

Document Version

Final published version

Licence

Dutch Copyright Act (Article 25fa)

Citation (APA)

Capitão, N., Zhang, Y., Zhao, Y., & Tao, Q. (2026). Physics-Informed Neural ODEs for Temporal Dynamics Modeling in Cardiac T1 Mapping. In J. C. Gee, D. C. Alexander, J. Hong, J. E. Iglesias, C. H. Sudre, A. Venkataraman, P. Golland, J. H. Kim, & J. Park (Eds.), *Medical Image Computing and Computer Assisted Intervention, MICCAI 2025 - 28th International Conference, 2025, Proceedings* (pp. 492-501). (Lecture Notes in Computer Science; Vol. 15960 LNCS). Springer. https://doi.org/10.1007/978-3-032-04927-8_47

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership. Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Physics-Informed Neural ODEs for Temporal Dynamics Modeling in Cardiac T_1 Mapping

Nuno Capitão^{1,2}, Yi Zhang¹, Yidong Zhao¹, and Qian Tao¹(✉)

¹ Department of Imaging Physics, Delft University of Technology, Delft, The Netherlands

up201709363@up.pt, {y.zhang-43,y.zhao-8,q.tao}@tudelft.nl

² University of Porto, Porto, Portugal

Abstract. Spin-lattice relaxation time (T_1) is an important biomarker in cardiac parametric mapping for characterizing myocardial tissue and diagnosing cardiomyopathies. Conventional Modified Look-Locker Inversion Recovery (MOLLI) acquires 11 breath-hold baseline images with interleaved rest periods to ensure mapping accuracy. However, prolonged scanning can be challenging for patients with poor breathholds, often leading to motion artifacts that degrade image quality. In addition, T_1 mapping requires a voxel-wise nonlinear fitting to a signal recovery model involving an iterative estimation process. Recent studies have proposed deep-learning approaches for rapid T_1 mapping using shortened sequences to reduce acquisition time for patient comfort. Nevertheless, existing methods overlook important physics constraints, limiting inter-pretability and generalization. In this work, we present an accelerated, end-to-end T_1 mapping framework leveraging Physics-Informed Neural Ordinary Differential Equations (ODEs) to model temporal dynamics and address these challenges. Our method achieves high-accuracy T_1 estimation from a sparse subset of baseline images and ensures efficient null index estimation at the test time. Specifically, we develop a continuous-time LSTM-ODE model to enable selective Look-Locker (LL) data acquisition with arbitrary time lags. Experimental results show superior performance in T_1 estimation for both native and post-contrast sequences and demonstrate the strong benefit of our physics-based formulation over direct data-driven T_1 priors.

Keywords: Quantitative Cardiac MRI · Physics-Informed Neural Networks · Modified Look-Locker Inversion Recovery

1 Introduction

Quantitative Magnetic Resonance Imaging (qMRI) is a widely adopted noninvasive technique for cardiac parametric mapping, where spin-lattice relaxation

N. Capitão and Y. Zhang—Contributed equally to this work.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2026
J. C. Gee et al. (Eds.): MICCAI 2025, LNCS 15960, pp. 492–501, 2026.
https://doi.org/10.1007/978-3-032-04927-8_47

time (T_1) is routinely used as a biomarker to clinically assess myocardial tissue and diagnose cardiomyopathies [5]. Cardiac T_1 maps are generated by fitting a parametric model to a series of baseline images that encode signal recovery governed by spin-lattice energy exchange after an inversion pulse.

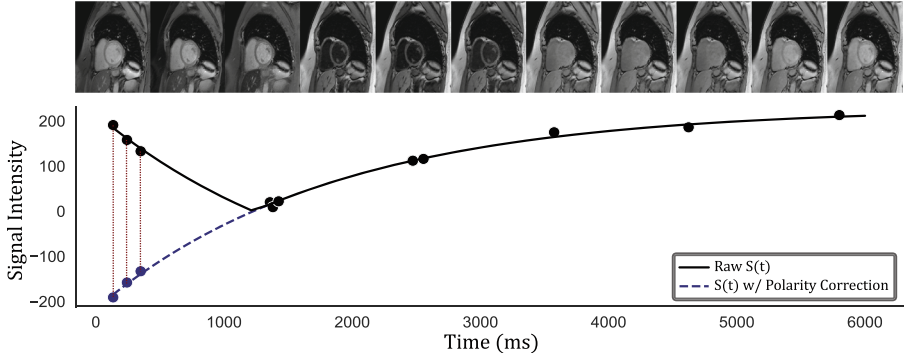


Fig. 1. MOLLI 3(3)3(3)5 reconstruction scheme illustrating signal recovery with and without polarity recovery. The corresponding temporally ordered MOLLI T_1 -weighted images are shown above.

