



# Multi-parameter molecular MRI quantification using physics-informed self-supervised learning



Alex Finkelstein <sup>1</sup>, Nikita Vladimirov <sup>1</sup>, Moritz Zaiss <sup>2,3</sup> & Or Perlman <sup>1,4</sup>

Biophysical model fitting plays a key role in obtaining quantitative parameters from physiological signals and images. However, the model complexity for molecular magnetic resonance imaging (MRI) often translates into excessive computation time, which makes clinical use impractical. Here, we present a generic computational approach for solving the parameter extraction inverse problem posed by ordinary differential equation (ODE) modeling coupled with experimental measurement of the system dynamics. This is achieved by formulating a numerical ODE solver to function as a step-wise analytical one, thereby making it compatible with automatic differentiation-based optimization. This enables efficient gradient-based model fitting, and provides a new approach to parameter quantification based on self-supervised learning from a single data observation. The neural-network-based train-by-fit pipeline was used to quantify semisolid magnetization transfer (MT) and chemical exchange saturation transfer (CEST) amide proton exchange parameters in the human brain, in an *in-vivo* molecular MRI study ( $n = 4$ ). The entire pipeline of the first whole brain quantification was completed in  $18.3 \pm 8.3$  minutes. Reusing the single-subject-trained network for inference in new subjects took  $1.0 \pm 0.2$  s, to provide results in agreement with literature values and scan-specific fit results.

Magnetic resonance imaging (MRI) plays a central role in clinical diagnosis and neuroscience. This modality is highly versatile and can be selectively programmed to generate a large number of image contrasts<sup>1</sup>, each sensitive to certain biophysical parameters of the tissue. In recent years, there has been extensive research into developing quantitative MRI (qMRI) methods that can provide reproducible measurements of magnetic tissue properties (such as:  $T_1$ ,  $T_2$ , and  $T_2^*$ ), while being agnostic to the scan site and the exact acquisition protocol used<sup>2</sup>. Classical qMRI quantifies each biophysical property separately<sup>3</sup>, using repeated acquisition and gradual variation of a single acquisition parameter under steady state conditions. This is followed by fitting the model to an analytical solution of magnetization vector dynamics<sup>4</sup>.

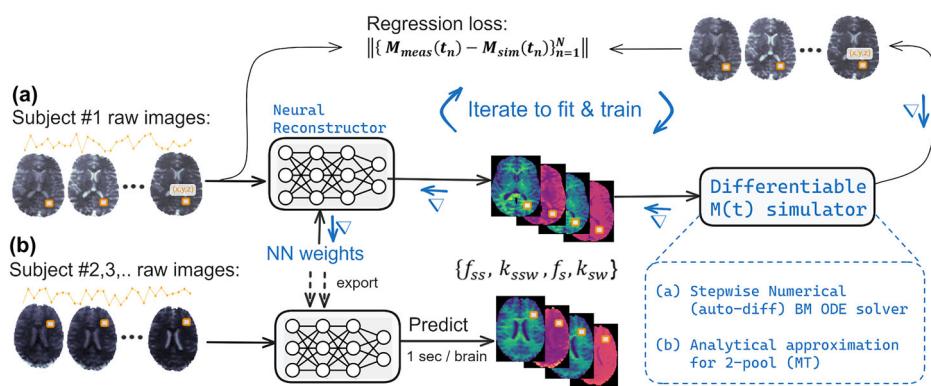
The exceedingly long acquisition times associated with the classical quantification pipeline have motivated the development of magnetic resonance fingerprinting (MRF)<sup>5</sup>, an alternative paradigm for the joint extraction of multiple tissue parameter maps from a single pseudorandom pulse sequence. Since MRF data are acquired under non-steady state conditions<sup>6</sup>, the corresponding magnetization vector can only be resolved numerically.

This comes at the expense of the complexity of the inverse problem, namely finding tissue parameters that best reconstruct the signal according to the forward model of spin dynamics. Since model fitting under these conditions takes an impractically long time<sup>7</sup>, MRF is commonly solved by dictionary matching, where a large number of simulated signal trajectories are compared to experimentally measured data<sup>8</sup>. Unfortunately, the size of the dictionary scales exponentially with the number of parameters (the “curse of dimensionality”<sup>9</sup>), which rapidly escalates the compute and memory demands of both generation and subsequent use of the dictionary for pattern matching-based inference.

Recently, various deep learning (DL)-based methods have been developed for replacing the lengthy dictionary matching with neural-network (NN)-based inference<sup>10–13</sup>. While this approach greatly reduces the parameter quantification time, networks still need to be trained using a comprehensive dictionary of synthetic signals. Since dictionary generation may take days<sup>12</sup>, it constitutes an obvious bottleneck for routine use of MRF, and reduces the possibilities for addressing a wide variety of clinical scenarios. Even with a faster generation, the transfer of synthetic

<sup>1</sup>Department of Biomedical Engineering, Tel Aviv University, Tel Aviv, Israel. <sup>2</sup>Institute of Neuroradiology, University Hospital Erlangen, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Erlangen, Germany. <sup>3</sup>Department of Artificial Intelligence in Biomedical Engineering, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Erlangen, Germany. <sup>4</sup>Sagol School of Neuroscience, Tel Aviv University, Tel Aviv, Israel. e-mail: [orperlman@tauex.tau.ac.il](mailto:orperlman@tauex.tau.ac.il)

**Fig. 1 | Schematic representation of the core neural Bloch McConnell fitting (NBMF) pipeline.** A quantitative parameter reconstructor parameterized as a multi-layer perceptron (MLP) and a differentiable Bloch-McConnell simulator are serially connected into a single computational graph. Single-subject Magnetic Resonance Fingerprinting (MRF) data serves both as the input and as the regression target for the reconstructor-simulator circuit. The network convergence (a) provides the fitted exchange parameter maps for the examined subject as well as a trained NN reconstructor; the latter can be used to extract parameter maps for new subjects within seconds (b). The simulator can be realized using the exact numerical Bloch McConnell ODE solver or using analytical approximations when available (e.g., for 2-pool semisolid Magnetization Transfer (MT) quantification<sup>33</sup>). While not shown in the diagram, auxiliary per-voxel data such as  $T_1$ ,  $T_2$ ,  $B_0$ , and  $B_1$  maps can be added as input to the neural reconstructor and the simulator. Furthermore, the pipeline main block can be serially repeated so that estimated semisolid MT volume fraction ( $f_{ss}$ ) and proton exchange rate ( $k_{ssw}$ ) maps inferred at the first stage are joined to the raw data used in a second reconstructor aimed to quantify the amide proton exchange parameters ( $f_s$ ,  $k_{sw}$ ).



data-trained NN to experimental data raises concerns about biased estimates.

The complexity and time constraints associated with the MRF pipeline are drastically exacerbated for molecular imaging applications that involve a plurality of proton pools, such as chemical exchange saturation transfer (CEST) MRI<sup>14</sup>. While CEST has demonstrated great potential for dozens of biomedical applications<sup>15–21</sup>, some on the verge of entering clinical practice<sup>22</sup>, the inherently large number of tissue properties greatly complicate analysis<sup>23</sup>. This has prompted considerable efforts to transition from CEST-weighted imaging to fully quantitative mapping of proton exchange parameters<sup>24–28</sup>. Early CEST quantification used the fitting of the classical numerical model (based on the underlying Bloch-McConnell equations) after irradiation at various saturation pulse powers ( $B_1$ )<sup>27</sup>. However, applying this approach in a pixelwise manner in-vivo is unrealistic because both the acquisition and reconstruction steps may require several hours. Later, faster approaches, such as quantification of the exchange by varying saturation power/time and Omega-plots<sup>25,29–31</sup> still rely on steady-state (or close to steady state) conditions, and approximate analytical expressions of the signal as a function of the tissue parameters<sup>32,33</sup>. Unfortunately, a closed-form analytical solution does not exist for most practical clinical CEST protocols, which utilize a train of off-resonant radiofrequency (RF) pulses saturating multiple interacting proton pools. Similarly to the quantification of water  $T_1$  and  $T_2$ , incorporating the concepts of MRF into CEST studies provided new quantification capabilities<sup>28,34–37</sup> and subsequent biological insights, for example, in the detection of apoptosis after oncolytic virotherapy<sup>12</sup>. However, in order to further push the boundaries of CEST MRF research and expedite its progress, the long dictionary generation associated with each new application needs to be replaced by a rapid and flexible approach that adequately models multiple proton pools under saturation pulse trains.

Here, we describe a physics-based DL framework for rapid model fitting of the human brain tissue proton spin properties. While this approach is applicable for quantifying a variety of MRI parameters, we focus on a challenging CEST imaging scenario, involving multiple proton pools, a saturation pulse train, and non-steady-state MRF acquisition. The computational pipeline (Fig. 1) combines a spin physics simulator and a NN-based quantitative parameter reconstructor in a fully auto-differentiable manner<sup>38</sup>. Our system effectively solves and

inverts the Bloch-McConnell ordinary differential equations (ODEs), which govern the multi-pool exchange, saturation, and relaxation dynamics of molecular MRI. Hence, we refer to this approach as “neural Bloch McConnell fitting” (NBMF). Importantly, the network can be trained in a self-supervised manner, directly on the single-subject data of interest (inspired by related work on test-time-<sup>39</sup>, internal-<sup>40</sup>, and zero-shot-<sup>40–42</sup> learning). This circumvents the need for prior curation of a large training dataset, which is often inaccessible, especially for molecular MRI.

