NeuroImage 337 (2026) 122032

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

MicroKAN: Mapping human brain microstructure using diffusion MRI and adaptive nonlinear modeling

Yifei Chen [a,b,1], Zihan Li [b,1], Yuanhan Wang [b], Yijin Li [b], Jialan Zheng [b], Hongjia Yang [b], Mingxuan Liu [b], Tolga Cukur [c], Qiuyun Fan [d,e,f], Ziyu Li [g,***], Jie Lu [a,**], Qiyuan Tian [b,*]

[a] Department of Radiology and Nuclear Medicine, Xuanwu Hospital, Capital Medical University, Beijing, China
[b] School of Biomedical Engineering, Tsinghua Medicine, Tsinghua University, Beijing, China
[c] Department of Electrical and Electronics Engineering, Bilkent University, Ankara, Turkey
[d] Academy of Medical Engineering and Translational Medicine, Tianjin University, Tianjin, China
[e] Tianjin Key Laboratory of Brain Science and Neuroengineering, Tianjin, China
[f] Haihe Laboratory of Brain-Computer Interaction and Human-Machine Interaction, Tianjin, China
[g] Oxford Centre for Integrative Neuroimaging, FMRIB, Nuffield Department of Clinical Neurosciences, University of Oxford, Oxford, United Kingdom

ARTICLE INFO

Keywords:
Deep learning
Self-supervised learning
Physics-informed neural network
DTI
NODDI

ABSTRACT

Diffusion magnetic resonance imaging (dMRI) provides powerful insights into brain microstructure, but conventional microstructural modeling methods require long acquisition times for covering sufficient diffusion directions and are computationally intensive. While deep learning has shown promise in reducing the direction requirement and accelerating the modeling, traditional architectures such as CNNs often struggle to capture the highly nonlinear relationships between multi-shell diffusion signals and microstructural properties. We present MicroKAN, a novel framework built upon Kolmogorov-Arnold Networks with adaptive spline-based activations, specifically designed to represent complex biophysical models with enhanced flexibility and efficiency. MicroKAN supports both supervised and self-supervised paradigms: the supervised variant learns mappings from data to reference metrics, while the self-supervised variant estimates model parameters directly by reconstructing signals through the forward diffusion process, eliminating the need for ground-truth labels. Evaluated on diffusion tensor imaging (DTI) and neurite orientation dispersion and density imaging (NODDI) across multiple datasets, MicroKAN substantially accelerates acquisition and improves the fidelity of microstructural parameter estimation. Beyond supervised training, its self-supervised formulation shows strong robustness to distribution shifts, enabling reliable performance even without annotations. Furthermore, transfer learning with minimal labeled data preserves high accuracy, underscoring the framework's adaptability to diverse scenarios. These advances establish MicroKAN as a versatile and efficient tool for dMRI analysis, offering new opportunities to accelerate neuroscience research and expand the clinical utility of microstructural imaging. Our source code is available at https://github.com/JustlfC03/MicroKAN.

1. Introduction

Diffusion MRI (dMRI) is an important tool for mapping human brain microstructure non-invasively. By measuring and modeling the diffusion patterns of water molecules in brain tissue, dMRI enables in vivo characterization of brain tissue microstructural properties (Le Bihan et al., 1986). Various computational models have been developed to infer tissue information from dMRI signals. For instance, the diffusion tensor model (Le Bihan et al., 2001a), provides a signal representation that quantifies apparent diffusivity, characterizing white matter integrity (Kubicki et al., 2005), myelination (Fujiyoshi et al., 2016), and related properties. Moreover, the neurite orientation dispersion and density

* Corresponding author at: Center for Biomedical Imaging Research, Tsinghua University, 30 Shuangqing Road, Haidian District, Beijing, China PA100084.
** Corresponding author at: Department of Radiology and Nuclear Medicine, Xuanwu Hospital, Capital Medical University, 45 Changchun Road, Xicheng District, Beijing, China PA100053.
*** Corresponding author at: FMRIB, John Radcliffe Hospital, Headington, Oxford, UK OX3 9DU.
E-mail addresses: ziyu.li@ndcn.ox.ac.uk (Z. Li), imaginglu@hotmail.com (J. Lu), qiyuantian@tsinghua.edu.cn (Q. Tian).
[1] Co-first author.

https://doi.org/10.1016/j.neuroimage.2026.122032
Received 28 November 2025; Received in revised form 8 May 2026; Accepted 2 June 2026
Available online 7 June 2026

imaging (NODDI) model (Zhang et al., 2012) quantifies white matter axon density and dispersion. These biophysical models yield biomarkers that reflect cognitive function (Chiari-Correia et al., 2023; Jokinen et al., 2013), aging processes (Yeatman et al., 2014), and mental health characteristics (Al-Sharif et al., 2025; Jiang et al., 2024). Consequently, these capabilities have established dMRI as an essential tool in neuroscientific research and clinical practice (Van Everdingen et al., 1998).

Nonetheless, dMRI modeling is computationally intensive, requiring nonlinear optimization methods that are iterative and must be applied to hundreds of thousands to millions of voxels per subject. This process can take hours, depending on model complexity, imaging resolution, and hardware capabilities, rendering real-time processing and large-scale imaging studies (e.g., HCP (Van Essen et al., 2013), UKB (Alfaro-Almagro et al., 2018; Miller et al., 2016), and ADNI (Jack et al., 2008)) challenging. Moreover, the sheer computational burden limits the feasibility of deploying these models in clinical workflows or large-scale population studies. While methods such as accelerated optimization algorithms and GPU-based parallelization have been proposed to speed up processing, diffusion model fitting remains a significant challenge (Hernández et al., 2013; Xu et al., 2012). Additionally, optimization results are highly sensitive to implementation details, such as initialization strategies and solver settings often exhibiting instability across different acquisition conditions or datasets. This lack of robustness necessitates model-specific reconfiguration and careful retuning for each model or dataset, thereby slowing the development and practical deployment of emerging techniques (Daducci et al., 2015). These limitations highlight the urgent need for alternative approaches, such as data-driven methods, that can offer faster, more scalable, and diffusion-model-agnostic solutions for diffusion MRI analysis.

Fortunately, deep learning (DL) techniques have achieved substantial progress in medical imaging, demonstrating remarkable potential for improving and accelerating dMRI analysis, particularly microstructure model fitting. In contrast to traditional fitting methods that rely on iterative, voxel-wise optimization with model-dependent parameter tuning, DL frameworks using neural networks (NNs) transform the optimization process into a direct signal-to-parameter mapping through a single forward propagation, thereby significantly accelerating model parameter estimation. Importantly, after training, the NN model can be directly applied to new data without voxel-wise iterative fitting or case-specific adjustment of fitting parameters, which further reduces the computational burden during inference. Moreover, these frameworks can automatically learn local and non-local spatially contextualized features, enabling them to represent complex signal patterns effectively. Prior studies have successfully employed NNs, including multilayer perceptron (MLP) (Aliotta et al., 2019; Bertleff et al., 2017; Chen et al., 2023; De Almeida Martins et al., 2021; Golkov et al., 2016; Karimi et al., 2021b; Ye, 2017) and convolutional neural networks (CNNs) (Gibbons et al., 2019; Karimi et al., 2021a; Li et al., 2019; Liang et al., 2023; Lin et al., 2019), to estimate dMRI model parameters accurately and robustly from reduced numbers of diffusion-weighted image (DWI) volumes, highlighting inherent redundancy in q-space data. For example, the q-space deep learning (q-DL) framework using MLPs effectively reduced the number of DWIs required for reliable parameter estimation by 12-fold (Golkov et al., 2016). DeepDTI utilized a CNN to reliably recover both scalar and vector diffusion tensor imaging (DTI) metrics from only six DWIs, the physical limitation for DTI (Tian et al., 2020).

Despite these impressive advances, existing MLP- and CNN-based methods still present practical limitations. Although CNNs are theoretically universal function approximators, their practical reliance on predefined activation functions and linear convolution kernels often necessitates a substantial number of layers to achieve such approximation capability. This structural constraint limits their efficiency in capturing the highly complex and heterogeneous nonlinear relationships inherent in diffusion MRI signal modeling. Recently, Kolmogorov-Arnold Networks (KANs) have emerged as a novel neural network paradigm capable of capturing complex nonlinearities more effectively through the use of learnable activation functions (Liu et al., 2025). This property is particularly relevant for diffusion MRI microstructure modeling. For instance, the NODDI model adopts a three-compartment formulation (intracellular, extracellular, and isotropic compartments) and incorporates a Watson distribution to characterize neurite orientation dispersion. These components jointly induce highly heterogeneous nonlinear relationships between diffusion signals and biophysical parameters, resulting in highly nonlinear signal-to-parameter mapping.

KANs embed adaptive spline-based activation functions directly into their architectures, dynamically adjusting shapes during training to precisely capture nonlinear data distributions. This flexibility enables KANs to represent nonlinear relationships with compact functional forms, improving parameter efficiency and potentially attenuating the risk of overfitting. The central innovation of KANs is inspired by the Kolmogorov-Arnold theorem, which asserts that continuous multivariate functions can be represented using compositions and superpositions of univariate functions. KANs leverage this theorem by incorporating spline-based activation functions within the network architecture. Splines offer remarkable flexibility due to the piecewise polynomial nature, enabling them to adaptively approximate complex functional relationships with fewer parameters than CNNs (Chen et al., 2025). The introduction of KANs demonstrated that spline-based approaches provide compelling alternatives to standard MLP architectures by enhancing representational power while reducing parametric complexity.

Nonetheless, initial KAN architectures faced challenges, including substantial computational overhead and an extensive number of trainable parameters. Subsequent modifications and improvements to KANs have been proposed to address these issues, focusing primarily on enhancing computational efficiency and model compactness. Building upon these developments, recent research has explored extending KANs beyond MLP structures into convolutional layers (Drokin, 2024; Ge et al., 2026). Traditional CNNs utilize linear convolution operations paired with fixed activation functions, thus limiting their flexibility in modeling data complexity. By adapting KAN principles to convolutional structures, convolutional KANs (Conv-KANs) introduce nonlinear convolution operations facilitated by spline-based activations. This adaptation results in significantly improved flexibility, enabling these networks to better capture and represent the subtle features inherent in the complex data.

Motivated by these promising advances, we introduce MicroKAN, a novel deep learning framework that leverages Conv-KANs to improve diffusion MRI model fitting. MicroKAN comprises two integral components designed to address specific challenges in diffusion MRI analysis. The first component, Connect-KAN, is a supervised framework that harnesses adaptive spline-based nonlinear convolutions integrated within a skip-connection architecture. This innovative design enables Connect-KAN to effectively model intricate nonlinear relationships between diffusion MRI signals and microstructure properties. Recognizing the prevalent challenge for obtaining sufficient reference data in diffusion MRI, the second component, DIMOND-KAN, extends Connect-KAN into a self-supervised learning paradigm. By incorporating the physics-informed principle from our recently proposed DIMOND framework (Li et al., 2024), DIMOND-KAN effectively leverages the diffusion data of a particular subject for subject-specific training to achieve accurate microstructural modeling without any reference data as the training target. Extensive experimental validation demonstrates that MicroKAN, through both its supervised Connect-KAN and self-supervised DIMOND-KAN modules, significantly outperforms conventional optimization-based models and existing DL methods. These results highlight the framework's potential applicability for accelerated, accurate, and scalable clinical diffusion MRI analyses, particularly beneficial for large-scale neuroimaging studies.

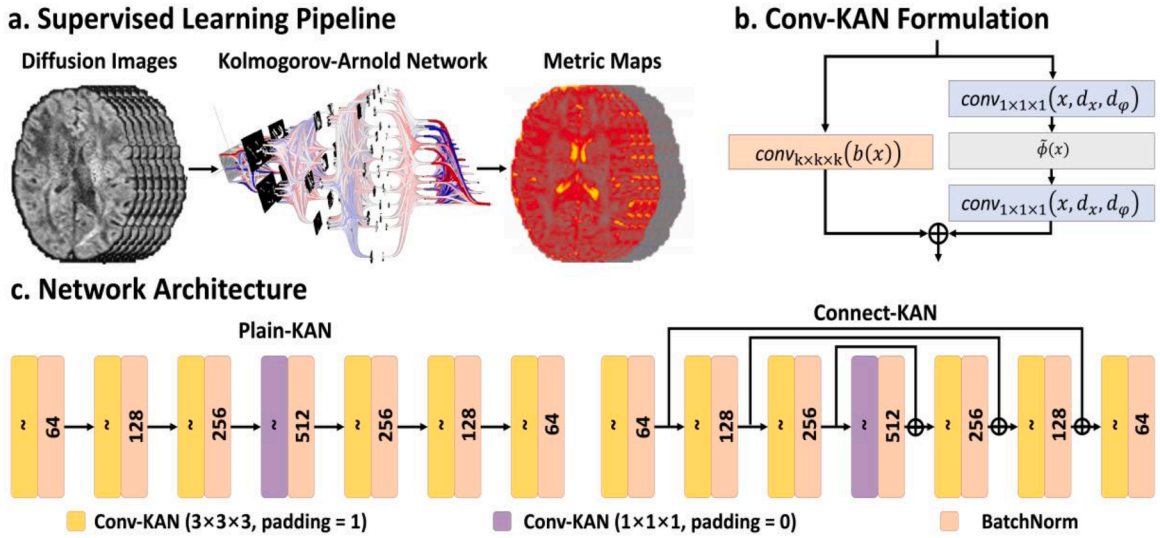

Fig. 1. Supervised learning architecture. (a) Overview of the supervised learning pipeline illustrates the workflow from diffusion-weighted MRI images to microstructural metric maps. (b) Mathematical formulation of the Conv-KAN operation, detailing the nonlinear convolutional mechanism based on spline activations. (c) Detailed depiction of the network architecture, contrasting the Plain-KAN structure for comparison with the Connect-KAN architecture, which incorporates skip connections to effectively enhance feature propagation and preserve spatial information without the use of pooling operations.

