Under the principle of

independent causal mechanisms (ICM), the causal mechanism $P(Y|X)$ is invariant under changes in the distribution of the causes $P(X)$, so that it can be applied to any causes with different distributions [45, 46]. On the other hand, the anti-causal direction $P(X|Y)$, in which we are interested in many practical tasks, is not invariant with respect to $P(Y)$, since $X$ causes $Y$ rather than the opposite direction. With this property, learning anti-causal direction is hard since identifying causes over a large search space is expensive as [18] says that "inverses of well-behaved functions indeed can have larger execution complexity". [27] suggests a method to exploit a causal model for an efficient search strategy in the anti-causal direction. The anti-causal model finds causal factors by repeating the process of predicting possible causes from input data and getting feedback from the causal model that checks whether it actually generates the input. This leads to training the anti-causal model with strong generalization with invariant features for making predictions. [1, 51] apply encoders and decoders for the anti-causal model and the causal model, respectively. Also, [27] argues that the observed data is generated from independent causes that are entangled through causal mechanisms. Thus, the anti-causal model should disentangle factors for identifying causes from the observed effect and [11, 31] disentangle factors using causal generative models of the data.

Another highly related line of research is physics-informed machine learning [26] where a physics-based prior is embedded in the learning process. One of the typical examples is physics-informed neural networks (PINNs) [47] which are usually trained with both a data fitting loss and a PDE loss. Combining a physics prior in designing neural networks was shown to be beneficial to a broad range of applications [4, 13, 25, 39, 43]. PINNs are able to generalize well with a limited number of training data. In a similar spirit, we aim to incorporate the physics model that describes the MRI process into our neural networks to improve its data efficiency and generalization.

### 2.2 Model-based and learning-based MRF approaches

As tissue properties result in magnetic responses through the MRI dynamics, quantifying tissue's properties via QMRI/MRF is an typical anti-causal task. The core idea of MRF is based on fact that for each specific tissue, a pseudo-random pulse sequence leads to an unique magnetic response (i.e. magnetizition along the temporal dimension) which can serve as an identifiable signal signature, analogue to a "fingerprint", for the corresponding tissue. Once the magnetic responses are obtained, estimation of tissue properties from responses reduces to a pattern recognition problem. In the original MRF work [38], this is addressed via dictionary matching (DM) which finds the best matching entry in a pre-computed dictionary for each inquiry magnetic response. Accordingly, the best matching dictionary entry leads to multiple tissue properties directly via a look-up-table (LUT) operation. More specifically, the pre-computed dictionary is composed of a number of magnetic responses for a variety of tissues characterized by the values of their intrinsic properties, such as T1, T2 relaxation times, etc. In this way, each dictionary entry is associated with a specific tissue and its properties. Thus, once the best matching entry (i.e. most correlated element with respect to the enquiry in terms of their inner product, a.k.a. $\ell_2$ distance) is found, it directly leads to multiple properties simultaneously through a LUT. However, high computation and storage burden makes DM-based approaches prohibitively time-consuming and memory-consuming when the number of types and values of tissue properties increases, because the size of the dictionary and lookup table increases exponentially accordingly. To alleviate such drawbacks, other model-based MRF approaches were proposed to improve the speed, incorporate additional useful priors, reduce the computational complexity [5, 9, 34, 41, 42, 50].

Model-based MRF methods tend to suffer from burdens of storing the (enormous) dictionary and computational overhead. To address these shortcomings, learning-based approaches have been proposed for fast MRF by eliminating/replacing the dictionary by a compact neural network [6, 7, 14, 15, 17, 24, 44, 48, 49]. In particular, motivated by the success of deep learning [19, 33] in a number of tasks [12, 16, 20, 21, 23, 28, 29], there is an emerging trend [6, 7, 14, 15, 17, 24, 44, 48, 49] that suggests to use a trained neural network as an alternative substitute for the MRF dictionary and LUT so that the time-consuming dictionary matching operation can be eliminated and replaced by an efficient inference through a trained network. In fact, a well designed and tuned neural network is capable of approximating arbitrarily complex functions, therefore should also be able to approximate the response-to-property mapping function. Although these learning-based MRF approaches demonstrated better performance in terms of speed and accuracy in comparison with model-based variants, they also exhibit some limitations, such as degraded robustness and generalizability to various data shifts and out-of-distribution (OOD) data samples. For example, these learning-based

MRF models tend to suffer from increased risk of overfitting, reduced robustness and generalizability to various data shifts and out-of-distribution (OOD) data samples. Moreover, most of them are designed empirically without taking into account the MRI physics underlying the imaging process. Even though [6] lightly touched the concept of incorporating the Bloch dynamics in an encoder-decoder framework, the decoder in their work, however, is actually a learned network rather than a solid physics model. Since the decoder network in [6] just simulates and approximates the Bloch equations, it is more like a standard decoding model than a physics-based causal decoding one.