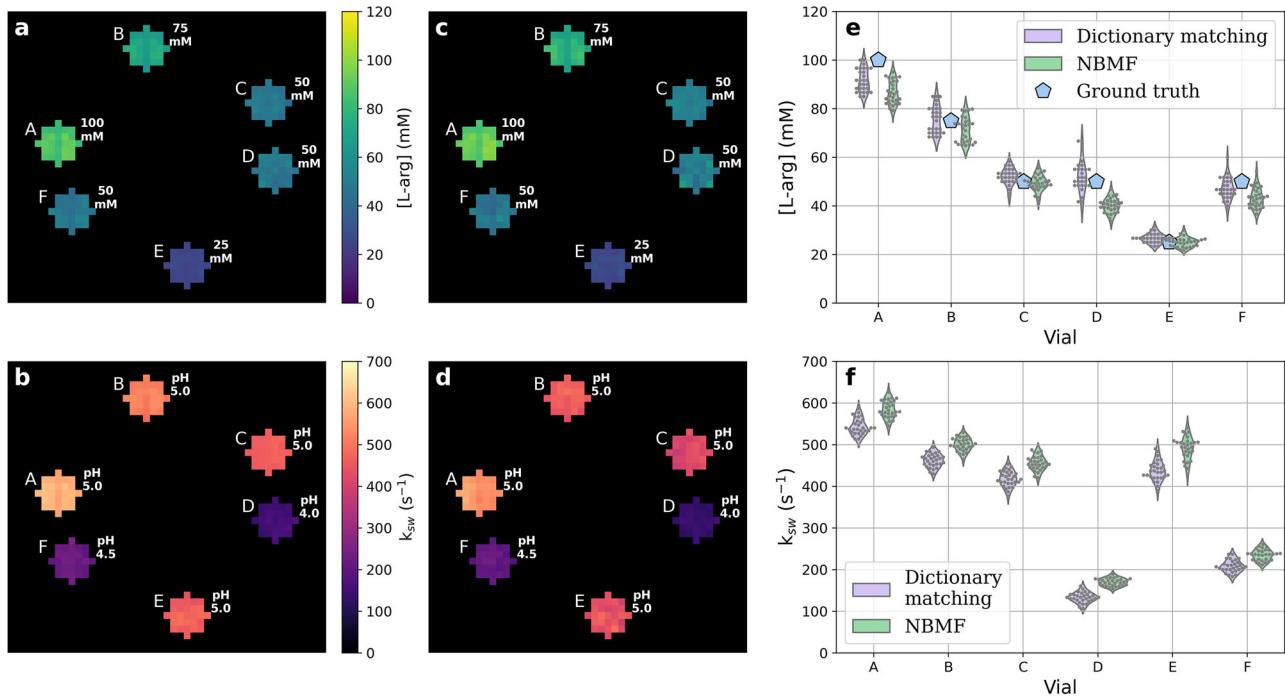
## Results and discussion

### In-vitro CEST quantification

A phantom composed of six vials with different combinations of L-arginine concentration and pH was assembled and scanned in a 3T clinical scanner (Prisma, Siemens Healthineers, Germany) using a previously published non-steady-state rapid CEST protocol<sup>12,43</sup>. Good agreement was obtained between the NBMF-estimated and known L-arginine concentrations (Fig. 2a, c, e): Pearson’s  $r = 0.986$ ,  $p = 3.0 \times 10^{-4}$ , root mean square error (RMSE) = 8.4 mM, mean absolute percentage error (MAPE) = 10.8%. The NBMF-reconstructed proton exchange rates were in good agreement with the corresponding values estimated by traditional MRF dictionary-matching (Fig. 2b, d, f): Pearson’s  $r = 0.999$ ,  $p = 1.6 \times 10^{-6}$ , RMSE = 41.0 s<sup>-1</sup>, MAPE = 13.2%. The pH dependence of the NBMF reconstructed exchange rate (Fig. 2b) was a good fit for a base-catalyzed proton exchange model ( $R^2 = 0.94$ ,  $p = 1.4 \times 10^{-3}$ ), as predicted by theory<sup>44</sup>. An additional in-depth comparison between traditional dictionary matching and NBMF is available in Supplementary Table 1, and Supplementary Figs. 1, 2.

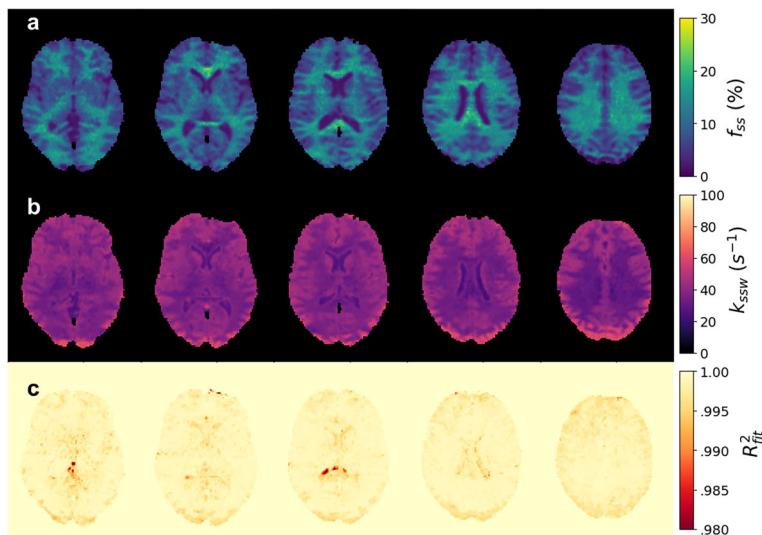
### Quantifying the semisolid-MT and CEST proton exchange parameters in the human brain

The NBMF pipeline used in-vitro was modified for semisolid-MT and amide proton exchange parameter mapping in the human brain. Two 3D and rapid acquisition pulse sequences were applied, as described previously<sup>12,43</sup>. The first sequence varied the saturation pulse frequency offset between 6 and 14 ppm (designed to encode semisolid-MT information), whereas the second sequence fixed it at 3.5 ppm (for amide proton parameter encoding). In both cases, a total of 31 raw information encoding

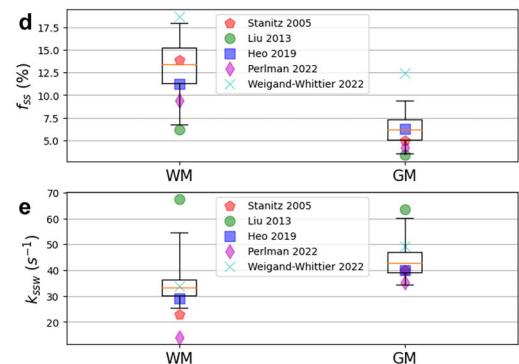


**Fig. 2 | In-vitro study.** L-arginine samples were imaged using a pulsed Chemical Exchange Saturation Transfer Magnetic Resonance Fingerprinting (CEST-MRF) protocol in a 3T clinical scanner. The neural Bloch McConnell fitting (NBMF)-based L-arginine concentrations (a) and proton exchange rates (b) were in good agreement with those obtained by dictionary-based pattern matching (c and

d, respectively). The ground truth L-arginine concentrations and pH values are mentioned in the white text next to each vial. The pixelwise distributions are further compared in (e, f). Each point in the swarm plot reflects a single  $1.8\text{ mm} \times 1.8\text{ mm} \times 5.4\text{ mm}$  voxel.



**Fig. 3 | In-vivo study: semisolid magnetization transfer (MT) imaging.** Results of neural Bloch-McConnell fitting (NBMF)-based quantification of the MT-related tissue parameters in the human brain scanned with a pulsed MT Magnetic Resonance Fingerprinting (MT-MRF) protocol are presented. Representative reconstructed parameter maps of the semisolid-MT proton volume fraction (a) and proton exchange rate (b), alongside a fidelity estimation (c) of the data-model agreement, computed as  $R^2 = 1 - \text{NMSE}$  (normalized mean square error).

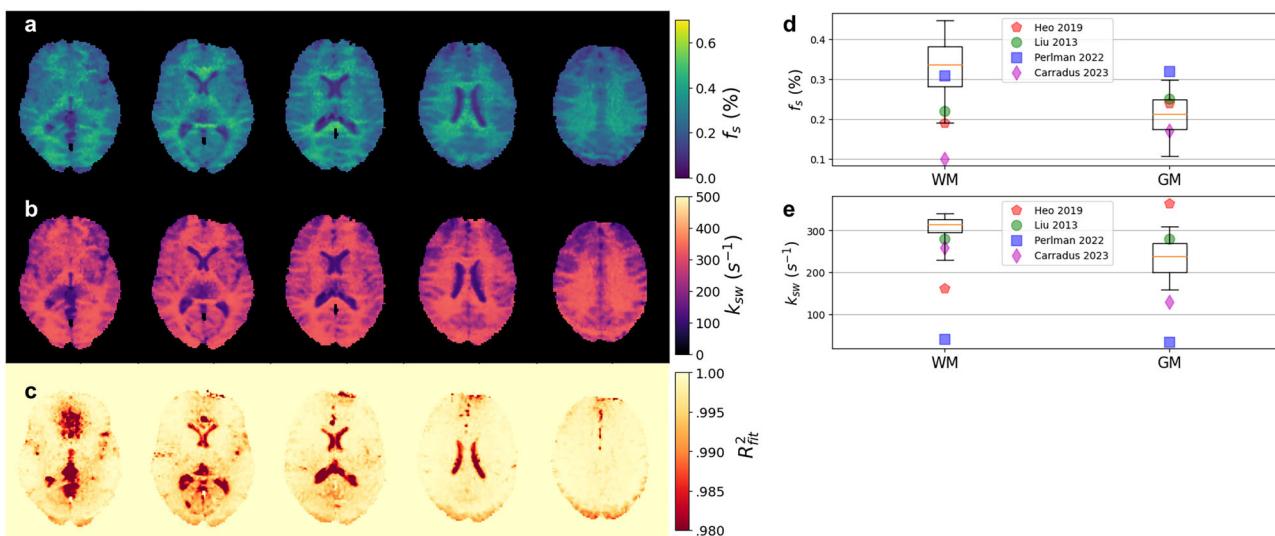


**d, e** Statistical analysis of the resulting proton exchange parameter values across the brain white matter and gray matter (WM/GM) regions of interest (box-plots,  $n=47,442/64,611$  voxels, respectively), compared to literature (colored markers)<sup>12,43,45,90,91</sup>. In the boxplots, the central horizontal lines represent median values, box size represents the two central (2nd, 3rd) quartiles, the whiskers represent the 90 central percentiles, and outliers are omitted.