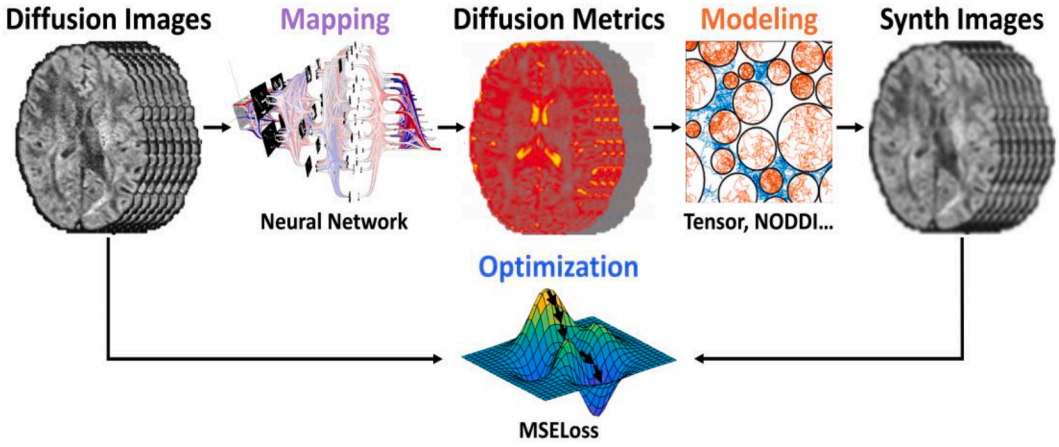

Fig. 2. Self-supervised learning architecture. The self-supervised learning framework DIMOND-KAN extends the Connect-KAN architecture to settings without external data and reference data for training. The model leverages inherent diffusion MRI signal relationships, avoiding explicit reliance on reference microstructural property maps and facilitating accurate microstructure estimation through self-supervision.

2. Methods

2.1. Supervised learning architecture

The proposed supervised learning architecture, Connect-KAN, employs Conv-KANs to capture the complex nonlinearities inherent in the transform from diffusion MRI data to microstructural properties (Fig. 1). Unlike conventional CNNs, which rely on linear convolution operations followed by fixed activation functions, Conv-KANs implement convolution operations based on the Kolmogorov-Arnold representation theorem. According to this theorem, any continuous multivariate function can be approximated using finite compositions and superpositions of univariate functions. Leveraging this theorem, Conv-KANs replace conventional linear convolutions with adaptive spline-based activation functions, significantly enhancing the representation power of convolutional operations.

Formally, the 3D Conv-KAN operation for input volumetric data $Y \in \mathbb{R}^{C_{\text{in}}\times D\times H\times W}$, where C_{in} denotes input channels and $D,\ H,\ W$ represent spatial dimensions, is defined as:

$$X_{c',d',h',w'} = \sum_{c=1}^{C_{\text{in}}} \sum_{a,b,e} \phi_{c,c',a,b,e}\left(Y_{c,d'+a,h'+b,w'+e}\right),$$

where $\phi_{c,c',a,b,e}(\cdot)$ are nonlinear spline activation functions individually learned for each spatial offset $(a,\ b,\ e)$ within the convolution kernel, and channel pair (c, c'). Each spline activation $\phi(z)$ is explicitly parameterized as a linear combination of spline basis functions $B_k(z)$:

$$\phi(z) = \sum_{k=0}^{K} \alpha_k B_k(z),$$

with α_k as trainable spline parameters and K denoting the number of spline knots.

To effectively reduce the computational complexity introduced by numerous spline parameters, the 3D Bottleneck-Conv-KAN (Drokin, 2024) was adopted, extending the Bottleneck design into three dimensions. Each Bottleneck-Conv-KAN layer consists of three distinct steps. First, linear compression via a $1 \times 1 \times 1$ convolution reduces

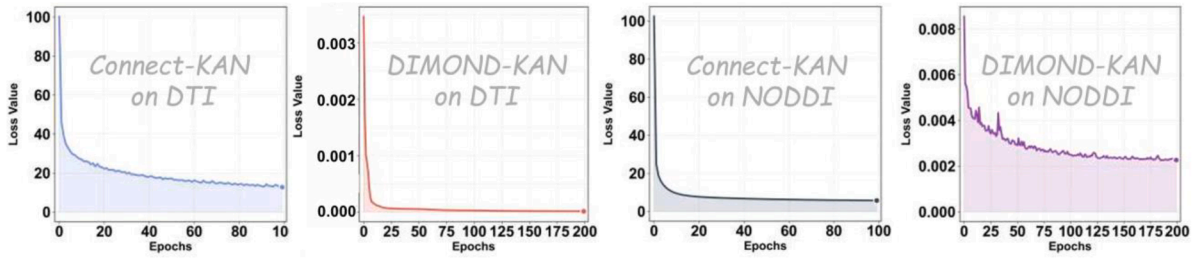

Fig. 3. Training loss curves for assessing convergence stability under supervised and self-supervised learning. The curves show the evolution of training loss across epochs for each training paradigm, illustrating the convergence behavior and training stability under the adopted experimental settings and hyperparameter configurations.

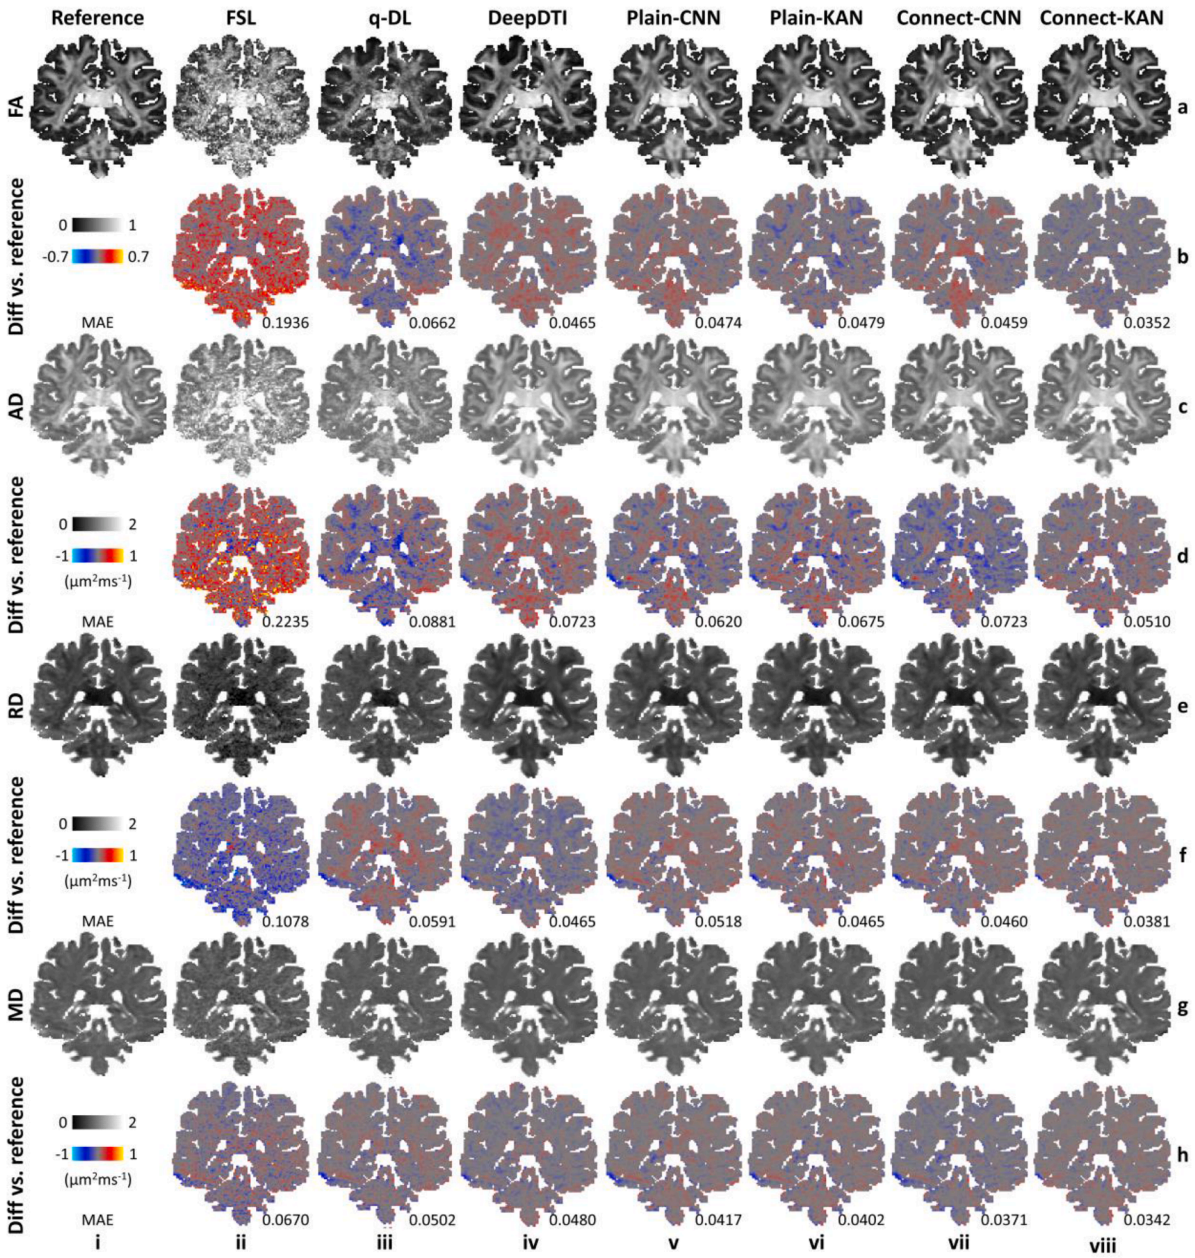

Fig. 4. Maps of DTI metrics from supervised learning training on HCP dataset. Representative coronal maps of DTI metric volumes, including fractional anisotropy (FA), axial diffusivity (AD), radial diffusivity (RD), and mean diffusivity (MD) estimated using supervised Connect-KAN and other methods for comparison on the sub-sampled data from a representative subject are shown (a, c, e, g), along with their residual maps (b, d, f, h) compared to the reference maps (column i). The mean absolute error (MAE) of each map compared to the reference is displayed at the bottom right corner of the residual map.

Table 1
Accuracy of DTI metrics for supervised learning on the HCP dataset. Group mean absolute errors (MAE ± standard deviation) across 10 HCP subjects between resultant and reference DTI metrics, including fractional anisotropy (FA), axial diffusivity (AD), radial diffusivity (RD), and mean diffusivity (MD) for Connect-KAN and other comparison methods using a data subset of N = 6 diffusion directions, and for the FSL OLS implementation using a data subset of N = 32 diffusion directions. Relative MAE decrease is computed as the percentage reduction relative to the baseline OLS implementation of FSL using N = 6 diffusion directions. The red text highlights the lowest MAEs while the blue text highlights the second lowest MAEs.