The Modified Look-Locker Inversion Recovery (MOLLI) [13] sequence has become the standard protocol in cardiac qMRI due to its high precision and reduced motion-induced artifacts. During a single breath hold, MOLLI acquires T_1 -weighted images at multiple inversion times using distinct Look-Locker (LL) experiments interleaved with rest periods. In a MOLLI 3(3)3(3)5 sequence, three LL experiments collect images in sets of 3, 3, and 5, each separated by three heartbeats. These acquisitions form a image series of 11 baseline images that span the full relaxation curve, as illustrated in Fig. 1. While multiple LL experiments at different inversion times improve T_1 estimation accuracy, they also extend the breath-hold to 17 heartbeats, which can be strenuous for patients with respiratory issues. Alternative schemes such as Shortened MOLLI (ShMOLLI) [14] and Saturation Recovery Single-Shot Acquisition (SASHA) [3] reduce the total number of acquisitions to accelerate imaging, but inherently trade off precision and accuracy [16, 18].

Traditional T_1 map estimation requires voxel-wise nonlinear curve fitting to a signal recovery model derived from MRI physics. This process is further complicated by the need to identify signal polarity transitions (null points) in MOLLI sequences, requiring additional iterative calculations for each pixel. The entire procedure is not only sensitive to image noise but also computationally intensive - processing a complete 3D volume over time has $\mathcal{O}(n^4)$ complexity, resulting in significant processing delays in clinical settings without specialized acceleration techniques [20].

Recent studies have shown that deep learning methods can accurately and efficiently perform cardiac T_1 mapping using as few as three to five baseline

images along the inversion recovery curve, achieving results comparable to conventional MOLLI. [4] introduced MyoMapNet, a fast myocardial T_1 mapping model that uses a fully connected neural network (FCNN) to directly map pairwise LL readouts to T_1 values. However, it neglects signal polarities and is constrained to a data-driven T_1 prior. In contrast, T1Net [11] employs a bidirectional long short-term memory (LSTM) network [6] and adheres to the signal recovery model to enforce consistency between acquired MRI signals and predicted T_1 values through its loss function. Nonetheless, we hypothesize that T1Net remains limited for several reasons: First, estimating T_1 maps from a single LL experiment may require non-uniform time intervals to capture the full relaxation curve. Like other discrete recurrent neural networks, LSTMs struggle with irregular time gaps [1, 9, 17]. Furthermore, although enforcing the signal recovery model provides a stronger prior than direct T_1 estimation, the mechanism by which the model learns the relaxation rate remains a black box, limiting both interpretability and generalization.

These limitations highlight the need for a more principled approach that can handle irregular sampling patterns while maintaining physical consistency with the underlying MR signal dynamics. To address these challenges, we introduce a robust and accelerated T_1 mapping technique using physics-informed neural networks (PINNs) [15]. We employ a continuous-time LSTM-ODE [12] that can naturally handle LL sequences with non-uniform acquisition intervals, overcoming the time-gap limitations of traditional recurrent networks. Our Neural Ordinary Differential Equation (NODE) [2] approach explicitly models the temporal dynamics of the relaxation process, providing a continuous representation of signal evolution rather than discrete approximations. This dynamic modeling not only adheres to the signal recovery model but also explicitly formulates the recovery rate along the relaxation trajectory, ensuring that T_1 estimation follows the actual physics of spin-lattice interactions. Furthermore, the framework seamlessly incorporates signal polarity correction during inference, eliminating the need for separate processing steps.