images were generated, with the saturation pulse power randomly varied between 0 and  $4\text{ }\mu\text{T}$ <sup>12,44</sup>. Water relaxation ( $T_1$  and  $T_2$ ) and field maps ( $B_1$  and  $B_0$ ) were acquired separately and used as an extra input to the neural reconstructor described in Fig. 1 in order to avoid water-pool- and inhomogeneity-associated biases, respectively<sup>12</sup>. A two-step NBMF (see

“Methods” section) was used to fit the semisolid-MT and amide proton exchange parameters to the raw data.

Quantitative semisolid-MT and amide proton exchange parameter maps derived from a representative healthy volunteer are presented in Fig. 3 and Fig. 4, respectively. The resulting proton volume fractions and



**Fig. 4 | In-vivo study: amide proton transfer (APT) imaging.** Results of neural Bloch-McConnell fitting (NBMF)-based quantification of the APT-related tissue parameters in a human brain scanned with a pulsed Chemical Exchange Saturation Transfer Magnetic Resonance Fingerprinting (CEST-MRF) protocol. Representative reconstructed parameter maps of the amide proton volume fraction (a) and proton exchange rate (b), alongside a fidelity estimation (c) of the data-model agreement, computed as  $R^2 = 1 - \text{NMSE}$  (normalized mean square error).

**d, e** Statistical analysis of the resulting proton exchange parameter values across the brain white matter and gray matter (WM/GM) regions of interest (box-plots,  $n = 47442/64611$  voxels, respectively), compared to literature (colored markers)<sup>12,45,91,92</sup>. In the boxplots, the central horizontal lines represent median values, box size represents the two central (2nd, 3rd) quartiles, the whiskers represent the 90 central percentiles, and outliers are omitted.

exchange rates were in agreement with the literature (although the large variability in previous reports is noted; see Fig. 3d, e and Fig. 4d, e). The mean values obtained for white/gray matter (WM/GM) were:  $f_{ss} = 13.09 \pm 3.44\%$ ,  $k_{ssw} = 34.7 \pm 7.8\text{ s}^{-1}$ ,  $f_s = 0.33 \pm 0.08\%$ ,  $k_{sw} = 305.1 \pm 34.0\text{ s}^{-1}$  for white matter and  $f_{ss} = 6.28 \pm 1.88\%$ ,  $k_{ssw} = 44.2 \pm 7.5\text{ s}^{-1}$ ,  $f_s = 0.21 \pm 0.06\%$ ,  $k_{sw} = 235.9 \pm 46.0\text{ s}^{-1}$  for gray matter.

The joint fit and training of the NBMF produced a neural reconstructor, optimized on a single subject. We then re-applied the trained reconstructors to additional subjects in a fast inference mode. A representative example comparing the parameter maps obtained from single-subject NBMF with those obtained by a rapid reconstructor reuse is shown in Fig. 5. The resulting agreement metrics (Fig. 5e) were as follows; NRMSE:  $7 \pm 1\%$ ,  $12 \pm 3\%$ ,  $7 \pm 1\%$ , and  $18 \pm 1\%$ ; Intraclass correlation coefficient ICC(2,1):  $0.87 \pm 0.03$ ,  $0.82 \pm 0.04$ ,  $0.86 \pm 0.03$ ,  $0.86 \pm 0.03$ ; SSIM:  $0.93 \pm 0.02$ ,  $0.87 \pm 0.07$ ,  $0.94 \pm 0.01$ ,  $0.90 \pm 0.03$ , for the  $f_{ss}$ ,  $k_{ssw}$ ,  $f_s$ , and  $k_{sw}$ , respectively. Additional analysis is provided in Supplementary Fig. 4.

#### Computational complexity, timing, and comparison with alternative approaches

The combined NBMF training and fitting procedure for all four semisolid-MT and amide proton volume fraction and exchange rate parameter maps from the whole brain of a single subject (169K–194K voxels) took  $18.3 \pm 8.3$  min on a standard GPU-equipped (GeForce RTX 3060) desktop workstation, of which, the two-pool quantification of the semisolid MT pool parameters took  $3.0 \pm 0.4$  min. Re-applying the trained reconstructors for whole-brain parameter mapping on a new subject took  $1.0 \pm 0.2$  s. Overall, the complete quantification pipeline takes less than 30 min for NBMF, compared to at least several hours using previously reported implementations of traditional Bloch-Fitting<sup>45</sup>, or dictionary-based preparation and supervised neural network training<sup>12,37</sup> (Table 1).

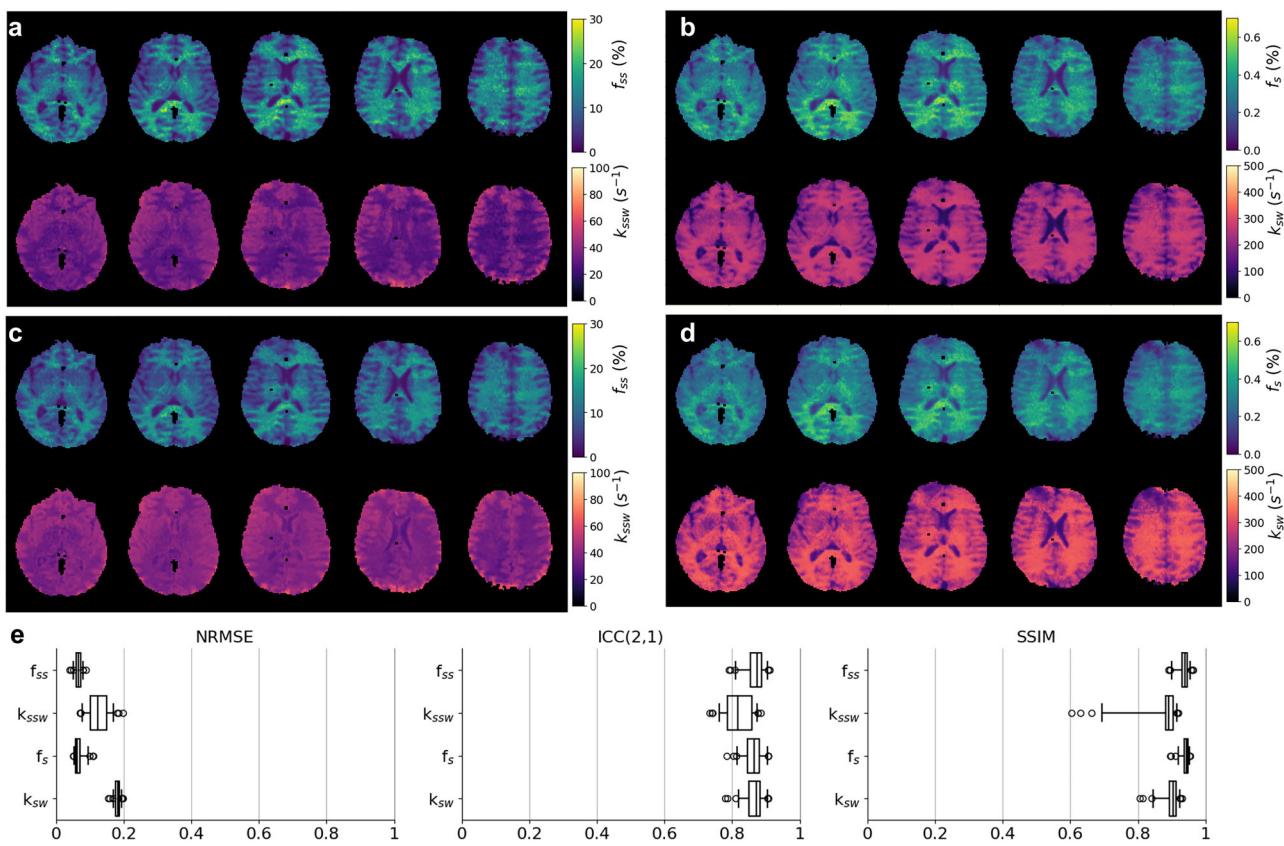
Next, we performed unified benchmarking of dictionary generation, matching, and supervised learning, using the accelerated approach developed as part of this work (GPU-based JAX formulation of the Bloch McConnell numerical solution); see additional implementation details in Supplementary Note 2. Notably, this yielded comparable timing to self-supervision (Table 2), given that a non-cartesian sampling grid is used for dictionary generation. The benefit of nonetheless using the self-supervised

approach compared to supervised training is highlighted in Supplementary Notes 1–3 in the context of consistency with the raw acquired data and per-subject discrepancy minimization. In general, by unlocking rapid direct fitting (via automatic differentiation) and coupling it with self-supervised learning, NBMF constitutes an alternative way to address the ill-posed in-vivo quantification challenge. It contributes to an improved consistency of the quantitative parameter estimates with the raw data given the model, compared to dictionary-based supervised learning (Supplementary Fig. 3). By combining the explicit objective of minimal data-model discrepancy with implicit neural regularization, NBMF created smoother maps with visible contrast, while keeping the data-model agreement close to that achieved by dictionary matching (Supplementary Fig. 5).