### 2.3 MRF PROBLEM SETTING

**Generation model**. The data generation model of producing magnetic responses from tissue properties is based on the MRI physics model formulated by the Bloch equations [3]. Given the RF pulse sequence whose parameter setting $\Phi = \{FA, TR, TE\}$ consists of flip angles $FA \in \mathbb{C}^L$, repetition times $TR \in \mathbb{R}^L$ and echo times $TE \in \mathbb{R}^L$ across $L$ time points, the temporal signal evolution $X_n \in \mathbb{C}^L$ for each individual voxel $n$ is associated with tissue properties such as $\Theta_n = \{T1_n, T2_n\}$, through the Bloch differential equations $\mathcal{B}(\Theta_n) : \mathbb{R}^p \to \mathbb{C}^L, p = 2$.[1]

$$X_n = \mathcal{B}(\Theta_n|\Phi) \quad \forall n \in 1, \ldots, N$$

The Bloch equations represent a nonlinear mapping from per-voxel intrinsic tissue properties to the corresponding temporal signal evolution that records the magnetisation response of proton dipoles to dynamic excitations induced by the RF sequence. Tissues with different properties respond distinctively to RF excitations. QMRI/MRF rely on this principle to estimate quantitative tissue properties from the signal evolution. In our experiment setting, two properties including the longitudinal T1 and transverse T2 relaxation times, are simultaneously encoded for each voxel. This setting could be further extended to include other properties, e.g. proton density $\rho$, off-resonance frequencies, T2*, diffusion and perfusion.

**Inverse model**. Given the magnetic response $X_n \in \mathbb{C}^L$ for the $n$-th voxel, the inverse process aims to address an inverse problem that maps the response back to the corresponding tissue properties $\Theta_n$.

$$\Theta_n = g(X_n) \quad \forall n \in 1, \ldots, N$$

where $g$ denotes the inverse mapping function. Note that, estimation of tissue properties $\Theta$ from magnetic responses $X$ requires long enough sequences $L > p$ to create unique signal evolutions that distinguish different tissues. Hence, the magnetic responses live on a low-dimensional (nonlinear) sub-manifold of $\mathbb{C}^L$.

## 3 PROPOSED METHOD

#### 3.0.1 ANTI-CAUSAL REPRESENTATION LEARNING FOR MRF

Here, we present an anti-causal representation learning method to achieve fast and robust MRF. The proposed approach adopts a supervised auto-encoder framework where the encoder solves the anti-causal primary task that predicts tissue properties from input signatures, while the decoder solves the causal auxiliary task that reconstructs the inputs signatures from the estimated tissue properties based on a MRI physics model. We highlight that a sophisticated MRI physics model is tailored and exploited as the decoder which plays the role of causal regularization and help the encoder learn causal representations effectively. The rationale is based on the fact that causal representations that reflect the underlying causal mechanism tend to exhibit stronger generalization to out-domain distribution, and the domain-specific causal knowledge can help to reduce the solution space, thereby contributing to a better solution when solving anti-causal tasks [27, 32, 40].

[1] For the sake of illustration and simplicity, we focus on the most commonly used tissue properties T1 and T2 relaxation time. We perform normalization for each signal evolution to remove the impact of the proton density which plays the role of a scale factor in front of the Bloch equations.

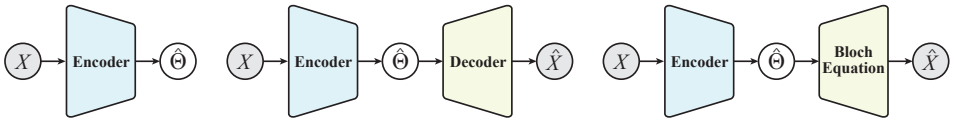


(a) Previsous method (b) Standard autoencoder (c) BlochNet: Physics-based autoencoder

Figure 2: Two baseline methods and our BlochNet. BlochNet exploits physics-based decoder for helping anti-causal tasks of the encoder.

#### 3.0.2 Encoder and Decoder

In the proposed approach, the encoder uses a neural network to predict T1, T2 parameters from input signatures. A three-layers fully connected neural network is applied for fixed length sequence inputs and a recurrent neural network(RNN) is for flexible length of inputs. Considering that the input signature is generated with strong temporal structure, RNN is adopted to capture the temporal information effectively from the input and leads to predicting tissue parameters precisely. In contrast, the decoder leverages the Bloch equation based MRF physics model, acting as a causal regularization, to perform supervised reconstruction to the input signatures from estimated tissue properties $\Theta$ as well as associated pseudo-random excitation pulse sequence settings $\Phi$, such as repetition time (TR), time of echo (TE), and radio frequency flip angle (FA) over time.