2 Methods

2.1 MR Physics-Informed Signal Modeling and Loss Functions

To develop our physics-informed approach, we first formalize the MR signal model that governs T_1 relaxation. In qMRI, N baseline images are collected via selective sampling, with each image capturing the signal at a different point along the relaxation curve. The signal intensity $S(x, y, z)$ at a given spatial coordinate (x, y, z) is described by a parametric function that reflects the underlying physics. In a MOLLI sequence, the signal recovery follows a 3-parameter model expressed as:

$$S_{(x,y,z)}(t_i) = c_{(x,y,z)} \cdot \left(1 - k_{(x,y,z)} \cdot \exp\left(-\frac{t_i}{T_{1,(x,y,z)}^*}\right)\right), \quad (1)$$

where $i \in \{1, 2, \dots, N\}$ and t_i is the inversion time of the i -th image. The parameter set $\{c, k, T_1^*\}$ is estimated at each voxel $v \in \mathbb{R}^{x \times y \times z}$ to determine T_1 biomarker using the following equation:

$$T_{1(x,y,z)} = T_{1,(x,y,z)}^* \cdot (k_{(x,y,z)} - 1). \quad (2)$$

While Eqs. (1) and (2) provide the foundation for conventional T_1 mapping, they represent a static view of the relaxation process. Our approach incorporates the underlying physical dynamics governing spin-lattice relaxation by explicitly modeling the temporal evolution of the signal. By differentiating Eq. 1, we derive a first-order differential expression that captures the instantaneous rate of change at each inversion time t_i :

$$\frac{dS_{(x,y,z)}(t_i)}{dt} = \frac{c_{(x,y,z)} \cdot k_{(x,y,z)}}{T_{1,(x,y,z)}^*} \cdot \exp\left(-\frac{t_i}{T_{1,(x,y,z)}^*}\right). \quad (3)$$

This differential formulation enables our model to learn not just the signal values but also their temporal dynamics, providing a more complete physical representation of the relaxation process.

To incorporate both data-driven priors and physics-based constraints into our deep-learning framework, we define a physics-informed loss function with two components: a T_1 consistency loss \mathcal{L}_{T_1} and a physics-based loss $\mathcal{L}_{\text{physics}}$.

$$\mathcal{L}_{T_1} = \frac{1}{HWD} \sum_{(x,y,z)} \|T_1 - \hat{T}_1\|_2^2, \quad (4)$$

where T_1 is derived from Eq. (2), enforcing consistency with our end-goal parametric mapping.

$$\mathcal{L}_{\text{physics}} = \frac{1}{NHWD} \sum_{(x,y,z)} \sum_{i=1}^N \|S(t_i) - \hat{S}(t_i)\|_2^2 + \lambda \left\| \gamma \frac{dS(t_i)}{dt} - \frac{d\hat{S}(t_i)}{dt} \right\|_2^2. \quad (5)$$

The first term ensures adherence to the signal recovery model (Eq. (1)), while the second term penalizes deviations from the expected signal dynamics (Eq. (3)). The scaling factor λ adjusts for the disparity in orders of magnitude observed at $\lim_{t \rightarrow 0} \frac{dS(t)}{dt}$ relative to the other loss terms. The factor γ is derived from the chain rule and accounts for the normalization ratio applied to signals and inversion times. By combining these complementary loss components, we formulate our final loss function as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{T_1} + \mathcal{L}_{\text{physics}}. \quad (6)$$

2.2 Continuous-Time Neural ODEs for Temporal Dynamics Modeling

To effectively capture the temporal dynamics of spin-lattice relaxation with non-uniform sampling intervals, we implement a continuous-time LSTM-ODE framework that naturally handles the irregular time gaps in MOLLI sequences. Let