#### Automatic differentiation of the Bloch-McConnell (BM) equation solutions

Quantification of semisolid MT/CEST proton exchange parameters under non-steady-state conditions is a notable example of a biophysical estimation where the forward model is perceived as too complex for direct inverse solution via fitting (requiring several days for a single whole brain reconstruction<sup>45</sup>). While the solution can be approximated via MRF, the large simulated signal dictionaries<sup>12</sup> associated with multi-pool imaging also demand significant computational resources<sup>46,47</sup>, limiting the development of new pulse sequences. A recently reported dictionary-free alternative<sup>48</sup> proposed unsupervised learning for semisolid-MT parameter quantification. However, this method assumes continuous pulse irradiation, which is not available in many clinical scanners, and also relies on analytical solutions, which are not compatible with multi-pool pulsed CEST imaging.

The dictionary-free method presented in this work overcomes all such limitations. Our approach is based on a fundamental insight: by proper formulation, ODE models considered only numerically solvable can become step-wise analytical, and thereby compatible with automatic differentiation-based optimization. Specifically, the suggested formulation enables GPU-based matrix inversion and exponentiation, which translates into efficient gradient descent via back-propagation. Combining this concept with a recently reported high-performing automatic differentiator<sup>38</sup> provides a new option for solving complex biophysical estimation tasks such as pulsed CEST quantification, demonstrated here. Compared to standard



**Fig. 5 | In-vivo study: rapid quantification by applying neural Bloch-McConnell fitting (NBMF) in “transfer mode”.** A comparison between the results of single-subject NBMF (a, b) and real-time quantification of the same subject by inferring the neural reconstructor trained while fitting another subject (c, d). A perceptually and quantitatively similar outputs were obtained for both semisolid (a, c) and amide (b, d) exchange parameters mapping. e Similarity analysis using normalized root-

mean-square (NRMSE), intraclass correlation coefficient (ICC(2,1)), absolute agreement-assessing variant), and structural similarity index measure (SSIM) across all ( $n = 50$ ) processed brain slices. In box plots, the central horizontal lines represent median values, box size represents the two central (2nd, 3rd) quartiles, whiskers represent  $1.5 \times$  the interquartile range above and below the upper and lower quartiles, and circles represent outliers.

model fitting, this approach avoids (i) computationally expensive and inaccurate purely-numerical derivatives computed via multiple evaluations, and (ii) explicit analytical approximations, which can only be applied to a limited subset of cases and lack generalization (e.g., unsuitable for a multi-proton-pool pulsed-RF saturation). Unlike MRF dictionary-trained networks<sup>10,28</sup>, the suggested approach can provide parameter estimates that allow the model to best describe the raw data (Supplementary Fig. 3).

Supervised NN training using a synthetic signal dictionary requires the estimation of the application-specific parameter distribution, which is often unknown in advance. The self-supervised (NBMF) approach circumvents this challenge by training on the in-vivo data itself, offering improved parameter distribution matching. When it comes to re-using the trained network on new unseen subjects, one drawback of this approach lies in its reliance on previously represented proton exchange parameters. Dictionary-based approaches, on the other hand, have the flexibility for representing the expected abnormality values (if they are known) or simply using a very broad parameter distribution that covers both the healthy and diseased states. A future patient cohort investigation is needed to examine the clinical utility of transferring the self-supervised quantification approach, when trained on healthy volunteers, for quantification in unseen pathology (e.g., small lesions).

As shown in Table 1, the most time-consuming steps for supervised dictionary-based learning are the dictionary generation step followed by neural network training. However, if the imaging scenario is a priori known (e.g., brain cancer treatment monitoring) and the acquisition protocol parameters are fixed and optimized, these steps can be done once without affecting the rapid inference time for each new subject. Self-supervised data-based learning (NBMF), on the other hand, offers the flexibility of

accommodating various imaging scenarios and is more easily adapted for new acquisition protocols and research directions. That being said, the biophysical modeling developed as part of this work (GPU-based JAX implementation of the Bloch McConnell numerical solution) can also accelerate both dictionary generation and supervised learning (Table 2), allowing the user to utilize and compare all different approaches.

The gradient of the forward model can be directly used for simple fitting of the unknown ODE coefficients corresponding to the physical parameters of interest (see voxelwise BMF in the “Methods” section, Supplementary Note 3, and Supplementary Fig. 6). However, when the core forward model automatic differentiator is also integrated into a self-supervised learning pipeline (NBMF, Fig. 1), a joint neural representation of the signal-to-parameter transformation can be trained and stored with little extra computational cost. This enables: (i) Faster convergence, which scales well with the number of voxels up to the full brain size, leveraging redundancy towards a spatially smoother solution (Supplementary Fig. 6). (ii) Later reuse for real-time inference on new subjects within a similar imaging scenario.

The human brain imaging results (Figs. 3–5, Supplementary Figs. 3–6, 8) reveal the potential for using an autodiff-compatible Bloch-McConnell solver for parameter quantification while training a simple multilayer perceptron (MLP) voxelwise. Combined with a 3D whole brain acquisition routine (which rapidly generates hundreds of thousands of voxels), the suggested system provides a rapid and efficient single-subject learning method. Notably, while this work presented a proof of concept for rapid inference by a network trained on a single subject, robustness and consistency of the transfer leaves a clear room for improvement (Fig. 5 and Supplementary Fig. 4). Future work could study different NN architectures

**Table 1 | NBMF computational complexity compared to previously reported studies of semisolid-MT and CEST quantification**

stage	Method	Numerical BM model direct fitting	Dictionary-based pattern matching	NN-based mapping (supervised learning)	NBMF (This work)
Preparatory steps using synthetic data + 1st subject mapping	None	Hours to days; e.g., 21.3 h (<56 nodes) for 79 M-entries dictionary generation	Hours; e.g., <4.5 h for 60K-entries dictionary generation and network training <sup>37</sup>	<30 min (fit & train)	<30 min (fit & train)
Nth subject tissue parameters quantification ( $N > 1$ )	Hours to days; e.g., 6 h recon of a $256 \times 256$ pixel slice <sup>45</sup>	Hours to days; e.g., 2.31 h for reconstructing a $128 \times 128$ pixel slice <sup>12</sup>	~1 s <sup>12,37</sup> (NN inference)	~1 s <sup>12,37</sup> (NN inference)	~1 s <sup>12,37</sup> (NN inference) or: <30 min (re-fitting)

The presented times refer to a whole brain reconstruction on a single desktop machine unless noted otherwise. NBMF: Neural Bloch-McConnell Fitting.

with spatial awareness (via convolutional or attention layers), as well as larger datasets composed of multiple subjects and a combination of both dictionary-based synthetic data and real-world scans. Subsequent efforts could also improve the modeling accuracy by taking under consideration the contributions of additional proton pools (such as amine and guanidinium) to the 3.5 ppm signal.

The proposed approach could be further exploited for other tasks across the CEST-MRF pipeline, such as accelerated dictionary synthesis<sup>49,50</sup> (as demonstrated in Supplementary Note 2 and Table 2) and pulsed-wave irradiation compatible CEST protocol discovery and optimization<sup>46,47,51–53</sup>. Furthermore, NBMF is directly applicable to anatomical-MRF (proton density,  $T_1$ ,  $T_2$ ) dictionary-free parameter quantification and conversely, to classical non-MRF molecular MRI, such as pulsed multi- $B_1$  Z-spectra fitting<sup>25,54</sup> (see Supplementary Notes 4, 5). While auto-differentiation of the Bloch equations has previously been leveraged for several MRI-related applications<sup>47,52,55–57</sup>, to the best of our knowledge this is the first report of utilizing this concept for a generalized Bloch-McConnell-fitting task. Beyond molecular MRI and MRF, this approach can also be applied to any other diagnostic and biophysical domains that involve dictionary-matching<sup>58</sup>.

### Learning to estimate ordinary differential equation (ODE) coefficients from observed data

The general strategy underlying NBMF can be applied to any inverse problem that involves fitting ODEs to observations of a dynamical system. In the biomedical realm alone, this includes cardiovascular<sup>59,60</sup>, pharmacokinetic<sup>61,62</sup>, and epidemiological<sup>63,64</sup> modeling, among many other tasks. In parallel to the exponential growth and improvement in AI performance, the last decade has witnessed a surge of interest in harnessing DL for physics-based problem solving. These efforts have most often been directed into two routes: (i) seeking a solution to a partial differential equation (PDE) as an output of a physics-informed neural network (PINN) that operates on spatial and temporal coordinates<sup>65–69</sup>; widely applied for spatially-resolved dynamics in solid<sup>70</sup> and fluid<sup>71</sup> mechanics, heat transfer<sup>72</sup>, power systems<sup>73</sup>, weather/climate<sup>74</sup>, and diffusion MRI<sup>75</sup>. (ii) Modeling parts of the equation with a NN, yielding a neural ODE/PDE<sup>76</sup>, often employed as a semi-parametric approach for model discovery<sup>77–79</sup>. Interestingly, the relatively simple direct inverse solution approach described here (Fig. 1), whereby a NN is trained to infer the coefficients of an ODE model from a few samples of the dynamics, has not received similar attention. This could open opportunities for the current work to inform new approaches to ODE-driven inverse problems across a multitude of tasks.