Given an enquiry signature $X_n \in \mathbb{C}^L$ for the $n$-th voxel, the encoder $E$ solves an anti-causal inverse problem and outputs predicted tissue properties $\hat{\Theta}_n = \{\hat{T1}_n, \hat{T2}_n\}$. This operation nonlinearly maps the input signature from a high dimensional manifold to a low dimensional manifold.

$$\hat{\Theta}_n = \mathcal{E}(X_n) \quad \forall n \in 1, \ldots, N$$

Given the RF sequence settings $\Phi$ and the estimated tissue properties $\hat{\Theta}_n$, the decoder reconstructs the input signature via solving Bloch equations using extended phase graph (EPG) formalism [22, 52].

$$\hat{X}_n = \mathcal{B}(\hat{\Theta}_n|\Phi) \quad \forall n \in 1, \ldots, N$$

where $\mathcal{B}$ denotes the Bloch equation based decoder (see Appendix for more details).

#### 3.0.3 Fast Bloch decoder based on efficient EPG

In MRI, the magnetic field is composed of a static magnetic field and a dynamic component which is manipulated through a radio frequency (RF) coil aligned with the x direction. The overall macroscopic dynamics of the net magnetization can be summarized by the Bloch equations which are composed of a set of linear ordinary differential equations:

$$\frac{d\vec{M}}{dt} = \vec{M} \times \gamma\vec{B} - \frac{\vec{M}_{xy}}{T2} - \frac{\vec{M}_z - \vec{M}_0}{T1} = \vec{M} \times \gamma\vec{B} - \begin{bmatrix} M_x/T2 \\ M_y/T2 \\ (M_z - M_0)/T1 \end{bmatrix}$$

where $\vec{M}$ is magnetization with $\vec{M}_{xy}$ and $\vec{M}_z$ as the transverse and longitudinal components, respectively. $\vec{M}_0$ is the equilibrium magnetization; $\vec{B}$ is the magnetic field; $\gamma$ is the gyromagnetic ratio.

Since there is no general analytic solution to the Bloch equations, numerical solutions such as EPG formalization are often adopted. However, a significant limitation of the released EPG code [52] is its slow computation speed in solving the Bloch equations. To circumvent this, recurrent neural network [36] and generative adversarial networks [53] have been applied as surrogates for the Bloch equation. However, these require a lot of training data and still may generate inaccurate sequences on unseen tissue parameters and RF pulse settings due to overfitting. Instead, we adapt the EPG code [52] to achieve a much more efficient implementation, making it practical to use the exact MRI physics model as a decoder in the training procedure. A key change involves incorporating the torch jit package, and using batch-wise computation for the 3 Bloch stages, leading to 500 times faster generation of magnetic responses for 1,000 sequences on CPU. [2] More details can be found in the appendix.

[2] Our code, including the fast and end-to-end solvable EPG-Bloch code, is released on our GitHub repository at *[text left blank during peer review]*.

#### 3.0.4 LOSS FUNCTION

The loss function consists of two parts: the mean squared error (MSE) between the ground truth and the predicted tissue properties, referred to as embedding loss, and the MSE between the input and the reconstructed signatures, referred to as reconstruction loss,

$$\mathcal{L} = \frac{1}{N}\sum_{n=1}^{N}\left(\frac{1}{2}\|\hat{\Theta}_n - \Theta_n\|_2^2 + \frac{1}{2}\|\hat{X}_n - X_n\|_2^2\right) \tag{1}$$

## 4 EXPERIMENT RESULTS

In this section, we perform evaluation on the proposed method and conduct comparison with other state-of-the-art MRF methods. We evaluate the generalization performance of all models across different data distributions and different RF pulse sequences.

### 4.1 DATA SETTINGS

#### 4.1.1 SYNTHETIC DATA

The synthetic signature data $X$ is generated by solving the Bloch equations using the adapted fast EPG formulation for a set of ground-truth tissue properties $\Theta = \{T1, T2\}$, given the fixed RF pulse sequences with setting $\Phi = \{FA, TR, TE\}$. Specifically, the tissue properties $\Theta$ are combinatorial pairs of $T1$ values that range from 0 to 5000, and $T2$ values that range from 0 to 2500, following the settings used in [42, 48].