DTI param	Methods Metrics	FSL (N = 6)	FSL (N = 32)	q-DL	DeepDTI	Plain-CNN	Plain-KAN	Connect-CNN	Connect-KAN
FA	**MAE**	0.2103	0.0417	0.0703	0.0482	0.0484	0.0496	0.0469	0.0370
	Decrease	±0.0150	±0.0031	±0.0037	±0.0013	±0.0014	±0.0023	±0.0014	±0.0026
		Baseline	80.17%	66.57%	77.08%	76.99%	76.41%	77.70%	82.41%
AD ($\mu m^2 ms^{-1}$)	**MAE**	0.2466	0.0826	0.0954	0.0730	0.0662	0.0709	0.0771	0.0556
	Decrease	±0.0184	±0.0039	±0.0037	±0.0026	±0.0025	±0.0020	±0.0043	±0.0020
		Baseline	66.50%	61.31%	70.40%	73.15%	71.25%	68.73%	77.45%
RD ($\mu m^2 ms^{-1}$)	**MAE**	0.1196	0.0597	0.0642	0.0528	0.0533	0.0500	0.0479	0.0415
	Decrease	±0.0078	±0.0032	±0.0035	±0.0028	±0.0031	±0.0026	±0.0025	±0.0029
		Baseline	50.08%	46.32%	55.85%	55.43%	58.19%	59.95%	65.30%
MD ($\mu m^2 ms^{-1}$)	**MAE**	0.0743	0.0611	0.0553	0.0476	0.0441	0.0435	0.0406	0.0379
	Decrease	±0.0039	±0.0036	±0.0037	±0.0032	±0.0027	±0.0024	±0.0025	±0.0026
		Baseline	17.77%	25.57%	35.94%	40.65%	40.24%	45.36%	48.99%

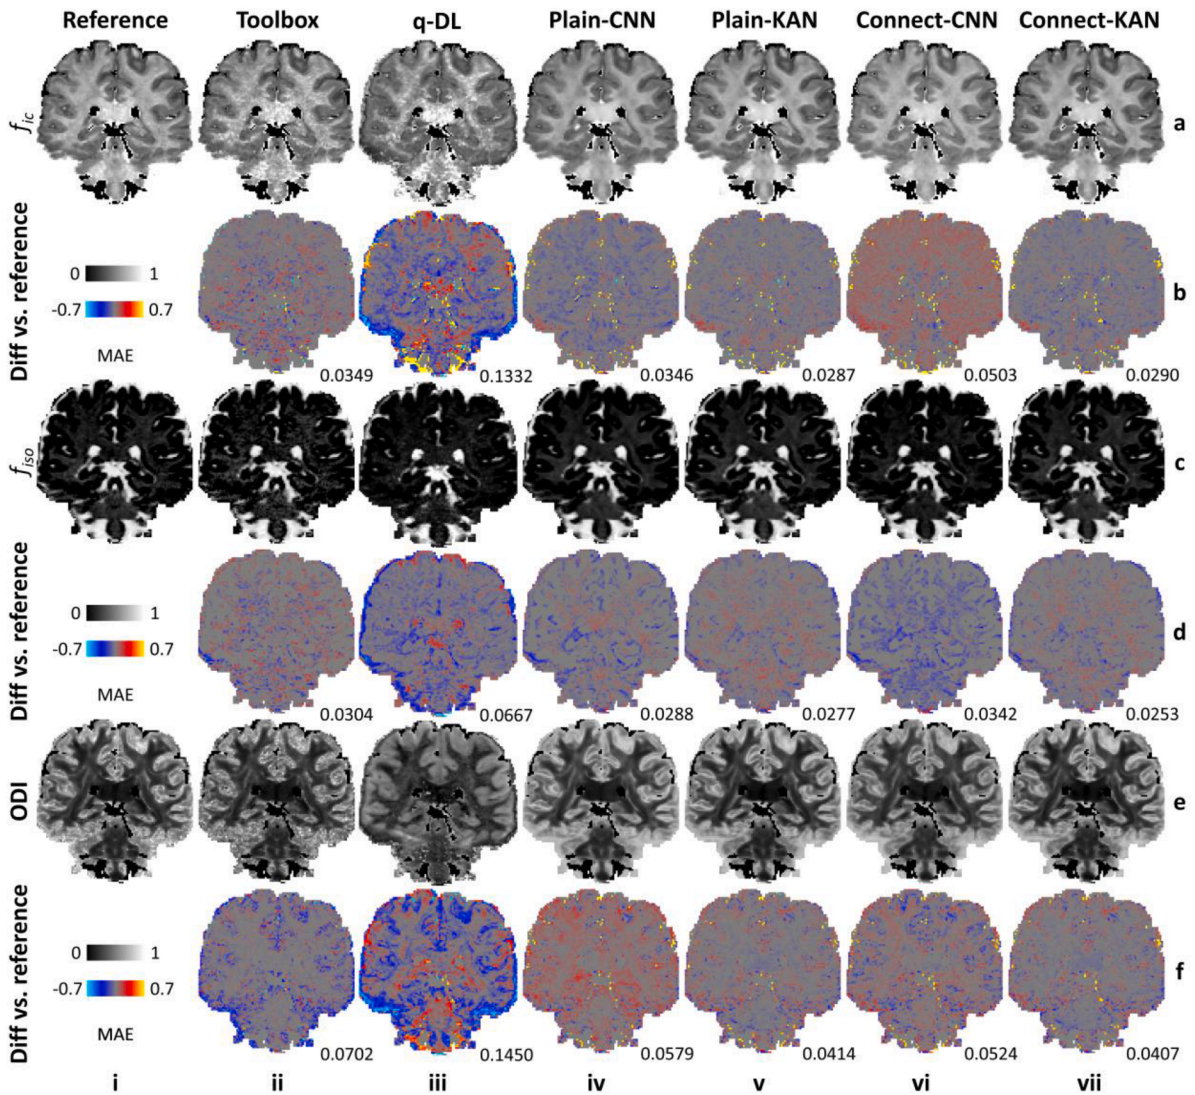

Fig. 5. Maps of NODDI metrics from supervised learning training on HCP dataset. Representative coronal maps of NODDI parameter volumes, including intracellular volume fraction (f_{ic}), isotropic volume fraction (f_{iso}), and orientation dispersion index (ODI) estimated using supervised Connect-KAN and other methods for comparison on the sub-sampled data from a representative subject are shown (a, c, e), along with their residual maps (b, d, f) compared to the reference maps (column i). The mean absolute error (MAE) of each map compared to the reference is displayed at the bottom right corner of the residual map.

channels from C_{in} to a bottleneck dimension C_b:

Table 2
NODDI metrics accuracy for supervised learning training on HCP dataset. Group mean absolute errors (MAE ± standard deviation) across 10 HCP subjects between resultant and reference NODDI metrics, including intracellular volume fraction (f_{ic}), isotropic volume fraction (f_{iso}), and orientation dispersion index (ODI) for Connect-KAN and other methods for comparison. The relative MAE value decrease is computed as the percentage decrease in MAE relative to the baseline NODDI-toolbox implementation. The red text highlights the lowest MAEs while the blue text highlights the second lowest MAEs.

NODDI param	Methods Metrics	Toolbox	q-DL	Plain-CNN	Plain-KAN	Connect-CNN	Connect-KAN
f_{ic}	**MAE**	0.0395±0.0024	0.1370±0.0077	0.0495±0.0016	0.0338±0.0073	0.0357±0.0018	0.0342±0.0074
	Decrease	**Baseline**	−246.89%	−25.32%	14.43%	9.62%	13.42%
f_{iso}	**MAE**	0.0321±0.0018	0.0695±0.0024	0.0358±0.0018	0.0299±0.0023	0.0304±0.0013	0.0294±0.0024
	Decrease	**Baseline**	−116.40%	−11.53%	6.85%	5.30%	8.41%
ODI	**MAE**	0.0671±0.0041	0.1510±0.0081	0.0528±0.0013	0.0435±0.0014	0.0573±0.0016	0.0430±0.0016
	Decrease	**Baseline**	−125.01%	21.31%	35.17%	14.61%	35.92%

$$Z_{b,d,h,w} = \sum_{c=1}^{C_{\text{in}}} W_{b,c} Y_{c,d,h,w}.$$

Then, nonlinear spline activations are voxel-wise applied to compressed features $Z_{b,d,h,w}$:

$$\widetilde{Z}_{b,d,h,w} = \phi_b(Z_{b,d,h,w}) = \sum_{k=0}^{K} \alpha_{b,k} B_k(Z_{b,d,h,w}).$$

Last, linear expansion via another 1 × 1 × 1 convolution expands bottleneck features to the desired number of output channels C_{out}:

$$X_{c',d,h,w} = \sum_{b=1}^{C_b} V_{c',b} \widetilde{Z}_{b,d,h,w},$$

where $W_{b,c}$, $V_{c',b}$ are learnable weights of linear convolutions. By using these Bottleneck-Conv-KANs in the supervised learning architecture, the number of parameters and computational load are further reduced while still maintaining strong nonlinear approximation.

Additionally, our supervised Connect-KAN employs skip-connections between multiple stacked Bottleneck-Conv-KAN layers, facilitating efficient gradient propagation and detailed spatial information retention, ultimately enhancing prediction accuracy and training stability. The Plain-KAN architecture, excluding skip connections, was also implemented for ablation comparison.

2.2. Self-supervised learning architecture

The proposed self-supervised learning framework, DIMOND-KAN, trains the network using the diffusion MRI data from a single subject, i.e., subject-specific training (Fig. 2). For each subject, a Connect-KAN maps the input diffusion data volume Y to diffusion model parameters with network parameters optimized by minimizing the mean-squared error between the input volume Y and the synthesized volume $\widehat{Y}$ derived from feeding the network output $\widehat{P}$ through the forward diffusion model:

$$Loss = \mathbb{E}\|Y - \widehat{Y}\|_2^2 = \mathbb{E}\|Y - Modeling(\widehat{P})\|_2^2.$$

The forward diffusion model function $Modeling(\cdot)$ can be from any existing or new diffusion MRI modeling methods. Therefore, DIMOND-KAN is a physics-informed neural network (PINN). This study was evaluated with the two most widely used diffusion models, i.e., the tensor model and the NODDI model, following the implementations in our previous work (Li et al., 2024).

The tensor model is formulated as:

$$Tensor(S_0, D) = S_0 e^{-bAD},$$

where A and b represent diffusion tensor transformation matrix and b-values, both determined by the acquisition protocol, and D and S_0 denote the diffusion tensors and non-diffusion-weighted signals estimated by DIMOND-KAN. The diffusion tensor transformation matrix A is expanded as:

$$A = [\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6]^T,$$

$$\alpha_i^T = \left[g_{ix}^2, g_{iy}^2, g_{iz}^2, 2g_{ix}g_{iy}, 2g_{ix}g_{iz}, 2g_{iy}g_{iz}\right],\ i = 1, \ldots, N,$$

where each row vector α_i^T is determined solely by the diffusion encoding direction $g_i = \left(g_{ix}, g_{iy}, g_{iz}\right)^T$, and N is the number of diffusion-encoding directions (Kingsley, 2006; Le Bihan et al., 2001b).

The NODDI model is implemented as:

$$NODDI(S_0, f_{iso}, f_{ic}, \kappa, \mu) = S_0((1 - f_{iso})(f_{ic}A_{ic} + (1 - f_{ic})A_{ec}) + f_{iso}A_{iso}),$$

where $A_{ic}(\mu, \kappa)$, $A_{ec}(\mu, f_{ic})$, A_{iso} represent diffusion signals from the intracellular, extracellular, and isotropic compartment, respectively, while f_{ic}, 1- f_{ic}, and f_{iso}represent their volume fractions, respectively (Zhang et al., 2012). Additionally, μ represents the mean orientation and κ measures the orientation dispersion of the Watson distribution.

The intracellular term depends nonlinearly on κ through the Watson distribution, inducing non-uniform signal sensitivity across dispersion regimes. The extracellular component introduces coupled dependencies between f_{ic} and the orientation structure. Moreover, the multiplicative interaction among S_0, f_{iso}, and the compartmental signals further entangles parameter effects. These components jointly induce complex nonlinear relationships that vary significantly across the dynamic range of biophysical parameters, resulting in a highly nonlinear signal-to-parameter mapping. Such structured nonlinearities align well with KAN architectures, whose adaptive spline activations enable localized and flexible approximation of region-varying functional relationships.

2.3. Human connectome project dataset

Pre-processed diffusion MRI data from the Human Connectome Project (HCP) WU-Minn-Ox Consortium public database were utilized in this study. A cohort of 93 unrelated healthy subjects was selected, with 83 subjects allocated for training and 10 reserved for evaluation in the supervised learning framework. The diffusion data from the same 10 subjects were also employed in self-supervised learning experiments for subject-specific training and evaluation. Since detailed acquisition and preprocessing protocols have been thoroughly documented previously (Glasser et al., 2013; Sotiropoulos et al., 2013; Uğurbil et al., 2013), key parameters and essential processing steps are briefly summarized below.

The dMRI acquisition protocol comprised multiple b-values (b = 0, 1000, 2000, 3000 s/mm^2), with 18 non-diffusion weighted image (b = 0) and 90 DWI volumes collected for each non-zero b-value. The b = 0 vol were interspersed between every 15 DWI volumes. Notably, the sampling order for all DWIs was optimized to ensure that sequentially acquired diffusion-encoding directions maintained uniformly distributed across the sphere, even if scans were prematurely terminated. For comprehensive artifact correction, each DWI was acquired with both left-right and right-left phase-encoding directions. During preprocessing, susceptibility and eddy current induced distortions were corrected,

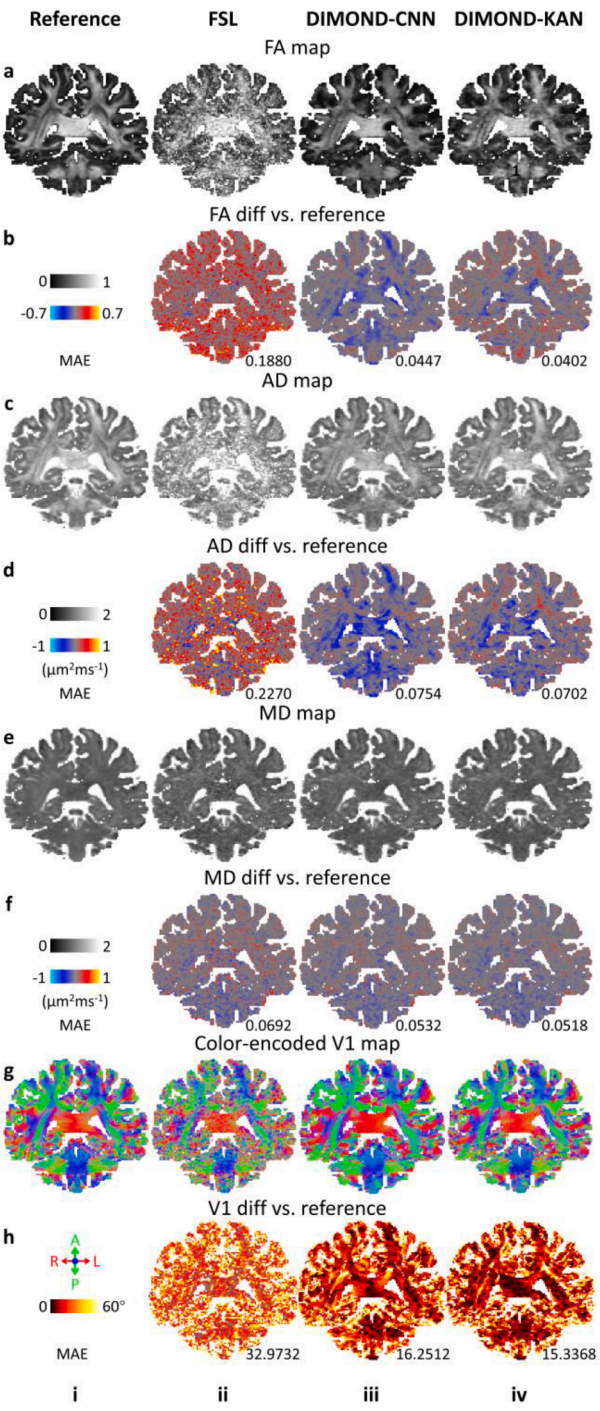

Fig. 6. Maps of DTI metrics from self-supervised learning training on HCP dataset. Representative coronal maps of DTI metric volumes, including fractional anisotropy (FA), axial diffusivity (AD), mean diffusivity (MD), and primary eigenvector (V1) estimated using self-supervised DIMOND-KAN and other methods for comparison on the sub-sampled data from a representative subject are shown (a, c, e, g), along with their residual maps (b, d, f, h) compared to the reference maps (column i). The mean absolute error (MAE) of each map compared to the reference is displayed at the bottom right corner of the residual map.

all images were co-registered, and DWI volumes from opposite phase-encoding directions were combined using the widely-used FSL software (Andersson et al., 2003; Smith et al., 2004).

