Bridging Physics and Deep Learning: An Autoencoder-Based Approach for Non-Invasive Estimation of Prostate Tissue Composition

Batuhan Gundogdu, Aritrick Chatterjee, Serene Bose, Gregory S. Karczmar and Aytekin Oto

Abstract:

Physics model-based approaches using Magnetic Resonance Imaging (MRI) have been developed to provide a non-invasive tool to estimate tissue composition by fitting the notoriously ill-conditioned sum-of-decaying-exponentials functions. However, even modest levels of noise can negatively impact the accuracy of such models. This study introduces a new paradigm that integrates the strengths of physics model-based and deep learning-based methods, while overcoming their respective weaknesses. We reformulated the MRI-based tissue composition estimation using a physics-informed autoencoder (PIA), transforming the problem from a least-squares fit into a robust deep learning framework. The PIA model consists of a trainable multi-head neural network encoder and a fixed, physics-informed decoder. Instead of treating the parameters as unknowns in a set of equations, the model considers them as latent variables within an autoencoder. Training is carried out in an unsupervised fashion, focusing on minimizing the squared error between the input and the "physics-informed output". The encoder is pre-trained with inputs of simulated MRI signals with added noise, similar to the denoising autoencoders, which in turn provided noise-robustness for in-vivo operation. Without a need for further fine-tuning for each new patient, the encoder's latent activations are utilized to derive prostate tissue composition estimates, such as volume fractions for each compartment along with their diffusivity and T2. PIA's estimates of tissue parameters are juxtaposed with the least-squares solution through both Monte Carlo simulations and the quantitative histology of in-vivo prostate scans. In-vivo evaluations, performed on 21 pathology-proven prostate cancer patients, have been validated with quantitative histology-based true tissue compositions. Both Monte Carlo and in-vivo evaluations reveal that PIA significantly outperforms the conventional least-squares solution on Pearson correlation and mean absolute error metrics. Notably, PIA was able to predict the diffusivity and T2 of each tissue compartment more accurately, whereas the least-squares method tends to pin these values to the extremities. More importantly, in in-vivo experiments PIA demonstrated a strong correlation with histology-based true tissue composition ($r=0.891$, $p<0.001$), outperforming the performance of the least-squares method ($r=0.820$, $p<0.001$), while providing a remarkable speed improvement of about 30,000x over the least-squares method. These findings position PIA as a promising non-invasive instrument for predicting prostate tissue composition from MRI, providing important new diagnostic markers, and laying the groundwork for a potential new quantitative MRI method.

1. Introduction:

1.1. Background: Prostate cancer represents a major health issue for men in the United States. It is the most common noncutaneous cancer and second leading cause of death among men in the United States [1]. Magnetic resonance imaging (MRI) is used for detection, characterization, and staging of prostate cancer as well as other pathological conditions as benign prostatic hyperplasia (BPH) [2]. Although MRI provides a tool for detecting the lesions, a large number of cancers and BPH are missed or overdiagnosed due to their similarity to the normal gland and benign conditions [3]. More importantly, MRI-based diagnosis is qualitative, and therefore the accuracy heavily depends on the experience of the radiologist. Therefore, histologic changes in the tissue composition, **based on pathology,** remain the reference standard for cancer detection. Specifically, estimation of volume fractions of the three compartments *epithelium*, *stroma*, and *lumen* provides the most robust means to predict cancer, and or other pathological conditions such as BPH.

In order to estimate the tissue composition non-invasively, physics-based models utilize sum-of-decaying-exponentials-form signal decay model to MRI signals using least squares fit [4]. However, these functions are notoriously ill-conditioned for least-squares solutions and the accuracy of such model fits is affected by even modest levels of noise. As an alternative to model-based approaches, deep learning-based cancer detection models require large amounts of annotated data, which is generally impractical [5]. Moreover, the performance of such models is questionable in new imaging environments with scanners from different vendors or images with different imaging parameters.