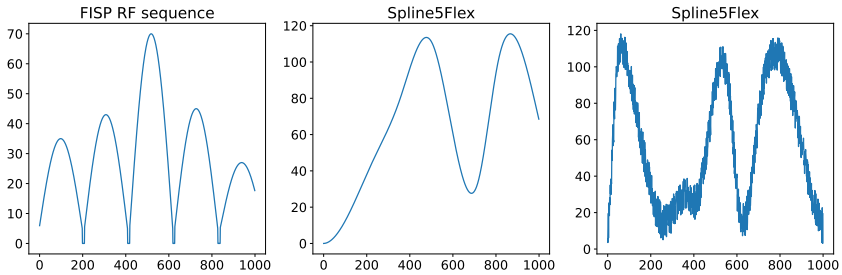


(a) Flip angles of three RF pulse sequences

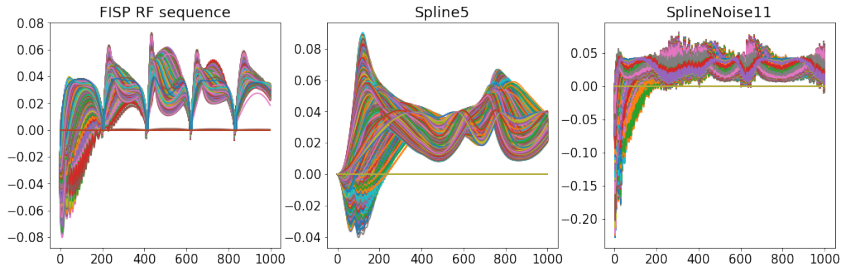


(b) Magnetic responses

Figure 3: Flip angles of three RF pulse sequences, including FISP, Spline5, Spline11Noisy, and corresponding magnetic responses generated using the Bloch equations

#### 4.1.2 PHANTOM MRI DATA

We exploit the fuzzy version of the brain phantom of the BrainWeb Brain Database [8] to construct a set of realistic, high-resolution T1 maps and T2 maps as ground truth, which faithfully exhibit spatial distribution for different tissue compositions, including CSF, Grey matter, White matter, etc. Then, we select a few slices from each pair of T1/T2 phantom volume across 10 subjects. After removing non-brain parts, these slices are then vectorized as column vectors and stacked as a 2D

matrix, leading to the ground truth tissue properties $\Theta \in \mathbb{R}^{N\times 2}$, where $N = 85,645$ denotes the total number of voxels from all slices. Given $\Theta$, signatures $X \in \mathcal{R}^{N\times L}$ are generated using the Bloch equation-based MRI physics model. (More details are provided in the appendix.)

#### 4.1.3 ANATOMICAL MRI DATA

The anatomical data is from [42, 48]. Specifically, brain MRI scans were first acquired from a healthy subject using GE Signa 3T HDXT scanner with Fast Imaging Employing Steady-state Acquisition (FIESTA) and Spoiled Gradient Recalled Acquisition in Steady State (SPGR) at four different flip angles $(3^\circ, 5^\circ, 12^\circ, 20^\circ)$. Then corrections [35] are implemented, followed by using DESPOT1 and DESPOT2 algorithms [10] to obtain T1, T2 maps of size of $128 \times 128$ for reference, leading to $\Theta \in \mathbb{R}^{N\times 2}$ with $N = 7,499$ voxels after removing non-brain parts. Then the signatures $X$ are generated from the Bloch equations with FISP RF pulse sequence.

### 4.2 BASELINE METHODS

We compare our approach with 6 representative state-of-the-art MRF methods, including dictionary matching (DM) [38], Fully-connected deep neural network (FC) [7], Hybrid deep learning (HYDRA) [48] as well as two auto-encoder methods with RNN encoder and RNN decoder(RNN-RNN) and FC encoder and FC decoder(FC-FC), respectively. The first three models are trained with only embedding loss, whereas auto-encoder models have an additional reconstruction loss for the decoder, apart from the embedding loss for the encoder. In constrast to the two auto-encoder models where the decoder is learned together with the encoder during training, our physics-based model exploits the Bloch equation based decoder and keeps the decoder fixed, while only the encoder is updated during training. The idea is to use the decoder as a causal regularization which gives reconstruction loss to guide the encoder to find more generalizable representations than baseline models.

| | Dictionary matching | FC | RNN | HYDRA | Autoencoder (FC-FC) | Autoencoder (RNN-RNN) | BlochNet (FC-Bloch) |
|---|---|---|---|---|---|---|---|
| Phantom data | 18.6652 | 0.0554 | 0.0486 | 0.1597 | 0.0519 | 0.0443 | **0.0409** |
| Anatomical data | 19.4883 | 0.0812 | 0.092 | 0.3088 | 0.0801 | 0.0889 | **0.0748** |

Table 1: Generalization performance across **different data distributions**: synthetic data for training while phantom (top row) and anatomical data (bottom row) for testing.

### 4.3 EXPERIMENTS OF EVALUATING GENERALIZATION PERFORMANCE

We evaluate the generalization performance of various models on two types of experiment settings: 1) across different data distributions, including synthetic, phantom and anatomical MRF data; 2) across different RF pulse sequences with different flip angles.