For DTI experiments, a subset of diffusion MRI data containing one b = 0 and the first six DWI volumes at b = 1000 s/mm^2 from each subject was selected as the highly accelerated DTI data. For NODDI experiments, a subset of diffusion MRI data containing three b = 0 and the first 15 DWI volumes from each of the b = 1000, 2000, and 3000 s/mm^2 shells from each subject was selected as the highly accelerated NODDI data. All 18 b = 0 image volumes and 90 b = 1000 s/mm^2 DWI volumes were both used for deriving reference DTI metrics. All 288 b = 0 and DWI volumes were both used for deriving robust and accurate reference NODDI metrics.

Table 3
DTI metrics accuracy for self-supervised learning training on HCP dataset. Group mean absolute errors (MAE ± standard deviations) across 10 HCP subjects between resultant and reference DTI metrics, including fractional anisotropy (FA), axial diffusivity (AD), mean diffusivity (MD), and primary eigenvector (V1) for DIMOND-KAN and DIMOND-CNN on a data subset of six diffusion encoding directions. The relative MAE value decrease is computed as the percentage decrease in MAE relative to the baseline OLS implemented in FSL using N = 6 diffusion encoding directions. The red text highlights the lowest MAEs while the blue text highlights the second lowest MAEs.

DTI param	Methods Metrics	FSL (N = 6)	DIMOND-CNN	DIMOND-KAN
FA	**MAE**	0.2103	0.0503	0.0498
	Decrease	±0.0150	±0.0046	±0.0045
		Baseline	76.08%	76.32%
AD	**MAE**	0.2466	0.0706	0.0704
(μm^{2}ms^{-1})	**Decrease**	±0.0184	±0.0043	±0.0042
		Baseline	71.37%	71.45%
MD	**MAE**	0.0743	0.0505	0.0501
(μm^{2}ms^{-1})	**Decrease**	±0.0039	±0.0034	±0.0037
		Baseline	32.03%	32.57%
V1 (°)	**MAE**	34.5187	16.7380±1.00	15.3470±0.95
	Decrease	±1.0753	51.51%	55.54%
		Baseline		

2.4. Tsinghua dataset

We further evaluated our method for its generalizability across datasets using a diffusion MRI dataset of 148 subjects from Tsinghua University, allocating 138 for training and 10 for evaluation. The same 10 subjects were further used in self-supervised experiments for subject-specific training and testing. All subjects underwent MRI examination with approval from Tsinghua IRB and written informed consent forms. Diffusion MRI data were collected using a product 2D simultaneous multi-slice pulsed gradient spin echo single-shot echo planar imaging (2D-SMS-PGSE-EPI) sequence with the following parameters: repetition time/echo time = 6600/66 ms, spatial resolution = 1.5 × 1.5 × 1.5 mm^3, SMS factor = 2, GRAPPA acceleration factor = 2, and partial Fourier factor = 6/8. DWIs were acquired at b-values of 0, 1000, and 2500 s/mm^2, incorporating two b = 0 vol of reverse phase-encoding direction (posterior-to-anterior), and 32 and 64 DWIs of uniformly distributed diffusion encoding directions with b-values of 1000 and 2500 s/mm^2 were collected. The total scan time is 12.5 min. The DWI volumes of each subject were corrected for susceptibility and eddy current-induced distortions using FSL software (Andersson et al., 2003; Smith et al., 2004).

For the DTI experiments, we established a highly accelerated acquisition scheme consisting of one b = 0 image and six DWI volumes at b = 1000 s/mm^2. Six diffusion directions that minimize the condition number of the diffusion transformation matrix were selected from all 32 directions. FSL's dtifit was then applied to estimate diffusion tensors. Based on these tensors, synthetic DWIs were generated along the same six diffusion encoding directions as the sub-sampled HCP dataset, thereby yielding equivalent image contrast and enabling effective cross-dataset generalization of HCP-pretrained models. Reference DTI metrics were derived from the complete set of two b = 0 images and thirty-two diffusion-weighted images at b = 1000 s/mm^2.

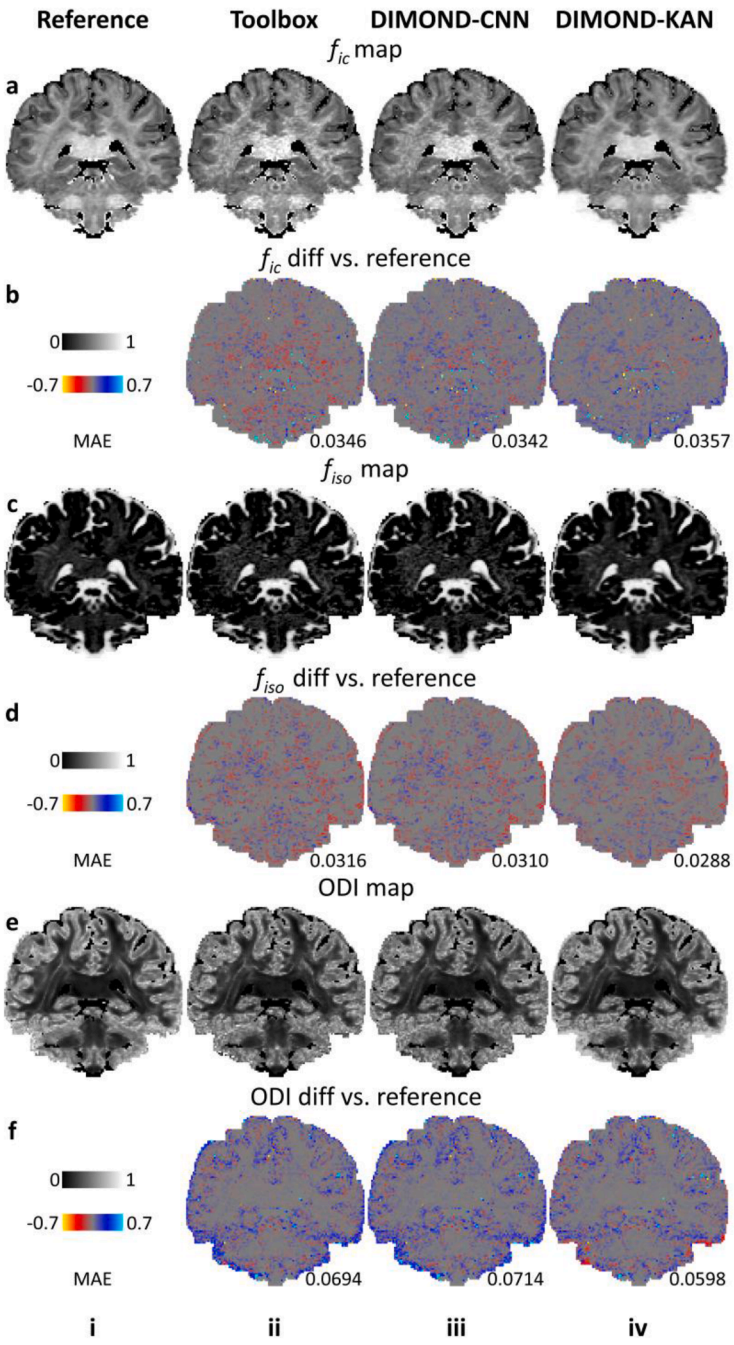

Fig. 7. Maps of NODDI metrics from self-supervised learning training on HCP dataset. Representative coronal maps of NODDI parameter volume, including intracellular volume fraction (f_{ic}), isotropic volume fraction (f_{iso}), and orientation dispersion index (ODI) estimated using self-supervised DIMOND-KAN and other methods for comparison on the sub-sampled data from a representative subject are shown (a, c, e), along with their residual maps (b, d, f) compared to the reference maps (column i). The mean absolute error (MAE) of each map compared to the reference is displayed at the bottom right corner of the residual map.

2.5. Simulated HCP dataset

High-quality diffusion MRI data of 83 subjects (1.25 mm isotropic) from the HCP dataset were used for simulation. Specifically, one b = 0 image volume and six DWI volumes at b = 1000 s/mm^2 were synthesized from the reference tensor model parameters calculated in the section "Human Connectome Project dataset" using the tensor model. Similarly, three b = 0 image volumes and 15 DWI volumes at b = 1000, 2000, 3000 s/mm^2 were synthesized from the reference NODDI model parameters using the NODDI model.

2.6. Conventional model fitting

For DTI, FSL's dtifit function with an ordinary least squares (OLS) regression implementation was used to fit diffusion tensors on the highly accelerated and fully sampled diffusion MRI data to obtain results as the baseline and the reference (Jenkinson et al., 2012). DTI metrics, including standard axial diffusivity (AD), mean diffusivity (MD), fractional anisotropy (FA), radial diffusivity (RD), and robust primary eigenvector (V1), were derived.

For NODDI, the MATLAB toolbox from NODDI inventors was used to fit NODDI's three-compartment model on the highly accelerated and fully sampled diffusion MRI data to obtain results as the baseline and the

Table 4
NODDI metrics accuracy for self-supervised learning training on HCP dataset. Group mean absolute errors (MAE ± standard deviation) across 10 HCP subjects between resultant and reference NODDI metrics, including intracellular volume fraction (f_{ic}), isotropic volume fraction (f_{iso}), and orientation dispersion index (ODI) for DIMOND-KAN and other methods for comparison. The relative MAE value decrease is computed as the percentage decrease in MAE relative to the baseline NODDI-toolbox implementation. The red text highlights the lowest MAEs while the blue text highlights the second lowest MAEs.

NODDI param	Methods Metrics	Toolbox	DIMOND-CNN	DIMOND-KAN
f_{ic}	**MAE**	0.0395	0.0393	0.0394
	Decrease	±0.0024	±0.0026	±0.0024
		Baseline	0.51%	0.25%
f_{iso}	**MAE**	0.0321	0.0303	0.0294
	Decrease	±0.0018	±0.0018	±0.0019
		Baseline	5.61%	8.41%
ODI	**MAE**	0.0671	0.0583	0.0576
	Decrease	±0.0041	±0.0033	±0.0044
		Baseline	13.11%	14.16%

reference using a point-wise grid search combined with gradient descent. NODDI metrics, including intracellular volume fraction (f_{ic}), isotropic volume fraction (f_{iso}), and orientation dispersion index (ODI), were derived from κ.

2.7. Network implementation

All experiments were conducted on a single NVIDIA A800 GPU, using Ubuntu 20.04.6 with Python 3.8 software. The MicroKAN models were implemented using the PyTorch 2.4.1 DL framework, specifically tailored for processing full 3D diffusion MRI data directly, eliminating the conventional need for slice-based preprocessing while preserving information continuity between slices. This model effectively captured complex spatial features intrinsic to diffusion MRI data by leveraging spline-based nonlinear convolutional operations, facilitating the accurate prediction of diffusion metrics.

In the HCP dataset, 10 subjects were randomly selected as an independent test set, while the remaining 83 subjects were designated as available training samples. In the supervised DTI task, training was performed in two stages. The learning rate was initially set to 1×10^{-4} as a standard stable setting for the Adam optimizer to ensure efficient global convergence, which was reduced to 1×10^{-6} to allow fine-grained adjustment of tensor-related parameters and prevent oscillations near local minima in the relatively well-conditioned DTI task. Compared with DTI, NODDI involves stronger nonlinear coupling and higher parameter sensitivity. Therefore, the learning rate for supervised NODDI task was initially set to 1×10^{-4} in the first stage and decreased to 1×10^{-5} in the second stage to balance convergence stability and optimization efficiency without excessively slowing training. For both DTI and NODDI tasks, each stage lasts 100 epochs and about 5 h. For self-supervised tasks, training and validation processes were carried out for each of the 10 evaluation subjects. Specifically, the two-stage learning rates for the self-supervised DTI task were 1×10^{-4} and 1×10^{-5}, respectively. The two-stage learning rates for the self-supervised NODDI task were 1×10^{-3} and 1×10^{-4}. For both DTI and NODDI tasks, each stage lasts 200 epochs and about 2 h.