1.2. Purpose: The purpose of this study is to utilize unsupervised deep learning framework to improve physics model-based tissue composition estimation from MRI, thereby providing a noise-robust and non-invasive "virtual pathology" methodology.

2. Methods:

This study introduces a new paradigm that integrates the strengths of physics model-based and deep learning-based methods, while overcoming their respective weaknesses, including the need for large amounts of annotated data. To estimate the volume fractions (*v*) of epithelium, stroma, and lumen, as well as their diffusivities (*D*) and T2 relaxation times, used in physics model-based solutions, we propose a novel approach called *physics-informed autoencoders (PIA)*[6]. **This model treats these parameters not as unknowns of an equation set but as latent variables of an autoencoder.**

2.1. Structure of Physics-Informed Autoencoder (PIA): We have reimagined the tissue composition estimation using PIA, transforming the problem from a least-squares fit into a robust deep learning framework. The PIA model consists of a trainable multi-head neural network *encoder* and a fixed, physics-informed *decoder*.

2.1.1 Encoder: The encoder is a critical part of the PIA, comprising a multi-head deep neural network. The design of this neural network is focused on yielding accurate parameter estimates for each compartment, including volume fractions, T2, and ADC. The intricate architecture ensures that the underlying physics of the physics-model is integrated into the learning process, thereby creating a well-informed and robust encoding mechanism. The input to the encoder is a multidimensional MRI signal measurements for each voxel. The first few layers of the encoder extract an embedding *(e)* from these measurements. This embedding is then fed to its multiple heads, each of which are deep neural networks, estimating a different parameter set, namely *v*, *D* and *T2* for the 3 compartments: epithelium, stroma and lumen.

$$\begin{aligned} \mathrm{e} &= \mathrm{encoder}(\mathrm{S}) \\ \mathrm{v} &= \mathrm{softmax}\big(volume\_model(e)\big) \\ \mathrm{D_n} &= D\_model(e) \\ \mathrm{T2_n} &= T2\_model(e) \end{aligned} \qquad (1)$$

2.1.2 Decoder: The decoder is uniquely set as nontrainable within the PIA architecture. Instead of learning from data, it employs a predefined analytical function that generates theoretical MRI signal values at the output. This aspect preserves the direct relationship between the encoded parameters and their corresponding theoretical values, ensuring that the deep learning model remains rooted in established physical laws.

$$\hat{\mathrm{S}} = \sum_{n=1}^{n=3} \mathrm{v_n} \times \exp\left(-D_n \times b - \frac{TE}{T2_n}\right), \qquad (2)$$

where gradient strength (b) and echo time (TE) are the imaging parameters used in acquiring the signal S. The overall flow of the model is presented in Figure 1.

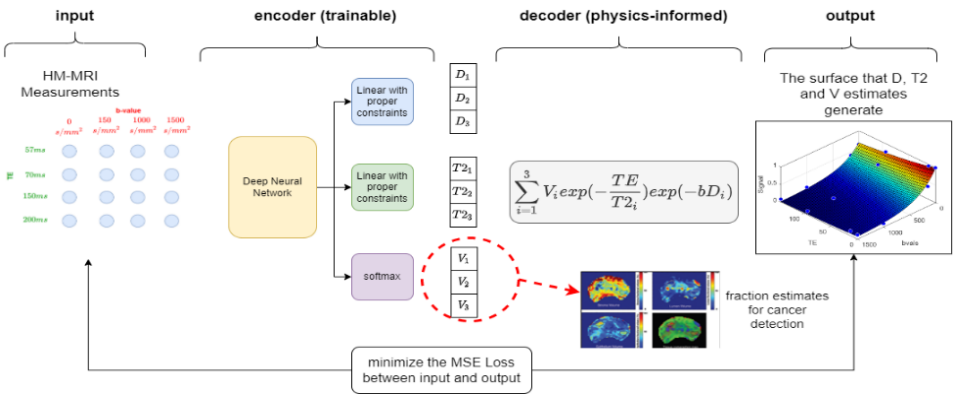

Figure 1. Physics informed autoencoder architecture.