#### 4.3.1 GENERALIZATION ACROSS DIFFERENT DATA DISTRIBUTIONS

In practice, acquiring anatomical MRF data with ground truth T1, T2 values is time-consuming and expensive. Due to limited labeled anatomical data, it is common practice to use a large amount of synthetic data to train models to avoid potential overfitting, and then perform validation on anatomical data [5, 9, 17, 34, 41, 42, 48, 50]. Following the same routine, we perform model training on synthetic MRF data, followed by model testing on phantom and anatomical data, in order to evaluate the generalization performance of trained models across different data distributions. In this experiment, we perform model training on synthesized MRF data introduced in Section 4.1.1 for our

| | Dictionary matching | FC | RNN | HYDRA | Autoencoder (FC-FC) | Autoencoder (RNN-RNN) | BlochNet (FC-Bloch) |
|---|---|---|---|---|---|---|---|
| Phantom data | 27.1955 | 0.8415 | 0.7973 | 0.3593 | 0.7018 | 0.6148 | **0.1574** |
| Anatomical data | 16.6151 | 0.7305 | 1.06603 | 0.695 | 0.4081 | 0.9385 | **0.2667** |

Table 2: Generalization performance across **different RF pulse sequences**: Spline5 and SplineNoisy11 in training, FISP in testing.

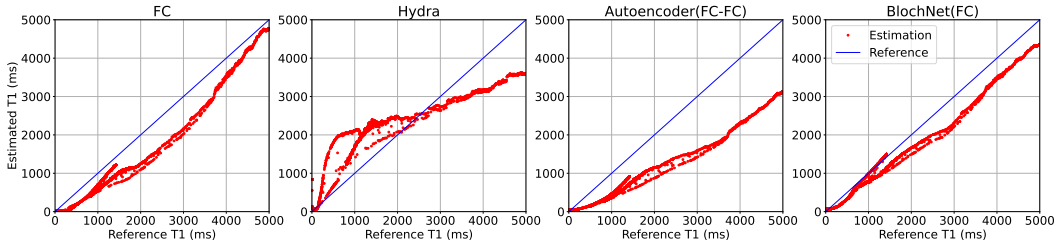


(a) T1 gold standard and predicted values for four models.

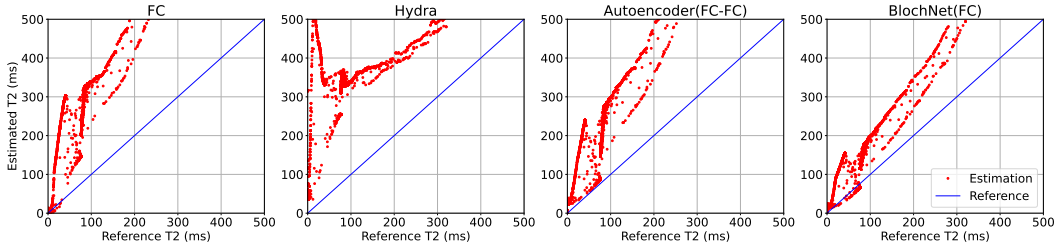


(b) T2 gold standard and predicted values for four models.

Figure 4: Comparison of the generalization performance across **different RF pulse sequences** for 4 models. Blue line: gold-standard. Red dots: predicted values for tissue properties T1 and T2.

model and baseline models, and then compare their performance on unseen phantom and anatomical data described in Section 4.1.2 and 4.1.3.

Table 1 includes the mean squared error (MSE) between the ground truth and predicted tissue properties for seven approaches on phantom and anatomical data. Note that these errors are computed in log-scale. As shown in the table, the dictionary-matching approach gave the worst performance, because the pre-computed dictionary and LUT did not cover the OOD data samples that could be quite different from the already contained dictionary entries. This caused poor inner product based matching results, and restricted the generalization.

Interestingly, the results show that the reconstruction loss provides benefits to between-data generalization for autoencoder models(FC-FC or RNN-RNN), in comparison with non-autoencoder models(FC or RNN), respectively, on both phantom and anatomical MRF data. Furthermore, our BlochNet outperforms all other models, indicating that reconstruction loss from the physics-based decoder has the best regularization effect that contributes to improved encoder training. In particular, when the physics-based decoder reconstructs input signatures from estimated tissue properties using Bloch equations, the corresponding reconstruction loss serves as additional guidance and constraint for the encoder to effectively capture the underlying anti-causal mapping. This in turn leads to better generalization performance.

Figure 5 shows the predicted tissue properties using various models on anatomical MRF data. All models perform well in the middle range of tissue properties and lead to small errors. However, each individual models show different prediction characteristics. Specifically, HYDRA, as shown in the third column of 5, suffers from a higher loss at the rim region of the brain, and leads to larger errors than other models. Autoencoder(FC-FC) model, as shown in the forth column, demonstrates a better prediction of T2 values than the non-autoencoder(FC) model in the second column. The proposed BlochNet outperforms other comparison models with the least prediction error and most stable performance across the whole range of both T1 and T2 values, as shown in the fifth column.