To mitigate overfitting and enhance convergence stability across different supervision regimes, early stopping was employed in all experiments based on validation loss monitoring. Training was automatically terminated when the validation loss failed to improve for a predefined number of consecutive epochs, and the model parameters corresponding to the lowest validation error were retained for subsequent evaluation. Both supervised and self-supervised settings demonstrated smooth and consistent convergence trends, indicating effective control of overfitting risk under the adopted hyperparameter configurations (Fig. 3).

In the Tsinghua dataset, 10 subjects were randomly selected as an independent test set, while the remaining 138 subjects were designated as available training samples. For supervised tasks, three experimental settings were employed. First, the model weights pretrained on 83 subjects from the HCP dataset were directly applied to the test set of Tsinghua dataset without any additional training. Second, starting from the same pretrained weights, fine-tuning was performed using nine non-test subjects from the Tsinghua dataset, trained for 100 epochs with a learning rate of 1×10^{-4}, requiring approximately 35 min. Finally, full supervised training was conducted from scratch using all 138 non-test subjects for 100 epochs with the same learning rate of 1×10^{-4}, which took about 9.5 h. For self-supervised tasks, one subject from the

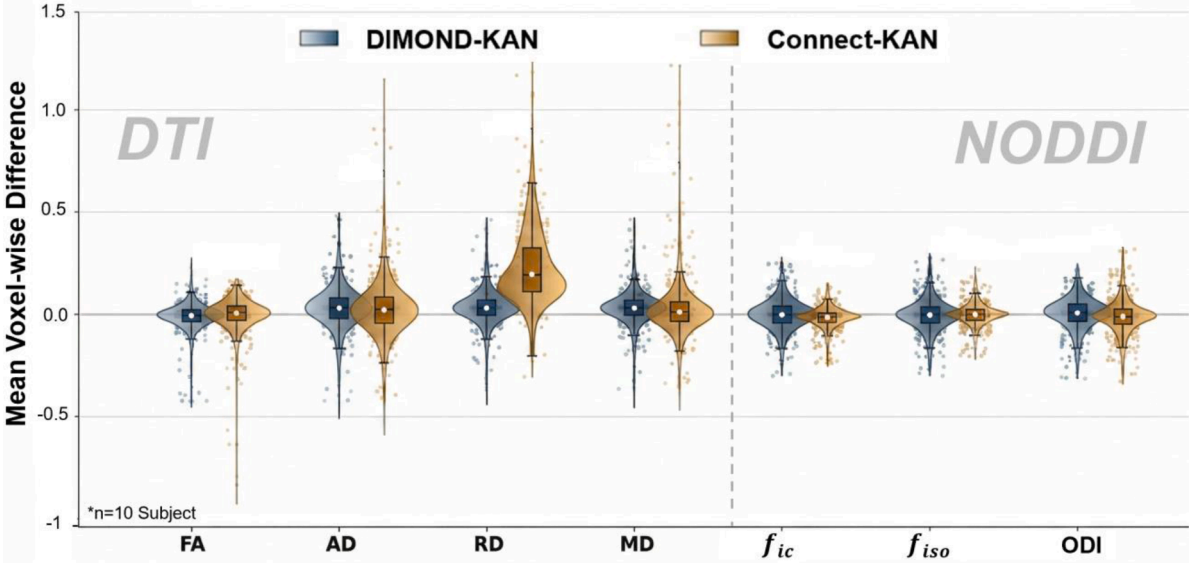

Fig. 8. Bias-variance analysis under supervised and self-supervised regimes. Violin plots show the distribution of mean voxel-wise signed errors between resultant and reference diffusion metrics across 10 HCP subjects (DTI metrics on the left, and NODDI metrics on the right, units for the DTI metrics (AD, RD, and MD) are μm²/ms). DIMOND-KAN and Connect-KAN are displayed in blue and orange, respectively. The horizontal zero line indicates the absence of systematic bias. Median values are indicated by horizontal lines, and the distribution spread reflects error variance. Both models exhibit near-zero median bias for most metrics, with performance differences primarily attributable to variance.

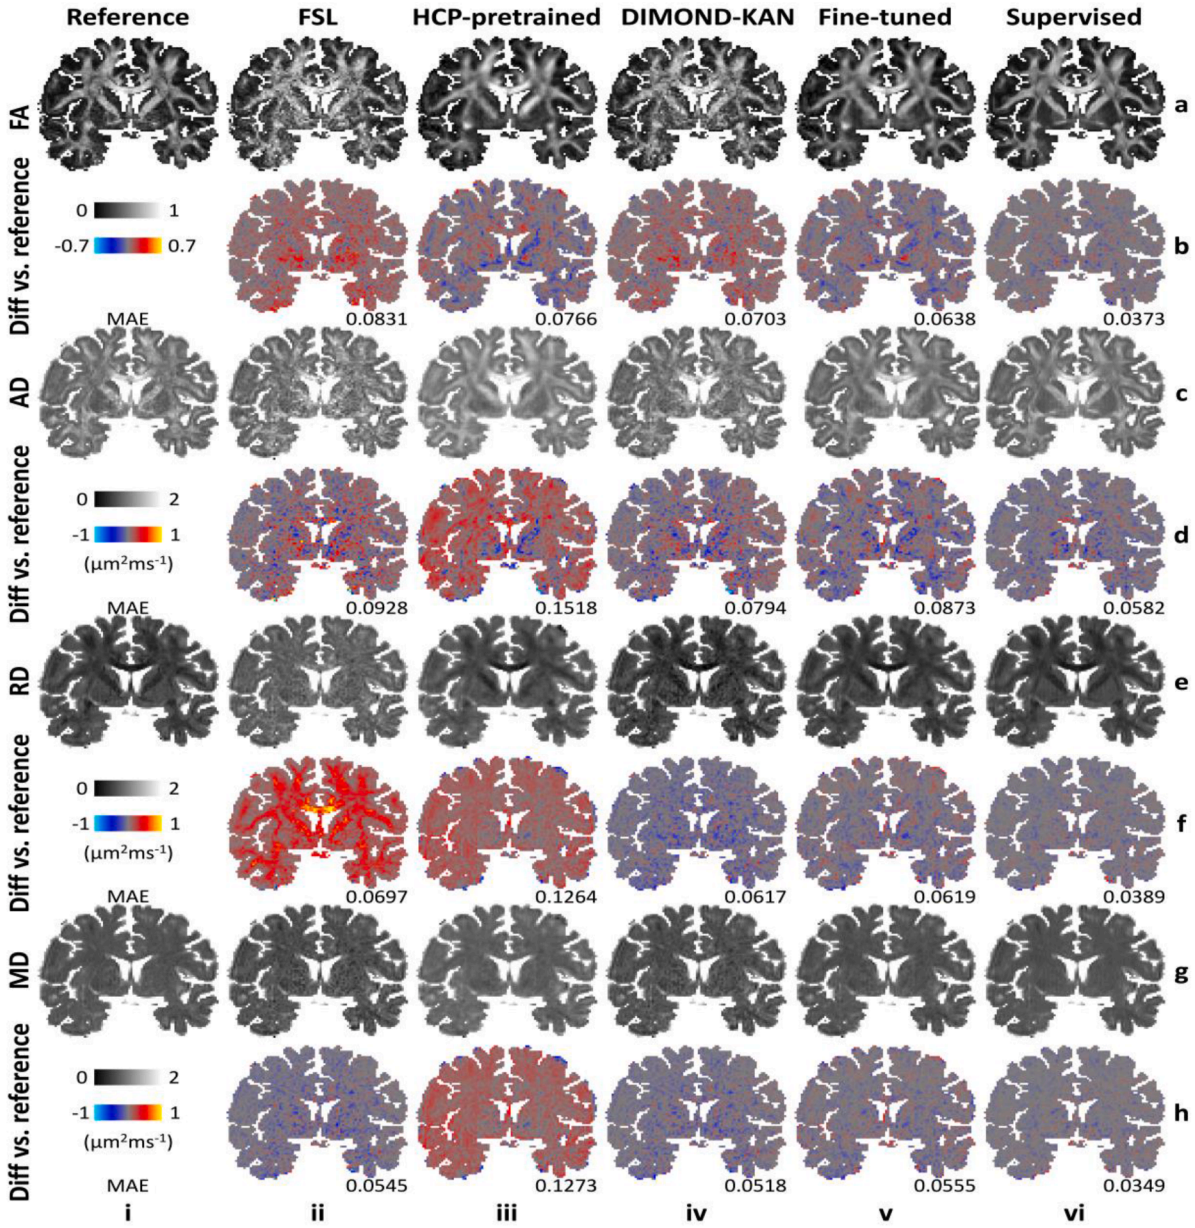

Fig. 9. DTI maps of five training strategies on the Tsinghua dataset. Representative coronal maps for fractional anisotropy (FA), axial diffusivity (AD), radial diffusivity (RD), and mean diffusivity (MD) of reference (i), and those generated on the sub-sampled subset using OLS implemented in FSL (ii), Connect-KAN pretrained on HCP dataset (iii), DIMOND-KAN (iv), HCP-pretrained Connect-KAN fine-tuned on the data of nine Tsinghua dataset subjects (v), and the Connect-KAN trained using randomly initialized parameters on the data of 138 Tsinghua dataset subjects (vi). The corresponding residual maps relative to reference maps are displayed (b, d, f, h) with the mean absolute error (MAE) reported.

Tsinghua test set was selected, trained for 40 epochs with a learning rate of 1×10^{-4}, with the training process lasting approximately 25 min.

Due to limited GPU memory, training was performed using blocks containing $64 \times 64 \times 64$ voxels, resulting in 12 blocks extracted per subject. Both training and validation were conducted using an A800 GPU (NVIDIA, Santa Clara, CA), with validation error closely monitored throughout to identify and save the model achieving the lowest validation error. For inference, the optimized network parameters were applied to the entire brain volume of each evaluation subject for performance assessment.

2.8. Quantitative comparison

The quality of diffusion model metrics was quantified by calculating the mean absolute errors (MAEs) between the estimated and reference metrics for DTI and NODDI. Specifically, DTI metrics, including the AD, MD, FA, RD, and V1, were evaluated within the brain tissue (i.e., excluding cerebrospinal fluid) mask. The error for V1 was assessed as the angular deviation between the estimated and reference orientation vectors. For NODDI, metrics including the f_{ic}, $f_{iso,}$ and ODI were evaluated within the brain mask. Notably, in voxels where f_{iso}approaches 1, f_{ic} and ODI values became arbitrary. Therefore, voxels with f_{iso}values exceeding 0.8 were excluded from the MAE calculations for f_{ic} and ODI. The bias and variance of diffusion model metrics were quantified through the voxel-wise signed errors between the estimated and reference metrics for DTI and NODDI, allowing for the disentanglement of their contributions to the overall error.

Table 5
Quantitative evaluation of four DTI scalar metrics under four training strategies. The table reports the mean absolute errors (MAE) of fractional anisotropy (FA), axial diffusivity (AD), radial diffusivity (RD), and mean diffusivity (MD), showing that DIMOND-KAN and Connect-KAN fine-tuned with nine training samples (Fine-tuned) achieve comparable accuracy, while Connect-KAN trained from scratch with 138 samples (Supervised) provides the best overall performance. The relative MAE value decrease is computed as the percentage reduction in MAE relative to the baseline OLS implemented in FSL using N = 6 diffusion directions. The red text highlights the lowest MAEs while the blue text highlights the second lowest MAEs.

DTI param	Methods Metrics	FSL (N = 6)	HCP-pretrained	DIMOND-KAN	Connect-KAN (Fine-tuned)	Connect-KAN (Supervised)
FA	**MAE**	0.0945±0.0059	0.0809±0.0023	0.0743±0.0021	0.0672±0.0018	0.0407±0.0014
	Decrease	**Baseline**	14.39%	21.38%	28.89%	56.93%
AD ($\mu m^2 ms^{-1}$)	**MAE**	0.1016±0.0064	0.1540±0.0096	0.0823±0.0024	0.0895±0.0029	0.0596±0.0026
	Decrease	**Baseline**	−51.57%	19.00%	11.91%	41.34%
RD ($\mu m^2 ms^{-1}$)	**MAE**	0.0780±0.0035	0.1234±0.0097	0.0635±0.0013	0.0632±0.0021	0.0410±0.0021
	Decrease	**Baseline**	−58.21%	18.59%	18.97%	47.44%
MD ($\mu m^2 ms^{-1}$)	**MAE**	0.0608±0.0026	0.1250±0.0106	0.0524±0.0034	0.0566±0.0024	0.0373±0.0033
	Decrease	**Baseline**	−105.59%	13.81%	6.91%	38.65%

2.9. Statistical analysis

A two-sided paired *t*-test was performed to assess differences in MAE across methods, with Holm correction applied to adjust for multiple comparisons. The 95% confidence interval (CI) and Cohen's dz were reported as measures of effect size. Adjusted P-values of <0.05 were considered statistically significant. All statistical analyses were conducted in Python 3.8.