2.2. Training Methodology and Implementation Details

Our model operates in an unsupervised manner. Training the PIA involves minimizing the squared error between its input *S* and the so-called "physics-informed output". The optimization is carried out using the adaptive moment estimation (adam) optimizer, where appropriate hyperparameters are selected after exhaustive experimentation. The entire framework is implemented using Pytorch 3.9 and is trained on a workstation with 32Gb GPU memory. The data preprocessing includes normalization and any required transformations to ensure that the data is suitable for training the neural network.

3. Results:

Evaluations were conducted via both Monte Carlo simulations under various noise conditions as well as in-vivo scans of patients with biopsy-verified prostate cancer.

3.1. Simulations:

For the Monte Carlo simulations, MRI signals are generated by using the multi-compartment signal rule in Equation (2) and randomly sampling from physically plausible epithelium, lumen and stroma tissue parameters. PIA was trained on millions of such synthetic MRI signals and by introducing additive Gaussian noise in the input. This application was key to obtain noise-robustness of PIA, as in denoising autoencoders. Upon training PIA, a separate set of 2500 synthetically generated and noisy MRI signals are used to estimate the underlying tissue parameters using both PIA and the least-squares-based model. It should be noted that the PIA uses only the encoder for this estimation, and the least-squares model in [4] solves the tissue parameters for each new voxel using the Equation (2). Pearson correlation coefficients and mean absolute errors between the estimates of each model against the true tissue parameters are calculated as performance metrics. The Pearson correlation coefficients and mean absolute error metrics for PIA and the least squares-based baseline (marked as HMMRI – Hybrid multidimensional MRI) are given in Figure 2. This scatter plot on the right figure, demonstrating the true vs estimated parameters for both approaches demonstrate the superiority of PIA in predicting the diffusivity and T2 of the compartments. These parameters, which are on the exponent on the Equation (2), gets pinned to the boundary conditions with the least squares solution, whereas PIA yields a better performance for estimation of these parameters.

Pearson's Correlation Coefficient (higher is better)

	Epithelium		Stroma		Lumen	
	HMMRI	PIA	HMMRI	PIA	HMMRI	PIA
Volume	0.61	**0.81**	0.53	**0.74**	0.91	**0.97**
ADC	0.19	**0.4**	0.21	**0.43**	0.07	**0.12**
T2	0.4	**0.76**	0.23	**0.52**	0.03	**0.18**

Mean Absolute Error (lower is better)

	Epithelium		Stroma		Lumen	
	HMMRI	PIA	HMMRI	PIA	HMMRI	PIA
Volume	0.13	**0.08**	0.16	**0.1**	0.06	**0.03**
ADC	0.16	**0.09**	0.35	**0.22**	0.14	**0.08**
T2	14.3	**7**	21.9	**12.1**	236.9	**134.5**

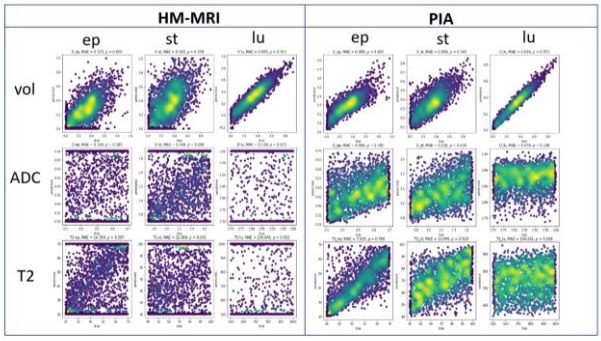

Figure 2: Pearson Correlation Coefficient and mean absolute error performances for each of the nine tissue parameters (left). PIA performs better than HMMRI on all metrics. The scatter plots on the right demonstrate the superiority of PIA against HMMRI, especially on ADC and T2 metrics, where HMMRI pins these values to the extremities. x axes show the true value and y axes show the models' prediction.