#### 4.3.2 Generalization across different RF pulse sequences

In this experiment, we perform model training on one RF pulse sequence and evaluate the trained models on another different RF pulse sequence. Specifically, we adopted 3 different RF pulse sequences, including FISP [38], Spline5 [36], Spline11Noisy [36] with their flip angles shown in Figure 3a. (More details are provided in Appendix.) FISP is used exclusively in the testing stage,

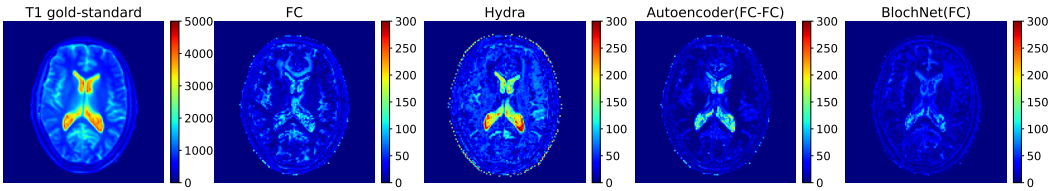


(a) Gold-standard T1 and errors between gold-standard and predicted T1 values of four models.

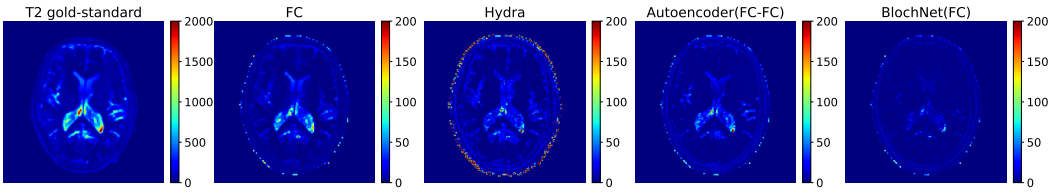


(b) Gold-standard T2 and errors between gold-standard and predicted T2 values of four models.

Figure 5: Generalization performance across different data distributions: data in training is synthetic while in testing is anatomical MRI for four different models.

while Spline5 and Spline11Noisy are used exclusively in the training stage. Under such settings, the performance of our BlochNet and other six models is compared in Table 2.

Both Table 2 show that the embedding loss for all models are much higher than Table 1 regardless of phantom and anatomical MRF data. This is because the signatures generated with Spline5 and Spline11Noisy RF pulse sequences have notable difference from signatures for FISP, as shown in Figure 3b, which are particularly challenging cases. In spite of degraded performance for all models, the results clearly show the advantage of autoencoder (FC-FC or RNN-RNN) models over non-autoencoder models(FC or RNN), which confirms the benefits of incorporating an decoder to derive the reconstruction loss as additional regularization. Furthermore, the proposed BlochNet demonstrates significant gains over the competing methods in such challenging cases on both phantom and anatomical MRF data. For example, the reconstruction loss of 7 baseline models is larger than 0.3 while our BlochNet gives 0.1574 on phantom MRI data. In all these settings, our approach leads to at least 50% smaller loss than the best competing approach.

These differences in MSE losses are shown in Figure 4. FC model (left) makes poor predictions on both T1 and T2 values with high variance, predicting different values even when the sequence is generated from the same tissue properties. This happens when the model cannot infer tissue properties from input signatures generated from different combinations of T1 and T2 values. Autoencoder(FC-FC) model (third column) shows more aligned and better inferences with lower variance, but still has high deviation between predicted and gold-standard values. In comparison, our BlochNet outputs predictions that are closest to the gold-standard with the lowest error. This confirms the benefits of our physics-based decoder that guides the encoder to learn the underlying anti-causal mechanism effectively.

## 5 Conclusion

We propose BlochNet, a physics-based auto-encoder model, to perform quantification of multiple tissue properties from their magnetic responses, a challenging ill-posed inverse problem in quantitative MRI. We believe this is the first practical approach to incorporate a Bloch equation based MRI physics model as the causal regularization for this anti-causal task. By its design, the approach captures the underlying obscure anti-causal mapping effectively. Experiments demonstrate that our method consistently outperforms competing methods with better robustness and generalizability.

In future work, we will consider the k-space subsampling circumstance that may lead to faster and more efficient QMRI/MRF. We will also explore more varied RF settings, for example, time-varying TR, TE and FA, as it has been reported [2] that a RF pulse sequence with multiple pseudo-random parameters may improve the identifiability of resulted magnetic responses.

## REFERENCES


[1] Michel Besserve, Arash Mehrjou, Rémy Sun, and Bernhard Schölkopf. Counterfactuals uncover the modular structure of deep generative models. *arXiv preprint arXiv:1812.03253*, 2018.

[2] Bhairav Bipin Mehta, Simone Coppo, Debra Frances McGivney, Jesse Ian Hamilton, Yong Chen, Yun Jiang, Dan Ma, Nicole Seiberlich, Vikas Gulani, and Mark Alan Griswold. Magnetic resonance fingerprinting: a technical review. *Magnetic resonance in medicine*, 81(1): 25–46, 2019.

[3] Felix Bloch. Nuclear induction. *Physical review*, 70(7-8):460, 1946.