2.10. Visualization of response function

The response functions of Connect-CNN and Connect-KAN was visualized using a controlled single-channel perturbation analysis. For each scalar value in the normalized input range, we created a zero-valued 3D input patch and assigned this value only to the central $3 \times 3 \times 3$ voxels of one selected input channel, while keeping all other voxels and channels fixed at zero. The perturbed inputs were then passed through the trained network with fixed parameters, and the response was measured from the pre-activation output of the penultimate convolutional layer before BatchNorm and ReLU. Specifically, we averaged the activation of the first output channel over the central $3 \times 3 \times 3$ spatial region and plotted it against the input value. This procedure was repeated for representative input channels to obtain channel-wise response curves.

3. Results

3.1. DTI metrics from supervised Connect-KAN

Connect-KAN recovers remarkably high-quality DTI metric maps (Fig. 4). The outputs of Connect-KAN are visually almost indistinguishable from the fully sampled tensor-fitting reference, delineating major white matter pathways with exceptionally sharp boundaries and clear anatomical contrast. Its difference maps are largely structure-free and exhibit noise-like patterns, indicating predominantly random residuals and minimal bias across the brain. In comparison, FSL produces DTI metrics that are heavily contaminated by noise because of the limited number of diffusion-encoding directions, and the degradation is most obvious in FA and AD. Plain-CNN, Plain-KAN, and Connect-CNN are able to recover fine anatomical detail, yet they introduce more pronounced bias. For example, Connect-CNN clearly overestimates FA in the corpus callosum (see FA difference map, Fig. 4b(vi)), and the boundaries between gray and white matter are blurred when compared with the reference. Compared with supervised deep learning baselines, including q-DL and DeepDTI, supervised Connect-KAN demonstrates superior estimation accuracy and structural fidelity. While q-DL partially recovers the anatomical contrast, its maps remain substantially noisier and less consistent with the reference under the present highly accelerated acquisition setting due to the absence of spatial information. DeepDTI mitigates these artifacts more substantially, yet Connect-KAN still yields sharper structures, less structured residuals, and improved agreement with the reference maps. Overall, Connect-KAN yields sharper and less noisy maps without structured residuals in the error maps, which is consistent with its lowest error among all the learned models.

Quantitatively, supervised Connect-KAN demonstrated significant improvements over baseline methods in predicting DTI metrics (Table 1). Connect-KAN's estimates of FA, AD, RD, and MD are in high agreement with the reference. The MAE for FA is approximately 0.037 in unitless FA values, indicating only minor and practically negligible deviation from the reference. For the diffusivity metrics, the errors are low, ranging from 0.037 to 0.057 $\mu m^2/ms$ for AD, RD, and MD. These MAE values account for only 1–3% of typical diffusivity values, underscoring the high fidelity of the predictions. By comparison, other models yield slightly higher errors. For example, a plain CNN produces an FA MAE of about 0.048 and exhibits considerably and notably larger errors for MD on the same data. Overall, supervised Connect-KAN achieves substantial relative improvements of 48 to 83 percent compared to the baseline methods in terms of MAE. Subject-level two-sided paired *t*-test further confirms the superior performance of Connect-KAN over all competing methods in terms of all supervised DTI metrics (Supplementary Table S6). Specifically, Connect-KAN outperformed FSL with 32 directions for FA (0.0370 ± 0.0026 vs. 0.0417 ± 0.0031, Holm-adjusted $p = 8.97 \times 10^{-12}$), Plain-CNN for AD (0.0556 ± 0.0020 vs. 0.0662 ± 0.0025, Holm-adjusted $p = 7.32 \times 10^{-14}$), and Connect-CNN for both RD (0.0415 ± 0.0029 vs. 0.0479 ± 0.0025, Holm-adjusted $p = 9.12 \times 10^{-12}$) and MD (0.0379 ± 0.0026 vs. 0.0406 ± 0.0025, Holm-adjusted $p = 9.12 \times 10^{-12}$).

3.2. NODDI metrics from supervised Connect-KAN

Connect-KAN consistently reconstructs high-quality NODDI metric maps from sub-sampled data (Fig. 5). The supervised Connect-KAN model produces maps that are highly consistent with the reference. In the f_{ic} map, it delineates clear contrast between white matter and surrounding cortical and ventricular regions. The f_{iso}map accurately highlights elevated isotropic fractions in cerebrospinal fluid spaces while remaining near-zero in brain tissue. The ODI map captures the expected spatial variation of dispersion, with higher values in regions of complex fiber configurations and lower values in coherent bundles. Overall, Connect-KAN outputs are sharp, physically plausible, and minimally deviate from the reference. In contrast, the NODDI-toolbox shows appreciable contrast but higher residual errors. q-DL recovers the overall contrast of the NODDI maps but exhibits less accurate tissue contrast and larger residuals, especially for f_{ic} and ODI. Plain-CNN reduces noise yet produces less consistent dispersion estimates. Plain-KAN achieves smoother and more accurate f_{ic} and f_{iso}than CNN, though subtle errors remain. Connect-CNN preserves structure but introduces bias and noise,

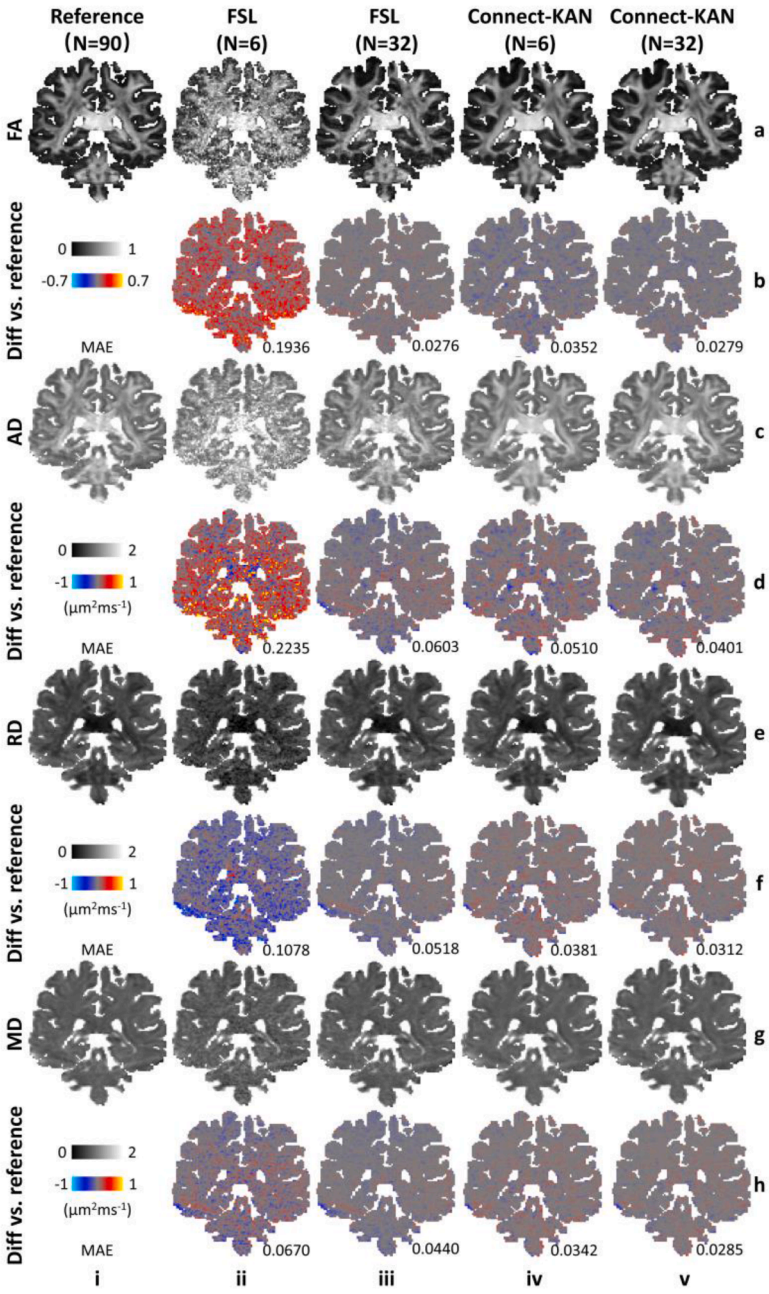

Fig. 10. DTI maps across different number of diffusion encoding directions on the HCP dataset. Representative coronal maps for fractional anisotropy (FA), axial diffusivity (AD), radial diffusivity (RD), and mean diffusivity (MD) of reference (i), and those generated using OLS implemented in FSL on the subset with six DWIs (ii), OLS implemented in FSL on the subset with 32 DWIs (iii), Connect-KAN on the subset with six DWIs (iv), and Connect-KAN on the subset with 32 DWIs (v) are shown. The corresponding residual maps relative to reference maps are displayed (b, d, f, h) with the mean absolute error (MAE) reported.

especially in ODI. Among all, Connect-KAN provides the lowest errors and the most coherent reconstruction across all three metrics.

Quantitatively, supervised Connect-KAN demonstrates clear advantages in predicting NODDI parameters (Table 2). Connect-KAN achieves the lowest errors across most metrics. For intracellular volume fraction f_{ic}, the MAE is 0.0342, corresponding to a 13.42% reduction relative to the baseline. For isotropic volume fraction f_{iso}, the error further decreases to 0.0294, yielding 8.41% improvement. The most pronounced gain is observed for the orientation dispersion index, where Connect-KAN attains an MAE of only 0.0430, marking a substantial 35.92% improvement compared with the NODDI-toolbox. Plain-KAN, although surpassing CNNs in certain cases, for example, a 35.17% reduction for ODI compared with Plain-CNN's 21.31%, still falls short of Connect-KAN. By contrast, Plain-CNN not only fails to improve but even degrades performance for f_{ic} and f_{iso}, showing error increases of 25.32% and 11.53% respectively. These results confirm that the incorporation of the Conv-KAN module and connectivity mechanisms is critical for capturing the complex microstructural features underlying NODDI parameter estimation. A two-sided paired t-test also corroborated the superior performance of Connect-KAN over Plain-KAN, the strongest competing baseline (Supplementary Table S7). Specifically, Connect-KAN achieved significantly lower MAE than Plain-KAN for f_{iso}(-Connect-KAN: Plain-KAN: 0.0294 ± 0.0024 vs. 0.0299 ± 0.0023, Holm-adjusted $p = 2.04 \times 10^{-3}$), and ODI (0.0430 ± 0.0016 vs. 0.0435 ± 0.0014, Holm-adjusted $p = 8.62 \times 10^{-15}$), whereas no statistically significant difference was observed for f_{ic} (0.0342 ± 0.0074 vs. 0.0338 ± 0.0073, Holm-adjusted $p = 0.987$).

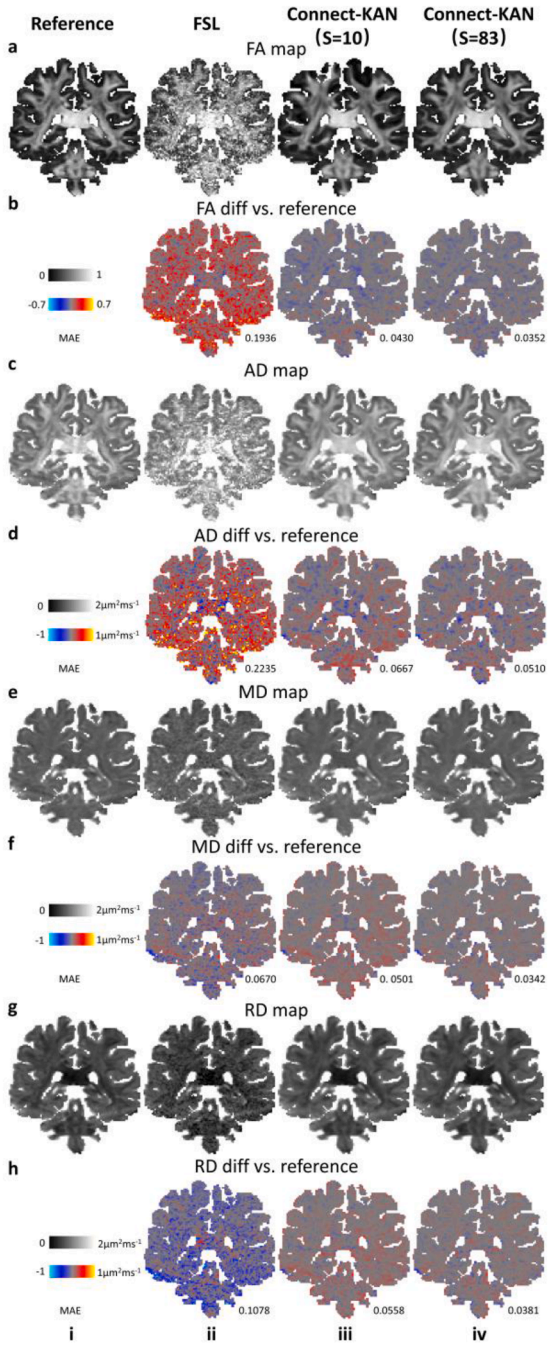

Fig. 11. DTI metric maps across different training dataset size. Representative coronal maps for fractional anisotropy (FA), axial diffusivity (AD), radial diffusivity (RD), and mean diffusivity (MD) of reference (i), and those generated from the subset of six DWIs using OLS implemented in FSL (ii), supervised Connect-KAN trained on 10 subjects (iii), and supervised Connect-KAN trained on 83 subjects (iv) are shown. The corresponding residual maps relative to reference maps are displayed (b, d, f, h) with the mean absolute error (MAE) reported.