### Conclusions

The NBMF framework enables rapid, dictionary-free, pulsed-saturation and multiple proton-pool-compatible semisolid MT/CEST-MRF quantification. By combining a GPU-accelerated auto-differentiable numerical ODE solver and self-supervised DL, the NBMF pipeline is able to match alternative AI-reconstruction based inference, while removing the need for dictionary synthesis. NBMF is three orders of magnitude faster than traditional Bloch fitting, and provides a one-stop-shop for reconstructing quantitative molecular MRI data. The underlying approach has potential applications across a wide variety of ODE-driven inverse problem tasks.

### Methods

#### CEST phantoms

A set of six 10 ml L-arginine (L-arg, chemical shift = 3 ppm, Sigma-Aldrich) phantoms was prepared at a concentration of 25, 50, 75, or 100 mM. The phantoms were titrated to different pH levels between 4.0 and 5.0 and placed in a 120 mm diameter cylindrical holder (MultiSample 120E, Gold Standard Phantoms, UK), filled with saline.

#### Human subjects

Four healthy volunteers (three females/one male, with average age  $23.75 \pm 0.83$ ) were scanned at Tel Aviv University (TAU), using a 3T MRI

**Table 2 | Unified benchmarking of all methods using the accelerated approach developed as part of this work**

Stage	Method	Autodiff-based (AD) model fitting ( <b>NBMF<sup>a</sup></b> )	Dictionary-based pattern matching	NN-based mapping (supervised learning)	Self-supervised <b>NN + AD (NBMF<sup>b</sup>)</b>
Preparation + 1st subject mapping		3 min	25 s <sup>b</sup> (generation) + 93 s (matching)	25 s (generation) + 58 s (training)	3 min (fit & train)
Nth subject mapping ( $N > 1$ )		3 min	93 s (matching)	1 s (NN inference)	1 s (inference)
Consistency with raw per-subject data		✓(up to convergence)	✓(up to grid resolution)	✗(empirically, in our test)	✓(up to prior)
Implicit NN-based smoothing prior		✗	✗	✓	✓

It utilizes a GPU-based JAX formulation of the Bloch-McConnell ODEs numerical/analytical solutions. The presented times refer to a whole brain reconstruction on a single GPU-equipped desktop machine. The comparison was based on a semisolid MT MRF protocol using a dictionary consisting of 400K entries and a whole-brain quantification task (194K voxels), see additional details in Supplementary Note 2.

Additional analysis and detailed discussion of the compared qualitative features are available in Supplementary Notes 1–3. NBMF neural Bloch-McConnell fitting, VBMF voxelwise Bloch-McConnell fitting.

<sup>a</sup>Methods introduced in this work are highlighted in **bold**.

<sup>b</sup>Exhibiting linear scaling: it took 80 min to generate the 78M entries of the CEST and semisolid MT cartesian grid dictionary used in ref. 12.

equipped with a 64-channel coil (Prisma, Siemens Healthineers). The research protocol was approved by the TAU Institutional Ethics Board (study no. 2640007572-2) and the Chaim Sheba Medical Center Ethics Committee (0621-23-SMC). All subjects gave written, informed consent before the study.

### MRI acquisition

All acquisition schedules were implemented using the Pulseq prototyping framework<sup>80</sup> and the open-source Pulseq-CEST sequence standard<sup>81</sup>. The MRF acquisition protocols were implemented as described previously<sup>12,43</sup>, with an unsaturated  $M_0$  image added at the beginning of each sequence. A spin lock saturation train ( $13 \times 100$  ms, 50% duty-cycle) was used for each one of the 30 additional iterations of the sequence, which varied the saturation pulse power between 0 and 4  $\mu$ T (average pulse amplitude). The saturation pulse frequency offset was fixed at 3 ppm for L-arginine phantom imaging<sup>44</sup>, 3.5 ppm for amide brain imaging<sup>12</sup>, or varied between 6 and 14 ppm for semisolid MT imaging<sup>43</sup>. The saturation block was followed by a 3D centric reordered EPI readout module<sup>82,83</sup>, providing a 1.8 mm isotropic resolution. The in-plane axial matrix size was  $116 \times 88$ , with 50 slices (169K–194K brain voxels) used per subject. The full sequences can be accurately reproduced using previously published Pulseq (.seq) files<sup>49</sup>. Each 3D MRF acquisition took 2:36 (min:s). The same readout module was used for acquiring additional  $B_0$ ,  $B_1$ ,  $T_1$ , and  $T_2$  maps, via WASABI<sup>84</sup>, saturation recovery, and multi-echo sequences, respectively. The total scan time per subject was 9 min. The WASABI sequence used a preparation scheme realized by a rectangular pulse of 5 ms and nominal  $B_1 = 3.7 \mu$ T. Twenty-four frequency offsets were equally spaced between  $-1.8$  ppm and  $1.8$  ppm with a recovery time of 4.5 s. An  $M_0$  image was taken at  $-300$  ppm with a recovery time of 12 s. The saturation recovery  $T_1$  mapping protocol used the following TR (s) values: 10, 6, 4, 3, 2, 1, 0.8, 0.5, 0.4, 0.3, 0.2, 0.1. The  $T_2$  mapping multi-echo sequence used the following echo times (s): 0, 0.01, 0.025, 0.03, 0.04, 0.05, 0.1, 0.2, 0.3, 0.5, 1.0 with a TR = 10 s.

### MRI data pre-processing

In vitro images with no L-arginine vials, partial vials, or severe image artifacts were removed. Regions of interest (ROIs) were defined using circular masks. In-vivo brain images were motion-corrected and registered using elastix<sup>85</sup>. WM/GM ROI segmentation of the  $T_1$  map was performed using statistical parameter mapping<sup>86</sup>. Quantitative reference CEST-MRF values (Fig. 2) were obtained using dot-product matching, as extensively described previously<sup>44,49</sup>.

### NBMF architecture for semisolid-MT and CEST quantification

The self-supervised learning framework comprises two main components (Fig. 1, Top):

(A) Reconstructor  $\mathcal{R}$  - a fully-connected multi-layer perceptron (MLP) NN, applied voxel-wise on the raw input data<sup>10,12,13,87</sup>. The NN is composed of three layers, with 256 neurons and ReLU activation in each hidden layer. The output layer consists of 5 neurons, encoding the estimates for the proton volume fraction and exchange rate of the compound of interest and their joint uncertainty expressed as noise covariance. It utilizes a sigmoid activation, with the output scaled to a predefined range of the parameter values, which effectively defines the prediction boundaries as follows: semisolid proton volume fraction  $f_{ss} \in [0, 30]$  (%), semisolid proton exchange rate  $k_{ssw} \in [0, 150]$  ( $s^{-1}$ ), amide proton volume fraction  $f_a \in [0, 1.2]$  (%), amide proton exchange rate  $k_{sw} \in [0, 400]$  ( $s^{-1}$ ), L-arginine concentration [ $L\text{-arg}$ ]  $\in [10, 120]$  (mM) and  $N_{\alpha}\text{-amine}$  (of L-arginine) proton exchange rate  $k_{sw} \in [100, 1400]$  ( $s^{-1}$ )<sup>44</sup>. Several auxiliary maps  $\mathbf{X}$ , including water relaxation  $T_1$ ,  $T_2$ , and  $B_0/B_1$  inhomogeneities, are appended to the MRF raw data  $\mathbf{D}$  to be used as inputs for the tissue parameter estimation:  $\mathbf{P} = \mathcal{R}(\mathbf{D}, \mathbf{X}, w)$  where  $w$  are the weights to be trained.

(B) Simulator  $\mathcal{F}$ —a differentiable multi-pool spin physics solver. A numerical simulation of the piecewise-constant coefficient Bloch-McConnell (BM) differential equations was implemented in the open-source JAX<sup>38</sup> computing framework, leveraging its strong auto-differentiation and GPU-

acceleration capabilities for matrix operations. The simulator concatenated and chained the calculations of the BM closed-form solution across all pulses and delays of the protocol. This was carried out by inversion and exponentiation of the  $9 \times 9$  BM-matrix  $\mathbf{A}$ , which expresses all precession, saturation, relaxation and exchange terms of the multi-pool magnetization vector ( $\mathbf{M}$ ) dynamics, as previously defined<sup>33</sup>:

$$\begin{aligned}\dot{\mathbf{M}} &= -\mathbf{AM} + \mathbf{C} \rightarrow \mathbf{M}_{eq.} = \mathbf{A} \setminus \mathbf{C}; \mathbf{M}(t + \Delta t) \\ &= e^{-\mathbf{A}\Delta t} (\mathbf{M}(t) - \mathbf{M}_{eq.}) + \mathbf{M}_{eq.}\end{aligned}\quad (1)$$

This solver is compatible with the rectangular pulse-train shape employed in this study and others<sup>12,43,81</sup>, while arbitrary pulse shapes can be supported using a simple matched-RMS approximation, or to any order through a Magnus expansion<sup>88</sup>. For two-pool imaging cases (such as semisolid MT data acquired using saturation pulses with a frequency offset higher than 6 ppm), additional acceleration was obtained by implementing the interleaved saturation-relaxation (ISAR2) approximate analytical solution<sup>33</sup> of the saturation stage. The RF pulses of spin-lock and readout flips were approximated as hard pulses generating precise flip angle rotations.