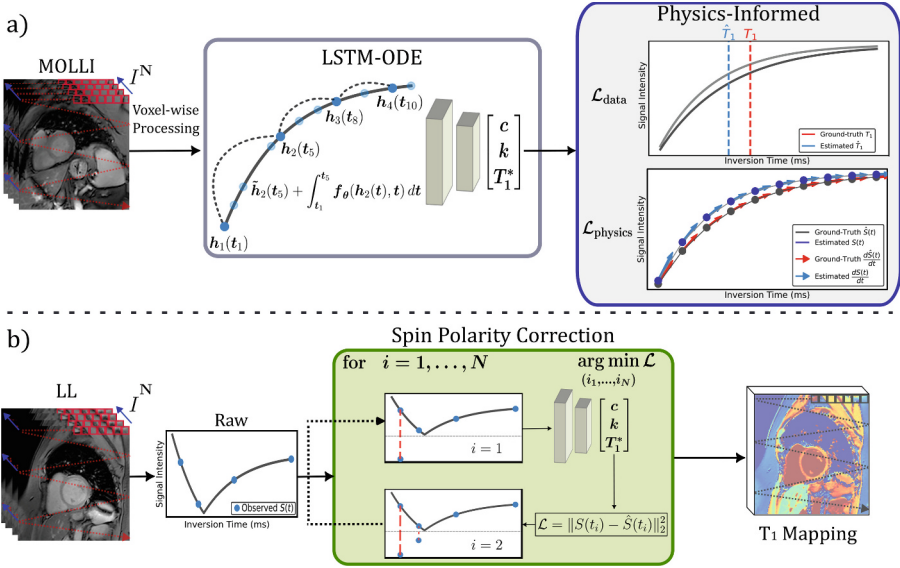


Fig. 2. Overview of the proposed physics-informed LSTM-ODE framework (a) **Training:** The model processes a subset of the 11 MOLLI baseline images in a voxel-wise manner through a continuous-time LSTM-ODE. The physics-informed loss is computed using $\{c, k, T_1^*\}$, estimated by the model. (b) **Inference:** For a series of I^N LL images, each voxel undergoes spin polarity correction. Every i -th image is sequentially inverted, and the resulting sequence is evaluated using the signal reconstruction loss. The index that minimizes the loss, $\arg \min \mathcal{L}$, determines the null index and the optimal parameters $\{c, k, T_1^*\}$.

$\{(S_i, t_i)\}_{i=1}^N$ represent a MOLLI voxel sequence, where S_i is the measured signal at inversion time t_i . At the i -th image, the LSTM’s discrete gating update is determined by

$$(\tilde{h}_i, c_i) = \text{LSTMCELL}((S_i, t_i), (h_{i-1}, c_{i-1})), \tag{7}$$

where h_{i-1} and c_{i-1} are the previous hidden and cell states, respectively, \tilde{h}_i denotes the candidate hidden state for the current image, and c_i represents the updated cell state. The candidate hidden state \tilde{h}_i is then evolved continuously over the interval $[t_{i-1}, t_i]$ by a NODE whose dynamics are governed by

$$\frac{dh(t)}{dt} = f_\theta(h(t), t), \tag{8}$$

where f_θ is a small, learnable neural network. We initialize the ODE at $h(t_{i-1}) = \tilde{h}_i$ and integrate forward using a black-box 5th-order Dormand-Prince solver [10]. The cell states $\{c_i\}_{i=1}^N$ remain discrete to preserve stable long-term gradient propagation [12]. The final hidden state, $h_N(t_N)$, provides a continuous-time encoding of the relaxation trajectory, capturing the temporal variations in signal

recovery across the N images. A multi-layer decoder then maps this encoding to the signal recovery model described by Eqs. (1-3). Formally, the mapping notation can be defined as $f_\theta : \{(S_i, t_i)\}_{i=1}^N \in \mathbb{R}^{N \times d_{\text{emb}} \times 2} \rightarrow \mathbb{R}^3$, where d_{emb} denotes the feature dimension of the embedded pairwise acquisitions, and the output represents the set of parameters $\{C, k, T_1^*\}$.

This continuous-time modeling approach enables our framework to accurately represent the underlying physics of relaxation processes while accommodating the practical constraints of clinical MRI acquisition protocols with varying inversion times.

2.3 Post-inversion Spin Polarity Correction

At test time, raw signal acquisitions retain their absolute magnitudes, which conflicts with our model’s assumptions about post-inversion signal recovery. However, since our method is formulated around the signal recovery model, it can also act as a proxy for retrieving the spin state polarity encoded in each image. To estimate the null index at each voxel, we evaluate all possible polarity permutations across the sequence by performing N forward passes, leveraging vectorized, batch-wise inference over the full field-of-view (FoV). Figure 2 illustrates our end-to-end T_1 parametric mapping framework, which includes the post-inversion spin polarity correction. During the polarity recovery step, each i -th signal is inverted, and the entire sequence undergoes a model evaluation to compute the signal reconstruction loss term in Eq. 5. For each voxel $v \in \mathbb{R}^{x \times y \times z}$, the $\text{arg min } \mathcal{L}$ across trials identifies the null index, providing the best estimate of $\{c, k, T_1^*\}$ for T_1 mapping.