3.3. DTI metrics from self-supervised DIMOND-KAN

DIMOND-KAN effectively recovers remarkably high-quality tensor-derived metrics in a self-supervised setting (Fig. 6). The predicted FA, AD, MD, and V1 maps show strong concordance with the fully sampled reference, faithfully and consistently preserving major white matter pathways and fine anatomical features. The FA maps clearly replicate the anisotropy patterns of white matter tracts, whereas FSL reconstructions appear heavily corrupted by noise and loss of contrast. The AD and MD maps from DIMOND-KAN maintain smooth tissue-dependent contrast and regional variations, while DIMOND-CNN displays residual structured errors and slight blurring at tissue boundaries. The V1 maps obtained with DIMOND-KAN accurately capture principal diffusion orientations, in contrast to the noisy and spatially inconsistent outputs of FSL and the more biased estimates of DIMOND-CNN. Residual discrepancies in DIMOND-KAN predictions are spatially incoherent and anatomically non-specific, indicating predominantly random noise rather than systematic bias. Taken together, DIMOND-KAN delivers denoised yet structurally precise outputs, highlighting the advantage of physics-informed self-supervision in recovering reliable tensor-derived metrics.

Quantitatively, self-supervised DIMOND-KAN achieves consistent advantages in predicting DTI metrics (Table 3). DIMOND-KAN yields the lowest errors across all metrics. For fractional anisotropy, the MAE is 0.0498, corresponding to a 76.32% improvement relative to the OLS in FSL. For axial diffusivity, the error is 0.0704 $\mu m^2/ms$, representing a 71.45% reduction. For mean diffusivity, DIMOND-KAN attains an MAE of 0.0501 $\mu m^2/ms$, with a 32.57% improvement. The most pronounced gain is observed in the primary eigenvector, where the error decreases to 15.35°, marking a 55.54% improvement over FSL. Compared with DIMOND-CNN, which also performs strongly, achieving a 76.08% improvement for FA, the proposed DIMOND-KAN further lowers errors in most cases. An exception is MD, where DIMOND-KAN shows a marginal advantage of 32.57% compared with 32.03% for DIMOND-CNN, likely reflecting the composite nature of this parameter and its reduced sensitivity to network architecture. Overall, DIMOND-KAN provides more reliable and accurate recovery of DTI metrics than competing methods.

3.4. NODDI metrics from self-supervised DIMOND-KAN

DIMOND-KAN extends its capability to reconstruct NODDI microstructure parameters with high fidelity (Fig. 7). The predicted f_{ic}, f_{iso}, and ODI maps closely match the reference maps obtained from conventional NODDI fitting, faithfully capturing key anatomical features. The f_{ic} maps reflect elevated intracellular volume fraction in major white matter tracts and reduced values in cortical gray matter and CSF-adjacent regions. The f_{iso}maps demonstrate strong agreement with the reference, accurately delineating regions of high isotropy such as ventricles and sulci. Likewise, the ODI maps reproduce the expected spatial distribution of fiber dispersion, with higher values in regions of complex crossing fibers and lower values in coherent bundles. The residual maps for DIMOND-KAN are of low amplitude, dominated by randomly distributed noise without structured artifacts, underscoring the accuracy and anatomical plausibility of its predictions. Overall, the results highlight the robustness of DIMOND-KAN in recovering reliable NODDI parameters under self-supervised learning.

Quantitatively, self-supervised DIMOND-KAN exhibits consistent gains in predicting NODDI parameters (Table 4). DIMOND-KAN achieves an MAE of 0.0394 for f_{ic}, representing a modest 0.25% improvement over the NODDI-toolbox, which is comparable to the 0.51% improvement obtained by DIMOND-CNN. More pronounced advances are observed for f_{iso}, where DIMOND-KAN reduces the error to 0.0294, corresponding to an 8.41% improvement and outperforming the 5.61% reduction achieved by DIMOND-CNN. The largest gain is seen for the orientation dispersion index, where DIMOND-KAN attains an MAE of 0.0576, marking a 14.16% improvement over the baseline and slightly surpassing DIMOND-CNN's 13.11%. The relatively small benefit for f_{ic} across both networks suggests that this parameter is less sensitive to architectural design. Overall, DIMOND-KAN delivers reliable performance with clear advantages for $f_{iso}$and ODI estimation.

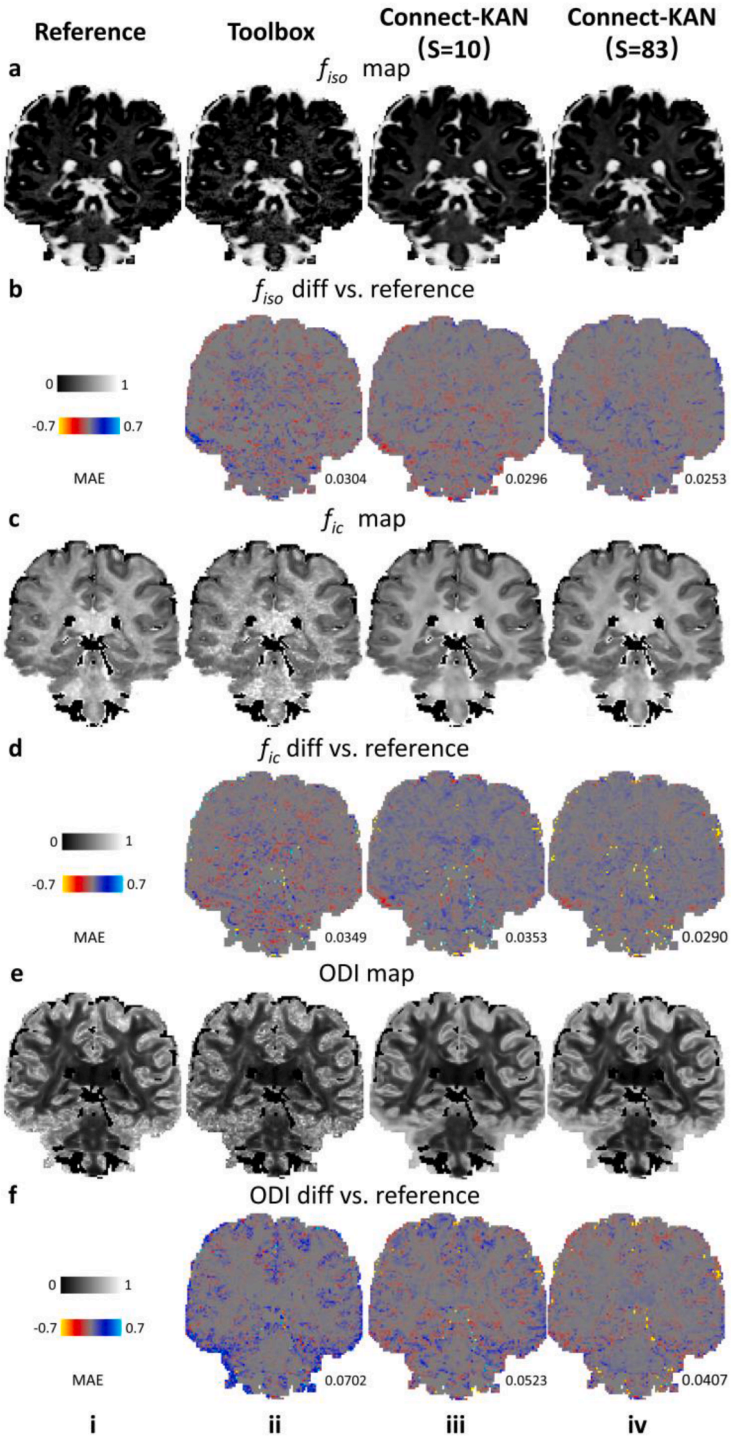

Fig. 12. NODDI metric maps across different training dataset size. Representative coronal maps for isotropic volume fraction (f_{iso}), intracellular volume fraction (f_{ic}), and orientation dispersion index (ODI) of reference (i), and those generated using NODDI-toolbox (ii), supervised Connect-KAN trained on 10 subjects (iii), and supervised Connect-KAN trained on 83 subjects (iv) are shown. The corresponding residual maps relative to reference maps are displayed (b, d, f, h) with the mean absolute error (MAE) reported.

3.5. Bias-variance analysis

For most DTI and NODDI metrics, both supervised and self-supervised models exhibit near-zero median error, indicating negligible systematic bias. The performance gains are primarily associated with reduced dispersion of the error distribution, reflecting decreased variance from the networks' denoising capability. Notably, for RD and ODI, the self-supervised regime demonstrates a slight systematic offset relative to the supervised counterpart, suggesting a trade-off between variance regularization and bias correction under subject-specific training (Fig. 8).

3.6. Cross-dataset generalization evaluation

Our methods demonstrate decent generalizability across datasets (Fig. 9). The conventional fitting method using FSL presents pronounced

Table 6
Effect of training set size on supervised DTI parameter estimation. Group mean absolute errors (MAE ± standard deviation) across 10 HCP test subjects are reported for Connect-KAN models trained using the configuration of (i) 10 subjects with six DWIs, (ii) 83 subjects with six DWIs, and (iii) 83 subjects with 32 DWIs, respectively. Relative MAE decrease is computed as the percentage reduction relative to the baseline OLS implemented in FSL using N = 6 diffusion directions. The red text highlights the lowest MAEs while the blue text highlights the second lowest MAEs.

DTI param	Methods Metrics	FSL (N = 6)	FSL (N = 32)	Connect-KAN (10 subjects) (N = 6)	Connect-KAN (83 subjects) (N = 6)	Connect-KAN (83 subjects) (N = 32)
FA	**MAE**	0.2103	0.0417	0.0449±0.0032	0.0370±0.0026	0.0298±0.0029
	Decrease	±0.0150	±0.0031	77.70%	82.41%	85.83%
		Baseline	80.17%			
AD ($\mu m^2 ms^{-1}$)	**MAE**	0.2466	0.0826	0.0678±0.0063	0.0556±0.0020	0.0457±0.0048
	Decrease	±0.0184	±0.0039	68.73%	77.45%	81.47%
		Baseline	66.50%			
RD ($\mu m^2 ms^{-1}$)	**MAE**	0.1196	0.0597	0.0546±0.0112	0.0415±0.0029	0.0362±0.0039
	Decrease	±0.0078	±0.0032	59.95%	65.30%	69.73%
		Baseline	50.08%			
MD ($\mu m^2 ms^{-1}$)	**MAE**	0.0743	0.0611	0.0491±0.0103	0.0379±0.0026	0.0339±0.0043
	Decrease	±0.0039	±0.0036	45.36%	48.99%	54.37%
		Baseline	17.77%			

Table 7
Effect of training dataset size on supervised NODDI parameter estimation. Group mean absolute errors (MAE ± standard deviation) across 10 HCP test subjects are reported for Connect-KAN models trained using data of 10 and 83 subjects, respectively. Relative MAE decrease is computed as the percentage reduction relative to the baseline NODDI-toolbox. The red text highlights the lowest MAEs while the blue text highlights the second lowest MAEs.

NODDI param	Methods Metrics	Toolbox	Connect-KAN (10 subjects)	Connect-KAN (83 subjects)
f_{ic}	**MAE**	0.0395	0.0400 ± 0.0069	0.0342±0.0074
	Decrease	±0.0024	−1.27%	13.42%
		Baseline		
f_{iso}	**MAE**	0.0321	0.0316 ± 0.0015	0.0294±0.0024
	Decrease	±0.0018	1.56%	8.41%
		Baseline		
ODI	**MAE**	0.0671	0.0540±0.0017	0.0430±0.0016
	Decrease	±0.0041	19.52%	35.92%
		Baseline		

Table 8
Accuracy of DTI parameter approximation on the Simulated HCP dataset. Mean absolute errors (MAE ± standard deviation) of DTI metrics, including fractional anisotropy (FA), axial diffusivity (AD), radial diffusivity (RD), and mean diffusivity (MD), between ground-truth values and those generated from Connect-CNN and Connect-KAN with comparable model capacity are listed. Relative MAE reduction is computed relative to the Connect-CNN baseline.

DTI param	Methods Metrics	Connect-CNN	Connect-KAN
FA	**MAE**	0.0515±0.0014	0.0384±0.0048
	Decrease	**Baseline**	25.44%
AD ($\mu m^2 ms^{-1}$)	**MAE**	0.0691±0.0024	0.0544±0.0110
	Decrease	**Baseline**	21.27%
RD ($\mu m^2 ms^{-1}$)	**MAE**	0.0507±0.0025	0.0501±0.0162
	Decrease	**Baseline**	1.18%
MD ($\mu m^2 ms^{-1}$)	**MAE**	0.0423±0.0029	0.0519±0.0161
	Decrease	**Baseline**	−22.70%

structural deviations and noise artifacts in regions of high anisotropy and at gray-white matter interfaces. The overestimation of FA and RD likely results from the extremely limited number of directions (i.e., six) that might exhibit a slightly biased distribution in q-space, even though we optimized the condition number. Even directly transferring the Connect-KAN weights pretrained on the HCP dataset without adaptation (HCP-pretrained) produces plausible DTI parameter fitting. However, systematic errors remain evident due to cross-dataset distributional shifts caused by the massively different imaging protocols (e.g., spatial resolution, TE, TR) between the Tsinghua dataset and the HCP dataset.