The model is designed to represent the whole sequence by simulating the Z-magnetization dynamics during the recovery, saturation, and readout stages, provided that spoilers are applied. For each of the two (semisolid MT/amide) sequences, the Nx31 non-steady-state MRF measurements from 169K-194K brain voxels were normalized using an unsaturated  $M_0$  reference image. Thus, the resulting acquired data  $\mathbf{D} = \{D_n\}_{n=1}^N \in [0, 1]$  is directly related to the magnetization vector governed by Eq. (1), at the end of the saturation pulse block. Therefore, given  $\tilde{\mathbf{P}}$ , an estimate of the sought parameters, the simulator provides a re-synthesis of the data as:  $\tilde{\mathbf{D}} = \mathcal{F}(\tilde{\mathbf{P}}, \mathbf{X}, \omega_{rf}, B_1)$ , where  $\omega_{rf}$  and  $B_1$  are the saturation pulse frequency offsets and powers implemented in the MRF protocol, and  $\mathbf{X}$  are any known tissue parameters.

The NBMF reconstruction of the semisolid-MT proton exchange parameters from the first (1) subject, was obtained by using the MT-MRF data  $\mathbf{D}_{ss}^{(1)}$  alongside independently quantified auxiliary parameter maps  $\mathbf{X}_{w,B}^{(1)} = \{T_{1w}, T_{2w}, B_1, B_0\}$ , for training the weights  $w_{ss}^{(1)}$  of a neural reconstructor  $\mathcal{R}_{MT}^{(1)}$ , designed to quantify the associated proton exchange parameters ( $\tilde{\mathbf{P}}_{ss}^{(1)} = f_{ss}, k_{ssw}$ ). To that end, the NBMF optimizes the following self-supervised objective of consistency with the biophysical model  $\mathcal{F}$ :

$$\begin{aligned}w_{ss}^{(1)} &= \operatorname{argmin}_w \|\tilde{\mathbf{D}} - \mathbf{D}_{ss}^{(1)}\| \\ &= \operatorname{argmin}_w \|\mathcal{F}(\mathcal{R}(\mathbf{D}_{ss}^{(1)}, \mathbf{X}_{w,B}^{(1)}, w), \mathbf{X}_{w,B}^{(1)}) - \mathbf{D}_{ss}^{(1)}\|\end{aligned}\quad (2)$$

The  $L_2$  norm was used as the regression loss. A cosine-decay learning rate schedule and simple early-stopping upon convergence (loss trend reaching plateau) were applied. Augmentation by noise was applied, twice: (a) Adding a  $\pm 0.1\%$  Gaussian noise to the raw samples (b) Adding a Gaussian noise to the  $f, k$  tissue parameters estimate, using covariance derived from extra outputs of the NN, inspired by a recent work<sup>89</sup>.

This process was repeated using the non-steady-state amide raw MRF data  $\mathbf{D}_s^{(1)}$  for NBMF quantification of the amide proton exchange parameters ( $\mathbf{P}_s = f_s, k_{sw}$ ). For human brain experiments, we also appended the semisolid MT pool parameter estimates  $\tilde{f}_{ss}$  and  $\tilde{k}_{ssw}$  (obtained from the semisolid MT NBMF procedure) to the auxiliary vector  $\mathbf{X}$ . This vector served as input for the amide reconstructor  $\mathcal{R}_s$  and the 3-pool biophysical model  $\mathcal{F}_s$ , so that:  $\mathbf{X}_{B,w,ss} = \{X_{B,w}, P_{ss}\} = \{T_{1w}, T_{2w}, B_1, B_0, f_{ss}, k_{ssw}\}$ <sup>12</sup>. For the two-pool L-arginine phantom experiments, the auxiliary parameters were assigned constant values based on previous reports ( $T_{1w} = 2800$  ms,  $T_{2w} = 1200$  ms)<sup>43</sup>.

Importantly, we obtain both the subject-specific proton exchange parameters  $\tilde{\mathbf{P}}^{(1)} = \mathcal{R}^{(1)}(\mathbf{D}^{(1)}, \mathbf{X}^{(1)})$  and the trained reconstructor  $\mathcal{R}$  at the

convergence of the NBMF. This enables ultra-fast quantification of the proton exchange parameters  $\tilde{\mathbf{P}}^{(2)} = \mathcal{R}^{(1)}(\mathbf{D}^{(2)}, \mathbf{X}^{(2)})$  from a new subject (2) (Fig. 1 bottom). Notably, this rapid inference is only applicable for new data drawn from the same distribution and cannot be applied to entirely new systems (such as muscle creatine quantification using brain-data trained NBMF).

As a natural ablation of the system by removing the neural component, the auto-diff simulator can be used for direct voxelwise parameter fitting:  $\tilde{\mathbf{P}}^{(1)} = \operatorname{argmin} \|\mathcal{F}(\mathbf{P}^{(1)}) - \mathbf{D}^{(1)}\|$ , referred to here as voxelwise Bloch-McConnell fitting (VBMF). This simpler process can be described in the context of Fig. 1 as stopping the gradients at the tissue parameters, which now assume the role of independent per-voxel variables. Apart from the obvious drawback of not yielding a neural reconstructor, VBMF's performance is inferior to NBMF for brain imaging (Supplementary Fig. 6), which we ascribe to the implicit smoothing regularization by the neural network. However, it is a viable direct method for in vitro analysis that is equally able to converge to the minimum of the modeling-error landscape (Supplementary Figs. 1, 2).

Finally, additional acceleration was achieved by parallelization of the computational graph across consecutive readout pairs  $\{D_{n-1}, D_n\}$ , decoupling the single-iteration simulators  $\mathcal{F}_n$ . Assuming that the  $D_{n-1}$  snapshot captures the preceding spin history evolution, the re-synthesis stage is now formulated as  $\tilde{\mathbf{D}} = \{\mathcal{F}_n(\tilde{\mathbf{P}}, D_{n-1})\}_{n=1}^N$  and embedded in Eq. (2) as such. See Supplementary Note 4 for further elaboration.

## Statistical analysis

The SSIM and ICC(2,1) were calculated using the open-source SciPy and Pingouin scientific computing libraries for Python. In slice-statistic box plots (Fig. 5e), the central horizontal lines represent median values, box size represents the two central (2nd, 3rd) quartiles, whiskers represent  $1.5 \times$  the interquartile range above and below the upper and lower quartiles, and circles represent outliers. In the voxel-statistic box plots (Figs. 3, 4) the central horizontal lines represent median values, box size represents the two central (2nd, 3rd) quartiles, the whiskers represent the 90 central percentiles and outliers are omitted. Numerical results in the text are presented as mean  $\pm$  SD.

## Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

## Data availability

The data for the phantom experiment and sample human brain datasets (a single slice for each of the  $n = 4$  subjects) are available at <https://github.com/momentum-laboratory/neural-fitting> and <https://doi.org/10.5281/zenodo.15021550>. The complete 3D human data cannot be shared due to subject confidentiality and privacy.

## Code availability

The code used in this work is available at <https://github.com/momentum-laboratory/neural-fitting> and <https://zenodo.org/records/15021550>.

Received: 10 November 2024; Accepted: 21 March 2025;

Published online: 15 April 2025

## References

1. Bernstein, M. A., King, K. F. & Zhou, X. J. *Handbook of MRI Pulse Sequences* (Elsevier, 2004).
2. Weiskopf, N., Edwards, L. J., Helms, G., Mohammadi, S. & Kirilina, E. Quantitative magnetic resonance imaging of brain anatomy and in vivo histology <https://doi.org/10.1038/s42254-021-00326-1> (2021).
3. Radunsky, D. et al. A comprehensive protocol for quantitative magnetic resonance imaging of the brain at 3 tesla. *PLoS ONE* **19**, e0297244 (2024).