[4] Shengze Cai, Zhicheng Wang, Sifan Wang, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks for heat transfer problems. *Journal of Heat Transfer*, 143 (6), 2021.

[5] Xiaozhi Cao, Congyu Liao, Zhixing Wang, Ying Chen, Huihui Ye, Hongjian He, and Jianhui Zhong. Robust sliding-window reconstruction for accelerating the acquisition of mr fingerprinting. *Magnetic resonance in medicine*, 78(4):1579–1588, 2017.

[6] Dongdong Chen, Mike E Davies, and Mohammad Golbabaee. Deep unrolling for magnetic resonance fingerprinting. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pp. 1–4. IEEE, 2022.

[7] Ouri Cohen, Bo Zhu, and Matthew S Rosen. MR fingerprinting deep reconstruction network (drone). *Magnetic resonance in medicine*, 80(3):885–894, 2018.

[8] D Louis Collins, Alex P Zijdenbos, Vasken Kollokian, John G Sled, Noor J Kabani, Colin J Holmes, and Alan C Evans. Design and construction of a realistic digital brain phantom. *IEEE transactions on medical imaging*, 17(3):463–468, 1998.

[9] Mike Davies, Gilles Puy, Pierre Vandergheynst, and Yves Wiaux. A compressed sensing framework for magnetic resonance fingerprinting. *SIAM Journal on Imaging Sciences*, 7(4): 2623–2656, 2014.

[10] Sean CL Deoni, Terry M Peters, and Brian K Rutt. High-resolution t1 and t2 mapping of the brain in a clinically acceptable time with despot1 and despot2. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 53(1):237–241, 2005.

[11] Guillaume Desjardins, Aaron Courville, and Yoshua Bengio. Disentangling factors of variation via generative entangling. *arXiv preprint arXiv:1210.5474*, 2012.

[12] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38(2):295–307, 2016.

[13] N Benjamin Erichson, Michael Muehlebach, and Michael W Mahoney. Physics-informed autoencoders for lyapunov-stable fluid flow prediction. *arXiv preprint arXiv:1905.10866*, 2019.

[14] Zhenghan Fang, Yong Chen, Mingxia Liu, Lei Xiang, Qian Zhang, Qian Wang, Weili Lin, and Dinggang Shen. Deep learning for fast and spatially constrained tissue quantification from highly accelerated data in magnetic resonance fingerprinting. *IEEE transactions on medical imaging*, 38(10):2364–2374, 2019.

[15] Zhenghan Fang, Yong Chen, Dong Nie, Weili Lin, and Dinggang Shen. Rca-u-net: Residual channel attention u-net for fast tissue quantification in magnetic resonance fingerprinting. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 101–109. Springer, 2019.

[16] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. *arXiv preprint arXiv:1705.03122*, 2017.

[17] Mohammad Golbabaee, Guido Buonincontri, Carolin M Pirkl, Marion I Menzel, Bjoern H Menze, Mike Davies, and Pedro A Gómez. Compressive mri quantification using convex spatiotemporal priors and deep encoder-decoder networks. *Medical Image Analysis*, 69:101945, 2021.

[18] Oded Goldreich. *Foundations of cryptography: volume 2, basic applications*. Cambridge university press, 2009.

[19] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.

[20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE Conf. Comput. Vision Pattern Recog*, pp. 770–778, 2016.

[21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Proc. Eur. Conf. Comput. Vision*, pp. 630–645. Springer, 2016.

[22] Jurgen Hennig. Echoes – how to generate, recognize, use or avoid them in mr-imaging sequences. part i: Fundamental and not so fundamental properties of spin echoes. *Concepts in Magnetic Resonance*, 3(3):125–143, 1991.

[23] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal processing magazine*, 29(6):82–97, 2012.

[24] Elisabeth Hoppe, Gregor Körzdörfer, Tobias Würfl, Jens Wetzl, Felix Lugauer, Josef Pfeuffer, and Andreas Maier. Deep learning for magnetic resonance fingerprinting: A new approach for predicting quantitative parameter values from time series. *Stud Health Technol Inform*, 243: 202–206, 2017.

[25] Truong Son Hy, Shubhendu Trivedi, Horace Pan, Brandon M Anderson, and Risi Kondor. Predicting molecular properties with covariant compositional networks. *The Journal of chemical physics*, 148(24):241745, 2018.

[26] George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

[27] Niki Kilbertus, Giambattista Parascandolo, and Bernhard Schölkopf. Generalization in anticausal learning. *arXiv preprint arXiv:1812.00524*, 2018.

[28] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Deeply-recursive convolutional network for image super-resolution. In *Proc. IEEE Conf. Comput. Vision Pattern Recog*, pp. 1637–1645, 2016.

[29] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.

[30] Trent Kyono, Yao Zhang, and Mihaela van der Schaar. Castle: regularization via auxiliary causal graph discovery. In *Advances in Neural Information Processing Systems*, volume 33, pp. 1501–1512, 2020.