3.2. In-vivo Experiments:

In-vivo evaluations were performed on 21 pathology-proven prostate cancer patients, using 43 lesions from cancers as well as healthy regions from peripheral and transition zones. The PIA estimates were validated with quantitative histology-based true tissue compositions. We also ran the same experiment with the least squares solution-based Hybrid Multidimensional MRI (HM-MRI) and compared HM-MRI's results with histology-based true tissue compositions. PIA yielded a significant correlation with the histology with r=0.911 for epithelium, r=0.851 for stroma and r=0.912 for lumen. HM-MRI also gives a good correlation with histology with r=0.908 for epithelium, r=0.753 for stroma and r=0.798 for lumen, which is not only inferior to PIA but also about 30,000 times slower. As per the estimates of diffusivity and T2, there's no straightforward way to validate the accuracy of the both methods' performance with in-vivo measurements. This could have been possible with microscopic MRI, but then it would have to have been ex-vivo. Nevertheless, as the Monte Carlo simulations also suggest, PIA yields more reasonable values for the diffusivities of stroma and epithelium. (ADC of lumen is almost always very close to that of water, i.e 3um^2/ms). While HM-MRI yields diffusivity values that are pinned to upper or lower bounds of the feasible values, PIA exhibits reduced diffusivities on cancer regions and a correlated diffusivity values for epithelium and stroma. The prostate tissue compartment fits and the diffusivity estimate of stroma and epithelium are demonstrated in Figure 3.

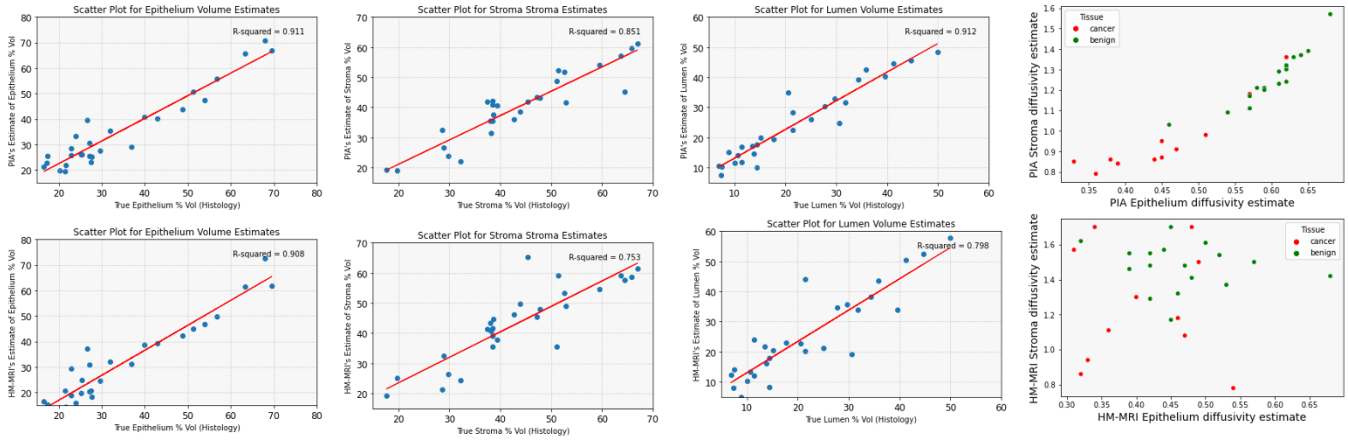

Figure 3: Top row shows the PIA and the bottom row show the HM-MRI estimates. From left to right scatter plots of estimated epithelium, stroma and lumen volume fractions compared to quantitative histology-based ground truths. The right image shows the epithelium vs stroma diffusivities, colored red and green to denote cancer and benign, respectively.