4. Taylor, A. J., Salerno, M., Dharmakumar, R. & Jerosch-Herold, M. T1 mapping: basic techniques and clinical applications. *JACC Cardiovasc. Imaging* **9**, 67–81 (2016).
5. Ma, D. et al. Magnetic resonance fingerprinting. *Nature* **495**, 187–192 (2013).
6. Panda, A. et al. Magnetic resonance fingerprinting—an overview. *Curr. Opin. Biomed. Eng.* **3**, 56–66 (2017).
7. Heo, H.-Y. et al. Quantifying amide proton exchange rate and concentration in chemical exchange saturation transfer imaging of the human brain. *Neuroimage* **189**, 202–213 (2019).
8. Poorman, M. E. et al. Magnetic resonance fingerprinting part 1: potential uses, current challenges, and recommendations. *J. Magn. Reson. Imaging* **51**, 675–692 (2020).
9. Hsieh, J. J. & Svalbe, I. Magnetic resonance fingerprinting: from evolution to clinical applications. *J. Med. Radiat. Sci.* **67**, 333–344 (2020).
10. Cohen, O., Zhu, B. & Rosen, M. S. Mr fingerprinting deep reconstruction network (drone). *Magn. Reson. Med.* **80**, <https://doi.org/10.1002/mrm.27198> (2018).
11. Kim, B., Schär, M., Park, H. W. & Heo, H. Y. A deep learning approach for magnetization transfer contrast MR fingerprinting and chemical exchange saturation transfer imaging. *Neuroimage* **221**, <https://doi.org/10.1016/j.neuroimage.2020.117165> (2020).
12. Perlman, O. et al. Quantitative imaging of apoptosis following oncolytic virotherapy by magnetic resonance fingerprinting aided by deep learning. *Nat. Biomed. Eng.* **6**, <https://doi.org/10.1038/s41551-021-00809-7> (2022).
13. Cohen, O. et al. Cest MR fingerprinting (CEST-MRF) for brain tumor quantification using EPI readout and deep learning reconstruction. *Magn. Reson. Med.* **89**, <https://doi.org/10.1002/mrm.29448> (2023).
14. Perlman, O. et al. Cest mr-fingerprinting: practical considerations and insights for acquisition schedule design and improved reconstruction. *Magn. Reson. Med.* **83**, 462–478 (2020).
15. Jones, K. M., Pollard, A. C. & Pagel, M. D. Clinical applications of chemical exchange saturation transfer (CEST) MRI. *J. Magn. Reson. Imaging* **47**, <https://doi.org/10.1002/jmri.25838> (2018).
16. Zhou, J., Heo, H. Y., Knutsson, L., van Zijl, P. C. & Jiang, S. Apt-weighted MRI: techniques, current neuro applications, and challenging issues. <https://doi.org/10.1002/jmri.26645> (2019).
17. Vinogradov, E., Keupp, J., Dimitrov, I. E., Seiler, S. & Pedrosa, I. CEST-MRI for body oncologic imaging: are we there yet? *NMR Biomed.* **36**, e4906 (2023).
18. Bricco, A. R. et al. A genetic programming approach to engineering MRI reporter genes. *ACS Synth. Biol.* **12**, 1154–1163 (2023).
19. Vladimirov, N. & Perlman, O. Molecular MRI-based monitoring of cancer immunotherapy treatment response. *Int. J. Mol. Sci.* **24**, 3151 (2023).
20. Wang, K. et al. Creatine mapping of the brain at 3T by CEST MRI. *Magn. Reson. Med.* **91**, 51–60 (2024).
21. Rivlin, M., Perlman, O. & Navon, G. Metabolic brain imaging with glucosamine CEST MRI: in vivo characterization and first insights. *Sci. Rep.* **13**, 22030 (2023).
22. Zhou, J. et al. Review and consensus recommendations on clinical apt-weighted imaging approaches at 3T: application to brain tumors. *Magn. Reson. Med.* **88**, <https://doi.org/10.1002/mrm.29241> (2022).
23. Wu, B. et al. An overview of CEST MRI for non-mr physicists. <https://doi.org/10.1186/s40658-016-0155-2> (2016).
24. Ji, Y., Zhou, I. Y., Qiu, B. & Sun, P. Z. Progress toward quantitative in vivo chemical exchange saturation transfer (CEST) MRI. *Isr. J. Chem.* **57**, 809–824 (2017).
25. Zaiss, M. et al. Quesp and quest revisited—fast and accurate quantitative CEST experiments. *Magn. Reson. Med.* **79**, <https://doi.org/10.1002/mrm.26813> (2018).
26. Meissner, J.-E. et al. Quantitative pulsed CEST-MRI using  $\omega$ -plots. *NMR Biomed.* **28**, 1196–1208 (2015).
27. Woessner, D. E., Zhang, S., Merritt, M. E. & Sherry, A. D. Numerical solution of the Bloch equations provides insights into the optimum design of PARCEST agents for MRI. *Magn. Reson. Med.* **53**, 790–799 (2005).
28. Perlman, O., Farrar, C. T. & Heo, H. Y. MR fingerprinting for semisolid magnetization transfer and chemical exchange saturation transfer quantification. *NMR Biomed.* **36**, <https://doi.org/10.1002/nbm.4710> (2023).
29. McMahon, M. T. et al. Quantifying exchange rates in chemical exchange saturation transfer agents using the saturation time and saturation power dependencies of the magnetization transfer effect on the magnetic resonance imaging signal (QUEST and QUESP): ph calibration for poly-l-lysine and a starburst dendrimer. *Magn. Reson. Med.* **55**, 836–847 (2006).
30. Meissner, J. E. et al. Quantitative pulsed CEST-MRI using omega-plots. *NMR Biomed.* **28**, <https://doi.org/10.1002/nbm.3362> (2015).
31. Wu, R., Xiao, G., Zhou, I. Y., Ran, C. & Sun, P. Z. Quantitative chemical exchange saturation transfer (qCEST) MRI—omega plot analysis of RF-spillover-corrected inverse CEST ratio asymmetry for simultaneous determination of labile proton ratio and exchange rate. *NMR Biomed.* **28**, <https://doi.org/10.1002/nbm.3257> (2015).
32. Zaiss, M. et al. A combined analytical solution for chemical exchange saturation transfer and semi-solid magnetization transfer. *NMR Biomed.* **28**, <https://doi.org/10.1002/nbm.3237> (2015).
33. Roeloffs, V., Meyer, C., Bachert, P. & Zaiss, M. Towards quantification of pulsed spinlock and CEST at clinical MR scanners: an analytical interleaved saturation-relaxation (ISAR) approach. <https://doi.org/10.1002/mrm.24560>. *NMR Biomed.* **28**, <https://doi.org/10.1002/nbm.3192> (2015).
34. Cohen, O., Huang, S., McMahon, M. T., Rosen, M. S. & Farrar, C. T. Rapid and quantitative chemical exchange saturation transfer (CEST) imaging with magnetic resonance fingerprinting (MRF). *Magn. Reson. Med.* **80**, <https://doi.org/10.1002/mrm.27221> (2018).
35. Zhou, Z. et al. Chemical exchange saturation transfer fingerprinting for exchange rate quantification. *Magn. Reson. Med.* **80**, <https://doi.org/10.1002/mrm.27363> (2018).
36. Heo, H.-Y. et al. Unraveling contributions to the z-spectrum signal at 3.5 ppm of human brain tumors. *Magn. Reson. Med.* **28**, 376–383 (2024).
37. Cohen, O. et al. Cest MR fingerprinting (CEST-MRF) for brain tumor quantification using EPI readout and deep learning reconstruction. *Magn. Reson. Med.* **89**, 233–249 (2023).
38. Bradbury, J. et al. JAX: composable transformations of Python + NumPy programs. GitHub. <http://github.com/google/jax> (2018).
39. Sun, Y. et al. Test-time training with self-supervision for generalization under distribution shifts. In *Proc. 37th International Conference on Machine Learning, ICML 2020*, vol. PartF168147-12 (PMLR, 2020).
40. Shocher, A., Cohen, N. & Irani, M. Zero-shot super-resolution using deep internal learning. In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. <https://doi.org/10.1109/CVPR.2018.00329> (2018).
41. Ulyanov, D., Vedaldi, A. & Lempitsky, V. Deep image prior. *Int. J. Comput. Vis.* **128**, <https://doi.org/10.1007/s11263-020-01303-4> (2020).
42. Yaman, B., Hosseini, S. A. H. & Akçakaya, M. Zero-shot self-supervised learning for MRI reconstruction. In *ICLR 2022—10th International Conference on Learning Representations* (ICLR, 2022).
43. Weigand-Whittier, J. et al. Accelerated and quantitative three-dimensional molecular MRI using a generative adversarial network. *Magn. Reson. Med.* **89**, <https://doi.org/10.1002/mrm.29574> (2023).
44. Cohen, O., Huang, S., McMahon, M. T., Rosen, M. S. & Farrar, C. T. Rapid and quantitative chemical exchange saturation transfer (CEST) imaging with magnetic resonance fingerprinting (MRF). *Magn. Reson. Med.* **80**, 2449–2463 (2018).

45. Heo, H. Y. et al. Quantifying amide proton exchange rate and concentration in chemical exchange saturation transfer imaging of the human brain. *NeuroImage* **189**, <https://doi.org/10.1016/j.neuroimage.2019.01.034> (2019).

46. Zhu, B., Liu, J., Koonjoo, N., Rosen, B. R. & Rosen, M. S. AUTOmated pulse SEQuence generation (AUTOSEQ) for MR spatial encoding in unknown inhomogeneous B0 fields. In *Proc. International Society for Magnetic Resonance in Medicine* Vol. 5489 (2019).

47. Perlman, O., Zhu, B., Zaiss, M., Rosen, M. S. & Farrar, C. T. An end-to-end ai-based framework for automated discovery of rapid CEST/MT MRI acquisition protocols and molecular parameter quantification (AUTOCEST). *Magn. Reson. Med.* **87**, <https://doi.org/10.1002/mrm.29173> (2022).

48. Kang, B., Kim, B., Schär, M., Park, H. W. & Heo, H. Y. Unsupervised learning for magnetization transfer contrast MR fingerprinting: application to CEST and nuclear overhauser enhancement imaging. *Magn. Reson. Med.* **85**, <https://doi.org/10.1002/mrm.28573> (2021).

49. Vladimirov, N. et al. Quantitative molecular imaging using deep magnetic resonance fingerprinting. *Nat. Protoc.* <https://doi.org/10.1038/s41596-025-01152-w> (2025).

50. Nagar, D., Vladimirov, N., Farrar, C. T. & Perlman, O. Dynamic and rapid deep synthesis of chemical exchange saturation transfer and semisolid magnetization transfer mri signals. *Sci. Rep.* **13**, 18291 (2023).

51. Loktyushin, A. et al. MRzero—automated discovery of MRI sequences using supervised learning. *Magn. Reson. Med.* **86**, <https://doi.org/10.1002/mrm.28727> (2021).

52. Kang, B., Kim, B., Park, H. & Heo, H.-Y. Learning-based optimization of acquisition schedule for magnetization transfer contrast mr fingerprinting. *NMR Biomed.* **35**, e4662 (2022).

53. Glang, F. et al. MR-double-zero—proof-of-concept for a framework to autonomously discover MRI contrasts. *J. Magn. Reson.* **341**, <https://doi.org/10.1016/j.jmr.2022.107237> (2022).