[31] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.

[32] Lei Le, Andrew Patterson, and Martha White. Supervised autoencoders: Improving generalization performance with unsupervised regularizers. *Advances in neural information processing systems*, 31, 2018.

[33] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436, 2015.

[34] Congyu Liao, Berkin Bilgic, Mary Kate Manhard, Bo Zhao, Xiaozhi Cao, Jianhui Zhong, Lawrence L Wald, and Kawin Setsompop. 3d mr fingerprinting with accelerated stack-of-spirals and hybrid sliding-window and grappa reconstruction. *Neuroimage*, 162:13–22, 2017.

[35] Gilad Liberman, Yoram Louzoun, and Dafna Ben Bashat. T1 mapping using variable flip angle spgr data with flip angle correction. *Journal of Magnetic Resonance Imaging*, 40(1):171–180, 2014.

[36] Hongyan Liu, Oscar van der Heide, Cornelis AT van den Berg, and Alessandro Sbrizzi. Fast and accurate modeling of transient-state, gradient-spoiled sequences by recurrent neural networks. *NMR in Biomedicine*, 34(7):e4527, 2021.

[37] Weiyang Liu, Zhen Liu, Liam Paull, Adrian Weller, and Bernhard Schölkopf. Structural causal 3d reconstruction. In *Proc. Eur. Conf. Comput. Vision*, 2022.

[38] Dan Ma, Vikas Gulani, Nicole Seiberlich, Kecheng Liu, Jeffrey L Sunshine, Jeffrey L Duerk, and Mark A Griswold. Magnetic resonance fingerprinting. *Nature*, 495(7440):187, 2013.

[39] Zhiping Mao, Ameya D Jagtap, and George Em Karniadakis. Physics-informed neural networks for high-speed flows. *Computer Methods in Applied Mechanics and Engineering*, 360: 112789, 2020.

[40] Andreas Maurer, Massimiliano Pontil, and Bernardino Romera-Paredes. The benefit of multi-task representation learning. *Journal of Machine Learning Research*, 17(81):1–32, 2016.

[41] Gal Mazor, Lior Weizman, Assaf Tal, and Yonina C Eldar. Low rank magnetic resonance fingerprinting. In *Engineering in Medicine and Biology Society (EMBC), 2016 IEEE 38th Annual International Conference of the*, pp. 439–442. IEEE, 2016.

[42] Gal Mazor, Lior Weizman, Assaf Tal, and Yonina C Eldar. Low-rank magnetic resonance fingerprinting. *Medical physics*, 45(9):4066–4084, 2018.

[43] George S Misyris, Andreas Venzke, and Spyros Chatzivasileiadis. Physics-informed neural networks for power systems. In *2020 IEEE Power & Energy Society General Meeting (PESGM)*, pp. 1–5. IEEE, 2020.

[44] Ilkay Oksuz, Gastao Cruz, James Clough, Aurelien Bustin, Nicolo Fuin, Rene M Botnar, Claudia Prieto, Andrew P King, and Julia A Schnabel. Magnetic resonance fingerprinting using recurrent neural networks. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pp. 1537–1540. IEEE, 2019.

[45] Judea Pearl. *Causality*. Cambridge university press, 2009.

[46] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.

[47] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

[48] Pingfan Song, Yonina C Eldar, Gal Mazor, and Miguel RD Rodrigues. Hydra: Hybrid deep magnetic resonance fingerprinting. *Medical physics*, 46(11):4951–4969, 2019.

[49] Refik Soyak, Ebru Navruz, Eda Ozgu Ersoy, Gastao Cruz, Claudia Prieto, Andrew P King, Devrim Unay, and Ilkay Oksuz. Channel attention networks for robust mr fingerprint matching. *IEEE Transactions on Biomedical Engineering*, 69(4):1398–1405, 2021.

[50] Zhe Wang, Hongsheng Li, Qinwei Zhang, Jing Yuan, and Xiaogang Wang. Magnetic resonance fingerprinting with compressed sensing and distance metric learning. *Neurocomputing*, 174: 560–570, 2016.

[51] Sebastian Weichwald, Bernhard Schölkopf, Tonio Ball, and Moritz Grosse-Wentrup. Causal and anti-causal learning in pattern recognition for neuroimaging. In *2014 International Workshop on Pattern Recognition in Neuroimaging*, pp. 1–4. IEEE, 2014.

[52] Matthias Weigel. Extended phase graphs: dephasing, rf pulses, and echoes-pure and simple. *Journal of Magnetic Resonance Imaging*, 41(2):266–295, 2015.

[53] Mingrui Yang, Yun Jiang, Dan Ma, Bhairav B Mehta, and Mark A Griswold. Game of learning bloch equation simulations for mr fingerprinting. *arXiv preprint arXiv:2004.02270*, 2020.