3 Experiments

Dataset: We used a cardiac MRI dataset of 50 subjects acquired with a 3.0T Ingenia MR scanner (Philips Healthcare). Each subject underwent both pre- and post-contrast MOLLI sequences, each consisting of 11 baseline images using the same 3-3-5 scheme provided by the manufacturer. One to three axial slices were acquired at the base, mid-ventricular, and apical levels, and the left-ventricle myocardium was manually annotated as the region of interest (ROI). To prevent data leakage, we performed a subject-wise random split: 30 subjects for training, 5 for validation, and 15 for testing. To evaluate out-of-domain generalization in signal recovery patterns at higher relaxation rates induced by a contrast agent, the training set included only native sequences, while the validation and testing sets contained both pre- and post-contrast sequences.

Implementation Details: We pre-train our LSTM-ODE model on all 11 MOLLI images and fine-tune on subsets of 3-5 images to simulate accelerated acquisition. During training, sparse readouts are strategically selected to ensure coverage across the relaxation curve: early phase ($\{t_1-t_3\}$), intermediate phase

($\{t_4-t_8\}$), and convergence phase ($\{t_9-t_{11}\}$). All models are trained using ground-truth $\{c, k, T_1^*\}$ parameters fitted to the 11-image data with the Levenberg-Marquardt (LM) algorithm. For the physics-based loss, we generate 1 000 interpolated points with 75% linearly spaced up to 2000 ms to emphasize relaxation dynamics. Our ODE solver uses tolerances of $\epsilon = 0.001$ and runs on an NVIDIA GTX 1070 GPU. We employ standard backpropagation with a derivative loss hyperparameter of $\lambda = 0.01$. The code is publicly available at <https://github.com/nunomiguelc18/pinn-node-cardiac-t1mapping>.

Table 1. Mean T_1 bias and SD values from Monte Carlo simulations, expressed as mean(std). Methods: TRF, MMNet (MyoMapNet), T1N (T1Net), P-MM (PINN MyoMapNet), P-LSTM (PINN LSTM-ODE). The best-performing methods are in **bold** and second-best are underlined. For bias values, * indicates no statistical significance ($p > 0.05$).

Method	Mean Bias (ms)			Fitting SD (ms)		
	LL ₃	LL ₄	LL ₅	LL ₃	LL ₄	LL ₅
Native						
TRF	-9.63(38.21)	-12.90(31.25)	15.64(18.80)	58.12 (37.87)	49.66 (27.16)	46.83 (24.12)
MMNet	51.50(11.44)	21.49(10.13)	11.29(9.26)	91.57(36.67)	67.49(26.94)	<u>50.47</u> (24.36)
T1N	<u>-11.28</u> (8.83)	<u>-8.77</u> (11.75)	<u>3.11</u> *(9.74)	94.00(34.81)	69.17(29.87)	53.33(26.82)
P-MM	0.15 *(6.62)	5.39*(11.70)	7.71(8.90)	<u>82.74</u> (37.10)	<u>62.40</u> (31.39)	51.19(27.60)
P-LSTM	7.65(9.93)	2.62 *(7.74)	0.75 *(8.68)	91.18(38.74)	70.04(30.09)	50.65(27.43)
Post-Gd						
TRF	16.68(24.57)	7.64(14.57)	11.83(12.65)	46.49 (37.58)	39.95 (17.48)	43.27(16.95)
MMNet	<u>1.18</u> (22.29)	14.23(13.86)	11.40(11.58)	<u>73.67</u> (17.32)	<u>58.34</u> (15.55)	42.65(20.45)
T1N	8.75 (10.62)	<u>8.75</u> (10.42)	20.44(7.82)	79.99(21.78)	53.97(17.18)	41.00(16.05)
P-MM	-6.52(18.47)	5.94(11.82)	<u>11.83</u> (9.65)	86.93(22.67)	59.24(18.94)	40.29 (16.93)
P-LSTM	-12.37(18.12)	-1.12 *(11.48)	3.38 *(11.11)	110.75(52.40)	70.55(18.75)	42.92(20.47)

Comparative Studies: We benchmark our baseline model against three methods: MyoMapNet, T1Net (both designed for accelerated T_1 mapping), and SciPy’s Trust Region Reflective (TRF) algorithm for bounded parameter constraints. To evaluate whether our physics-based priors improve performance relative to direct data-driven T_1 priors, we also integrate MyoMapNet’s architecture into our framework for direct comparison.