54. Zaiss, M., Schmitt, B. & Bachert, P. Quantitative separation of CEST effect from magnetization transfer and spillover effects by lorentzian-line-fit analysis of z-spectra. *J. Magn. Reson.* **211**, <https://doi.org/10.1016/j.jmr.2011.05.001> (2011).

55. Zhu, B., Liu, J., Koonjoo, N., Rosen, B. R. & Rosen, M. S. AUTOmated pulse SEQuence generation (AUTOSEQ) using Bayesian reinforcement learning in an MRI physics simulation environment. In *Proc. Joint Annual Meeting ISMRM-ESMRMB* (2018).

56. Lee, P. K., Watkins, L. E., Anderson, T. I., Buonincontri, G. & Hargreaves, B. A. Flexible and efficient optimization of quantitative sequences using automatic differentiation of Bloch simulations. *Magn. Reson. Med.* **82**, 1438–1451 (2019).

57. Luo, T., Noll, D. C., Fessler, J. A. & Nielsen, J. F. Joint design of rf and gradient waveforms via auto-differentiation for 3D tailored excitation in MRI. *IEEE Trans. Med. Imaging* **40**, <https://doi.org/10.1109/TMI.2021.3083104> (2021).

58. Ghodasara, S. et al. Quantifying perfusion properties with DCE-MRI using a dictionary matching approach. *Sci. Rep.* **10**. <https://doi.org/10.1038/s41598-020-66985-9> (2020).

59. Zenker, S., Rubin, J. & Clermont, G. From inverse problems in mathematical physiology to quantitative differential diagnoses. *PLoS Comput. Biol.* **3**, e204 (2007).

60. Linial, O., Ravid, N., Eytan, D. & Shalit, U. Generative ode modeling with known unknowns. In *Proc. Conference on Health, Inference, and Learning* 79–94 (2021).

61. Donnet, S. & Samson, A. A review on estimation of stochastic differential equations for pharmacokinetic/pharmacodynamic models. *Adv. drug Deliv. Rev.* **65**, 929–939 (2013).

62. Chou, W.-C. & Lin, Z. Machine learning and artificial intelligence in physiologically based pharmacokinetic modeling. *Toxicol. Sci.* **191**, 1–14 (2023).

63. Gumel, A. B., Iboi, E. A., Ngonghala, C. N. & Elbasha, E. H. A primer on using mathematics to understand COVID-19 dynamics: modeling, analysis and simulations. *Infect. Dis. Model.* **6**, 148–168 (2021).

64. Keeling, M. J. et al. Fitting to the UK COVID-19 outbreak, short-term forecasts and estimating the reproductive number. *Stat. Methods Med. Res.* **31**, 1716–1737 (2022).

65. Raissi, M., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, <https://doi.org/10.1016/j.jcp.2018.10.045> (2019).

66. Karniadakis, G. E. et al. Physics-informed machine learning. *Nat. Rev. Phys.* **3**, 422–440 (2021).

67. Cuomo, S. et al. Scientific machine learning through physics-informed neural networks: where we are and what's next. *J. Sci. Comput.* **92**, <https://doi.org/10.1007/s10915-022-01939-z> (2022).

68. Rajput, J. R. et al. Physics-informed conditional autoencoder approach for robust metabolic CEST MRI at 7T. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* Vol. 14227 LNCS <https://doi.org/10.1007/978-3-031-43993-3-44> (2023).

69. Kissas, G. et al. Machine learning in cardiovascular flows modeling: predicting arterial blood pressure from non-invasive 4D flow MRI data using physics-informed neural networks. *Comput. Methods Appl. Mech. Eng.* **358**, <https://doi.org/10.1016/j.cma.2019.112623> (2020).

70. Haghhighat, E., Raissi, M., Moure, A., Gomez, H. & Juanes, R. A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics. *Comput. Methods Appl. Mech. Eng.* **379**, 113741 (2021).

71. Cai, S., Mao, Z., Wang, Z., Yin, M. & Karniadakis, G. E. Physics-informed neural networks (PINNs) for fluid mechanics: a review. *Acta Mech. Sin.* **37**, 1727–1738 (2021).

72. Cai, S., Wang, Z., Wang, S., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks for heat transfer problems. *J. Heat Transf.* **143**, <https://doi.org/10.1115/1.4050542> (2021).

73. Huang, B. & Wang, J. Applications of physics-informed neural networks in power systems—a review. *IEEE Trans. Power Syst.* **38**, <https://doi.org/10.1109/TPWRS.2022.3162473> (2023).

74. Kashinath, K. et al. Physics-informed machine learning: case studies for weather and climate modelling. *Philos. Trans. R. Soc. A* **379**, 20200093 (2021).

75. Zapf, B. et al. Investigating molecular transport in the human brain from MRI with physics-informed neural networks. *Sci. Rep.* **12**, <https://doi.org/10.1038/s41598-022-19157-w> (2022).

76. Chen, R. T., Rubanova, Y., Bettencourt, J. & Duvenaud, D. K. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems 31*, (2018).

77. Bradley, W. & Boukouvala, F. Two-stage approach to parameter estimation of differential equations using neural odes. *Ind. Eng. Chem. Res.* **60**, 16330–16344 (2021).

78. Lai, Z., Mylonas, C., Nagarajaiah, S. & Chatzi, E. Structural identification with physics-informed neural ordinary differential equations. *J. Sound Vib.* **508**, <https://doi.org/10.1016/j.jsv.2021.116196> (2021).

79. Rudy, S. H., Brunton, S. L., Proctor, J. L. & Kutz, J. N. Data-driven discovery of partial differential equations. *Sci. Adv.* **3**, <https://doi.org/10.1126/sciadv.1602614> (2017).

80. Layton, K. J. et al. Pulseq: a rapid and hardware-independent pulse sequence prototyping framework. *Magn. Reson. Med.* **77**, 1544–1552 (2017).

81. Herz, K. et al. Pulseq-CEST: towards multi-site multi-vendor compatibility and reproducibility of CEST experiments using an open-source sequence standard. *Magn. Reson. Med.* **86**, <https://doi.org/10.1002/mrm.28825> (2021).

82. Mueller, S. et al. Whole brain snapshot CEST at 3T using 3D-EPI: aiming for speed, volume, and homogeneity. *Magn. Reson. Med.* **84**, 2469–2483 (2020).
83. Akbey, S., Ehses, P., Stirnberg, R., Zaiss, M. & Stöcker, T. Whole-brain snapshot CEST imaging at 7T using 3D-EPI. *Magn. Reson. Med.* **82**, 1741–1752 (2019).
84. Schuenke, P. et al. Simultaneous mapping of water shift and B1 (WASABI)-application to field-inhomogeneity correction of CEST MRI data. *Magn. Reson. Med.* **77**, 571–580 (2017).
85. Klein, S., Staring, M., Murphy, K., Viergever, M. A. & Pluim, J. P. Elastix: a toolbox for intensity-based medical image registration. *IEEE Trans. Med. Imaging* **29**, <https://doi.org/10.1109/TMI.2009.2035616> (2010).
86. Ashburner, J. & Friston, K. J. Unified segmentation. *NeuroImage* **26**, <https://doi.org/10.1016/j.neuroimage.2005.02.018> (2005).
87. Power, J. I. et al. In vivo mapping of the chemical exchange relayed nuclear overhauser effect using deep magnetic resonance fingerprinting. *iScience* **27**, 111209 (2024).
88. Blanes, S., Casas, F., Oteo, J. A. & Ros, J. The Magnus expansion and some of its applications <https://doi.org/10.1016/j.physrep.2008.11.001> (2009).
89. Glang, F. et al. Deepcest 3T: Robust MRI parameter determination and uncertainty quantification with neural networks-application to cest imaging of the human brain at 3T. *Magn. Reson. Med.* **84**, <https://doi.org/10.1002/mrm.28117> (2020).
90. Stanisz, G. J. et al. T1, T2 relaxation and magnetization transfer in tissue at 3T. *Magn. Reson. Med.* **54**, <https://doi.org/10.1002/mrm.20605> (2005).
91. Liu, D. et al. Quantitative characterization of nuclear overhauser enhancement and amide proton transfer effects in the human brain at 7 tesla. *Magn. Reson. Med.* **70**, <https://doi.org/10.1002/mrm.24560> (2013).
92. Carradus, A. J., Bradley, J. M. P., Gowland, P. A. & Mougin, O. E. Measuring chemical exchange saturation transfer exchange rates in the human brain using a particle swarm optimisation algorithm. *NMR in Biomedicine* **36**, e5001 (2023).

## Acknowledgements

The authors thank Tony Stöcker and Rüdiger Stirnberg for their help with the 3D EPI readout. The study was supported in part by MOONSHOT-MED, a joint grant program from Tel Aviv University and Clalit Innovation, the innovation arm of Clalit Health Services, and the Edmond J. Safra Center for Bioinformatics at Tel Aviv University. This work was supported by the Ministry of Innovation, Science and Technology, Israel, and a grant from the Blavatnik Artificial Intelligence and Data Science Fund, Tel Aviv University Center for AI and Data Science (TAD). This project was funded by the European Union (ERC, BabyMagnet, project no. 101115639). Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research

Council. Neither the European Union nor the granting authority can be held responsible for them.

## Author contributions

A.F. and O.P. conceived the computational framework with N.V. providing valuable input. M.Z., N.V., and O.P. developed and/or conducted the imaging protocols and data acquisition. A.F. was responsible for the AI design, optimization, and data analysis. A.F. wrote the manuscript. O.P. supervised the project. All authors reviewed the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s42005-025-02063-8>.

**Correspondence** and requests for materials should be addressed to Or Perlman.

**Peer review information** *Communications Physics* thanks Jianpan Huang and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025