Figure 4 shows one example from the in-vivo prostate scans along with histopathology of the corresponding slide, PIA-based tissue composition estimates and the predicted cancer, based on epithelium and lumen volumes. Specifically, a region is called cancer if volume fraction estimate of epithelium is greater than 40% and volume fraction estimate of lumen is less than 20%.

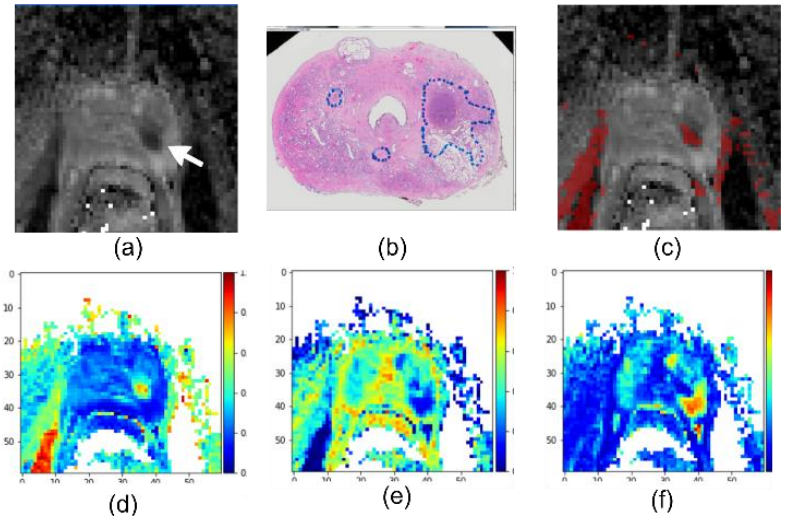

Figure 3: (a) ADC map of the cancer slice, arrow points at the cancer. (b) Histology of the corresponding slice. (c) Cancer map derived from volume fraction estimates. Volume fraction estimate map for (d) epithelium, (e)stroma and (f) lumen.

5. Conclusions:

These findings position PIA as a promising non-invasive instrument for predicting prostate tissue composition from MRI, providing important new diagnostic markers, and laying the groundwork for a potential new quantitative MRI method.

6. Acknowledgements:

PIA was first introduced by this lab in the AAPM 2023 Annual Conference. This work includes the histological validation with in-vivo patient data that the AAPM submission did not have.

7. References

[1] R. L. Siegel, K. D. Miller, N. S. Wagle, and A. Jemal, "Cancer statistics, 2023," *CA A Cancer J Clinicians*, vol. 73, no. 1, pp. 17–48, Jan. 2023, doi: 10.3322/caac.21763.

[2] T. Barrett, B. Turkbey, and P. L. Choyke, "PI-RADS version 2: what you need to know," *Clinical Radiology*, vol. 70, no. 11, pp. 1165–1176, Nov. 2015, doi: 10.1016/j.crad.2015.06.093.

[3] A. C. Westphalen *et al.*, "Variability of the Positive Predictive Value of PI-RADS for Prostate MRI across 26 Centers: Experience of the Society of Abdominal Radiology Prostate Cancer Disease-focused Panel," *Radiology*, vol. 296, no. 1, pp. 76–84, Jul. 2020, doi: 10.1148/radiol.2020190646.

[4] A. Chatterjee *et al.*, "Validation of Prostate Tissue Composition by Using Hybrid Multidimensional MRI: Correlation with Histologic Findings," *Radiology*, vol. 302, no. 2, pp. 368–377, Jan. 2022, doi: 10.1148/radiol.2021204459.

[5] H. Li, C. H. Lee, D. Chia, Z. Lin, W. Huang, and C. H. Tan, "Machine Learning in Prostate MRI for Prostate Cancer: Current Status and Future Opportunities," *Diagnostics*, vol. 12, no. 2, p. 289, Jan. 2022, doi: 10.3390/diagnostics12020289.

[6] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nat Rev Phys*, vol. 3, no. 6, pp. 422–440, May 2021, doi: 10.1038/s42254-021-00314-5.