Evaluation Metrics: We perform 100 Monte Carlo simulations per subject using random sampling to assess generalization across LL subsets. We evaluate two metrics: (1) mean T_1 bias relative to ground truth, and (2) fitting standard deviation (SD), which measures curve fitting quality in terms of T_1 at each voxel from the fitting residuals [7, 8, 19]. Both metrics are calculated within myocardial regions annotated by radiologists. Pixel-wise T_1 estimates are averaged over simulations to generate subject-level distributions, and systematic biases are evaluated using paired t -tests.

4 Results and Discussion

Table 1 presents our Monte Carlo simulation results, summarizing the mean T_1 bias and SD distributions for the Native and Post-Gd testing sets. Complementary qualitative comparisons for cardiac T_1 mapping are provided in Fig. 3. The PINN MyoMapNet model improved upon conventional MyoMapNet, demonstrating the benefits of incorporating physics during training. Our PINN LSTM-ODE method achieved the overall lowest mean bias for both Native and Post-Gd settings, indicating stronger generalization. However, its higher SD maps suggest that sparse representations may lead to inconsistent NODE-computed trajectories. The TRF algorithm performs subpar in terms of mean bias, highlighting the potential of deep learning for higher accuracy in T_1 mapping given a limited number of baseline images. Nevertheless, TRF reports the lowest SD map statistics, suggesting that deep learning methods still compromise precision compared to traditional nonlinear curve-fitting approaches.

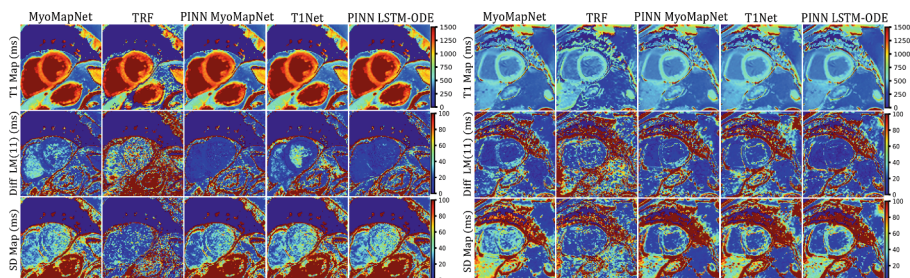


Fig. 3. Estimated mean values of T_1 and SD maps. Left: native pre-contrast sequence. Right: post-contrast sequence.

5 Conclusions

We proposed a novel deep learning framework to accelerate MOLLI-based cardiac T_1 mapping by integrating PINNs. Our experimental results show that the proposed PINN LSTM-ODE model outperforms alternative methods in both accuracy and precision, particularly with limited LL samples. Furthermore, we demonstrate that physics-informed formulations improve the performance of models that traditionally rely on direct data-driven T_1 priors, such as MyoMapNet. This highlights the benefits of embedding physics-based constraints into model training. The proposed framework enables accurate T_1 estimation and efficient FoV mapping from just 3 to 5 heartbeats, matching the performance of the conventional 17-heartbeat MOLLI sequence while significantly reducing acquisition time.

Acknowledgments. We gratefully acknowledge the Dutch Research Council (NWO) and Amazon Science for financial and computing support.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Baytas, I.M., Xiao, C., Zhang, X., Wang, F., Jain, A.K., Zhou, J.: Patient subtyping via time-aware lstm networks. In: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2017, pp. 65–74. Association for Computing Machinery, New York (2017). <https://doi.org/10.1145/3097983.3097997>
2. Chen, R.T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D.: Neural ordinary differential equations (2019), <https://arxiv.org/abs/1806.07366>
3. Chow, K., Flewitt, J.A., Green, J.D., Pagano, J.J., Friedrich, M.G., Thompson, R.B.: Saturation recovery single-shot acquisition (SASHA) for myocardial t(1) mapping. *Magn. Reson. Med.* **71**(6), 2082–2095 (2014)
4. Guo, R., et al.: Accelerated cardiac t1 mapping in four heartbeats with inline myomapnet: a deep learning-based t1 estimation approach. *J. Cardiovascular Magnet. Resonance* **24**(1), 6 (2022). <https://doi.org/10.1186/s12968-021-00834-0>, <https://doi.org/10.1186/s12968-021-00834-0>
5. Haaf, P., Garg, P., Messroghli, D.R., Broadbent, D.A., Greenwood, J.P., Plein, S.: Cardiac t1 mapping and extracellular volume (ecv) in clinical practice: a comprehensive review. *J. Cardiovascular Magnetic Resonance* **18**(1), 89 (2016). <https://doi.org/10.1186/s12968-016-0308-4>, <https://www.sciencedirect.com/science/article/pii/S109766472301030X>
6. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997). <https://doi.org/10.1162/neco.1997.9.8.1735>, <https://doi.org/10.1162/neco.1997.9.8.1735>
7. Kellman, P., Arai, A.E., Xue, H.: T1 and extracellular volume mapping in the heart: estimation of error maps and the influence of noise on precision. *J. Cardiovasc. Magn. Reson.* **15**(1), 1–12 (2013)
8. Kellman, P., Hansen, M.S.: T1-mapping in the heart: accuracy and precision. *J. Cardiovasc. Magn. Reson.* **16**, 1–20 (2014)
9. Kidger, P., Morrill, J., Foster, J., Lyons, T.: Neural controlled differential equations for irregular time series (2020). <https://arxiv.org/abs/2005.08926>
10. Konstantinov, M., Petkov, P., Christov, N.: Synthesis of linear systems with desired equivalent form. *J. Comput. Appl. Math.* **6**(1), 27–35 (1980). [https://doi.org/10.1016/0771-050X\(80\)90014-5](https://doi.org/10.1016/0771-050X(80)90014-5), <https://www.sciencedirect.com/science/article/pii/0771050X80900145>
11. Le, J.V., Mendes, J.K., McKibben, N., Wilson, B.D., Ibrahim, M., DiBella, E.V.R., Adluru, G.: Accelerated cardiac T1 mapping with recurrent networks and cyclic, model-based loss. *Med. Phys.* **49**(11), 6986–7000 (2022)
12. Lechner, M., Hasani, R.: Learning long-term dependencies in irregularly-sampled time series (2020). <https://arxiv.org/abs/2006.04418>
13. Messroghli, D.R., Radjenovic, A., Kozerke, S., Higgins, D.M., Sivananthan, M.U., Ridgway, J.P.: Modified look-locker inversion recovery (molli) for high-resolution t1 mapping of the heart. *Magnet. Resonance Med.* **52**(1), 141–146

- (2004). <https://doi.org/10.1002/mrm.20110>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.20110>
14. Piechnik, S.K., Ferreira, V.M., Dall'Armellina, E., Cochlin, L.E., Greiser, A., Neubauer, S., Robson, M.D.: Shortened modified look-locker inversion recovery (shmolti) for clinical myocardial t1-mapping at 1.5 and 3 t within a 9 heartbeat breathhold. *Journal of Cardiovascular Magnetic Resonance* **12**(1), 69 (Nov 2010). <https://doi.org/10.1186/1532-429X-12-69>, <https://doi.org/10.1186/1532-429X-12-69>
 15. Raissi, M., Perdikaris, P., Karniadakis, G.: Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019). <https://doi.org/10.1016/j.jcp.2018.10.045>, <https://www.sciencedirect.com/science/article/pii/S0021999118307125>
 16. Roujol, S., Weingärtner, S., Foppa, M., Chow, K., Kawaji, K., Ngo, L.H., Kellman, P., Manning, W.J., Thompson, R.B., Nezafat, R.: Accuracy, precision, and reproducibility of four T1 mapping sequences: a head-to-head comparison of MOLLI, ShMOLLI, SASHA, and SAPPHERE. *Radiology* **272**(3), 683–689 (2014)
 17. Schirmer, M., Eltayeb, M., Lessmann, S., Rudolph, M.: Modeling irregular time series with continuous recurrent units (2022). <https://arxiv.org/abs/2111.11344>
 18. Weingärtner, S., et al.: Myocardial t1-mapping at 3t using saturation-recovery: reference values, precision and comparison with molli. *J. Cardiovascular Magnet. Resonance* **18**(1), 84 (2016). <https://doi.org/10.1186/s12968-016-0302-x>, <https://www.sciencedirect.com/science/article/pii/S1097664723010256>
 19. Zhang, Y., Zhao, Y., Huang, L., Xia, L., Tao, Q.: Deep-learning-based groupwise registration for motion correction of cardiac t 1 mapping. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 586–596. Springer (2024)
 20. Zhang, Y., Zhao, Y., Shao, Y., Capitão, N.M.F., van den Bogert, F., Tao, Q.: Torch\$t_1\$: GPU-accelerated cardiac \$t_1\$ mapping with deep learning framework. In: *Medical Imaging with Deep Learning* (2024). <https://openreview.net/forum?id=8ONihf0Fh9>