Table 9
Accuracy of NODDI parameter approximation on the Simulated HCP dataset. Mean absolute errors (MAE ± standard deviation) of NODDI metrics, including intracellular volume fraction (f_{ic}), isotropic volume fraction (f_{iso}), and orientation dispersion index (ODI), between ground-truth values and those generated from Connect-CNN and Connect-KAN with comparable model capacity are listed. Relative MAE reduction is computed relative to the Connect-CNN baseline.

NODDI param	Methods Metrics	Connect-CNN	Connect-KAN
f_{ic}	**MAE**	0.0572±0.0034	0.0529±0.0034
	Decrease	**Baseline**	7.52%
f_{iso}	**MAE**	0.0555±0.0030	0.0515±0.0027
	Decrease	**Baseline**	7.21%
ODI	**MAE**	0.0605±0.0016	0.0580±0.0016
	Decrease	**Baseline**	4.13%

Fine-tuning the pretrained weights on nine subjects effectively mitigates such errors, with the details of FA and RD maps further aligned with high-quality references. The self-supervised DIMOND-KAN trained in a dataset-specific manner without requiring high-quality reference achieves consistently robust performance on the Tsinghua dataset compared to the HCP dataset. Compared to HCP pretrained net and fine-tuning results, DIMOND-KAN results appear noisier but less biased. This indicates the supervised methods are effective in reducing variance by harnessing the high training data quality, but might be more prone to bias due to cross-dataset difference. Self-supervised method shows more flexible adaptation across datasets with minimal bias, but the results might suffer from higher variance and noise level due to the lack of high-quality data for supervision. Full supervision with 138 samples trained from scratch delivers consistently excellent performance, achieving accurate estimation with high-SNR and tissue contrast that are similar to the reference parameter maps, serving as an upper-bound when sufficient high-quality training data are available.

The quantitative results on the method's generalizability agree with the qualitative results (Table 5). Direct transfer of HCP-pretrained Connect-KAN weights results in biased estimations, offering only a slight improvement in FA while markedly increasing the errors of AD, RD, and MD, which clearly reflects the adverse effect of cross-domain mismatch. Fine-tuning on nine annotated subjects substantially alleviates this bias, producing more balanced estimations across all metrics and reducing the FA error to 0.0672, a level notably lower than the baseline and close to the 0.0743 achieved by the self-supervised approach. DIMOND-KAN, trained without any annotations, demonstrates strong adaptability to distribution shifts, achieving relative MAE reductions of 21.4%, 19.0%, 18.6%, and 13.8% in FA, AD, RD, and MD, respectively. While its outputs tend to be noisier with higher variance,

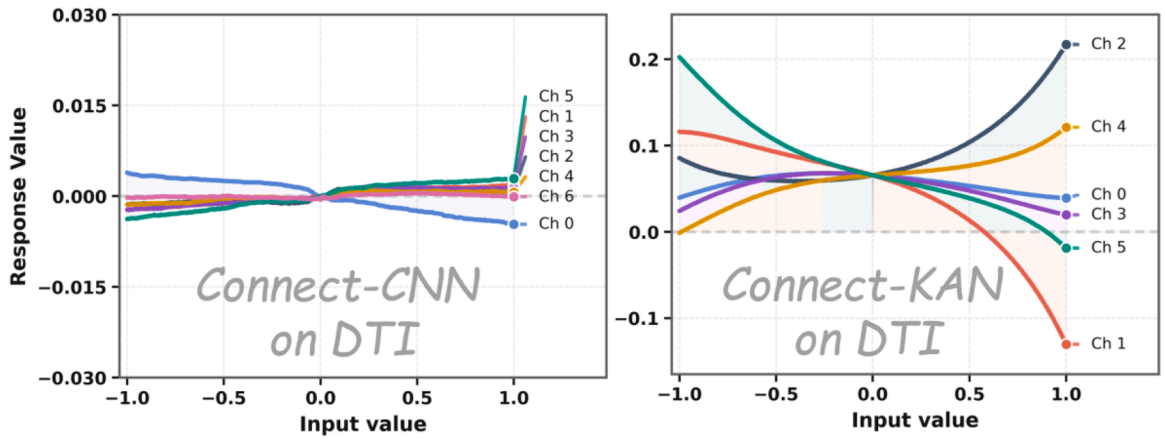

Fig. 13. Visualization of learned nonlinear response functions for DTI. Representative response functions from the final convolutional stage are shown for Connect-CNN and Connect-KAN in the supervised DTI task. The plotted curves illustrate the dependence of the learned response on the normalized input value for different channels associated with one output channel. Connect-CNN responses remain comparatively close to weakly nonlinear trends, whereas Connect-KAN responses exhibit more diverse and curved nonlinear shapes.

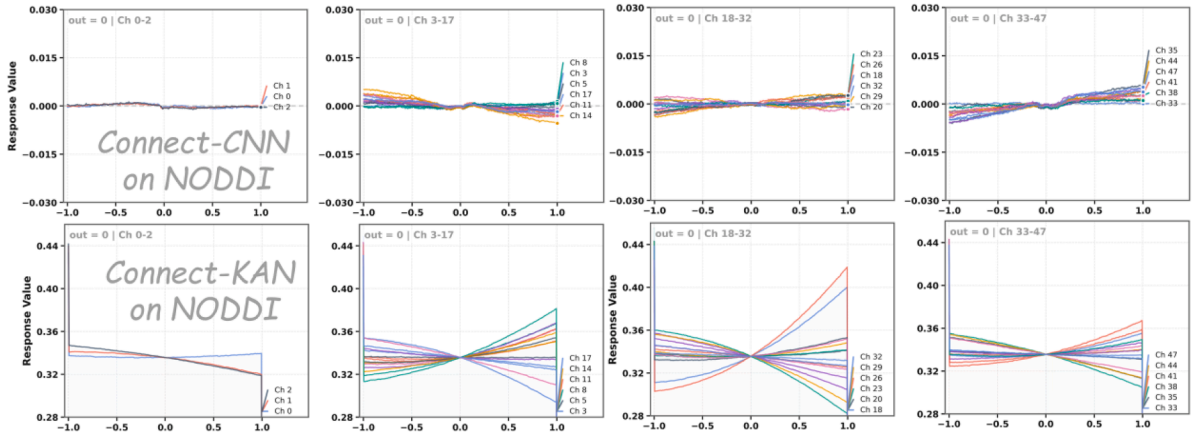

Fig. 14. Visualization of learned nonlinear response functions for NODDI. Representative response functions from the final convolutional stage are shown for Connect-CNN and Connect-KAN in the supervised NODDI task. Because the NODDI input contains 48 channels, the responses are displayed in grouped panels. For clarity, the three b = 0 channels and six representative channels from each of the b = 1000, 2000, and 3000 s/mm^2 shells are shown.

they remain largely unbiased. Finally, full supervision with 138 annotated samples consistently delivers the best performance, with FA reduced to 0.0407, corresponding to a 56.9% decrease relative to the baseline.

3.7. Impact of diffusion encoding directions and training dataset size

Compared with conventional tensor fitting using richer diffusion sampling, supervised Connect-KAN still demonstrates superior estimation accuracy under the highly under-sampled setting. When the number of diffusion directions is increased from N = 6 to N = 32, the quality of DTI metrics generated by OLS implemented in FSL improved substantially, exhibiting reduced noise and more stable parameter estimation (Fig. 10). Nevertheless, Connect-KAN trained on a subset of six DWIs produces parameter maps that are visually closer to the reference, exhibiting lower residual dispersion and superior preservation of anatomical structures. Importantly, Connect-KAN is not limited to sparse acquisitions. When scaled to the N = 32 directions, the network effectively exploits the richer angular information to further enhance fidelity and provide sharper delineations of complex white matter architectures.

Beyond its robustness to angular sampling patterns, Connect-KAN also exhibits remarkable data efficiency. Even when trained on a severely restricted cohort of only 10 subjects, Connect-KAN successfully recovers higher-quality DTI metrics compared to conventional OLS in FSL, underscoring the strong generalization capacity of its learned feature representations (Fig. 11). A similar trend is observed in NODDI model mapping, where Connect-KAN trained on just 10 subjects outperforms the conventional method implemented in the NODDI-toolbox (Fig. 12).

These advantages are further supported by the quantitative results (Table 6). Specifically, the model trained on 83 subjects achieves lower mean absolute errors (MAEs) than FSL (N = 32) across all DTI metrics, while the model trained on 10 subjects remains competitive and surpasses FSL (N = 32) in several metrics. Furthermore, increasing the number of diffusion encoding directions from N = 6 to N = 32 yields substantial performance gains for Connect-KAN. For instance, the MAE of FA decreases from 0.0370 to 0.0298, representing a 19.5% relative improvement. Similarly, for the NODDI model, Connect-KAN trained on 10 subjects yields a lower MAE than the NODDI-toolbox in terms of $f_{iso}$and ODI (Table 7).

3.8. Network approximation capability comparison on simulation data

Connect-KAN exhibited metric-dependent approximation advantages in noise-free simulation experiments. On the simulated HCP

dataset, Connect-KAN reduced the MAE of FA and AD from 0.0515 to 0.0384 and from 0.0691 to 0.0544, corresponding to relative reductions of 25.44% and 21.27%, respectively, while achieving a comparable RD error with a modest reduction of 1.18% (Table 8). However, its MD error was higher than that of Connect-CNN, indicating that spline-based nonlinear modeling does not uniformly improve all tensor-derived scalar metrics and may involve metric-specific trade-offs for averaged diffusivity measures. In contrast, for the more nonlinear NODDI model, Connect-KAN consistently outperformed Connect-CNN across f_{ic}, f_{iso}, and ODI, yielding MAE reductions of 7.52%, 7.21%, and 4.13%, respectively (Table 9). These results suggest that Connect-KAN enhances nonlinear approximation capability particularly for anisotropy-sensitive and compartmental diffusion metrics, while its benefits vary across derived scalar measures.

4. Discussion

This study introduces MicroKAN, a unified deep learning framework specifically designed for diffusion MRI microstructure modeling. By jointly supporting both supervised and self-supervised paradigms, MicroKAN adapts flexibly to varying data availability scenarios. When high-quality paired dMRI signals and diffusion metrics are available, Connect-KAN maps the input diffusion signals to target model parameters under full supervision. For single-subject or label-free data, DIMOND-KAN leverages physics-informed constraints and forward diffusion modeling to learn network parameters in a self-supervised manner. Across extensive experiments on the HCP dataset, Connect-KAN consistently outperformed conventional CNNs, the classical q-DL baseline, and tensor fitting, achieving lower MAEs across FA, AD, RD, and MD, as well as improved estimation of NODDI parameters. In the self-supervised setting, DIMOND-KAN also demonstrated substantial advantages, reducing FA and V1 errors by >76% and 56%, respectively, compared to traditional fitting, while improving $f_{iso}$and ODI estimation relative to DIMOND-CNN. These findings establish the superior performance and robustness of MicroKAN across both supervised and self-supervised regimes.

The efficacy and efficiency of MicroKAN is primarily derived from the Bottleneck-KAN—Conv modules, which provide compact yet highly expressive nonlinear modeling. Unlike conventional CNNs with fixed activations, spline-based activations dynamically learn nonlinear functional forms, enabling more accurate approximation of the complex nonlinear relationships inherent in dMRI signals without excessive network depth. Empirical ablations confirm that both skip-connections and spline activations contribute critically to performance (Table 1, and Table 2). The visualization of representative response functions further supports the stronger nonlinear mapping capacity provided by spline-based formulation (Figs. 13, and 14). Under the same normalized input perturbation range, Connect-KAN produced response functions with a much wider dynamic range and more evident curvature than Connect-CNN, as reflected by the markedly different y-axis scales. While the Connect-CNN responses remained close to zero and were nearly linear, the Connect-KAN responses showed larger response variations and channel-dependent nonlinear bending. Beyond the representational prowess, the Bottleneck-KAN—Conv module improves the fitting capacity while still harnessing the high computational efficiency of trained neural network. Connect-KAN reduces NODDI parameter estimation time from >12 h (toolbox-based fitting) to 1 min 22 s inference time. Even for the self-supervised DIMOND-KAN, which performs subject-specific training, the training and inference time remains within four hours per subject, still faster than standard voxel-wise numerical optimization-based NODDI implementations (Zhang et al., 2012).

The comprehensive generalization experiments on the Tsinghua dataset establish MicroKAN as a useful method in various scenarios where high-quality training data are insufficient, limited, or even unavailable. In the ideal case when high-quality training data are sufficient, supervised training using Connect-KAN consistently leads to the best performance on both HCP and Tsinghua datasets. Similarly, DIMOND-KAN produces consistent and robust performance on these datasets without requiring any high-quality labels. In the scenario where limited high-quality reference data are available, we introduce the strategy of transfer learning by fine-tuning the model pretrained on HCP using a limited number of subjects. Both DIMOND-KAN and transfer learning can be used in this scenario. The choice between them should be guided by the characteristics of the application. DIMOND-KAN, characterized by minimal systematic bias but higher variance, is better suited for high-SNR data where unbiased estimations are critical. Conversely, transfer learning typically yields lower variance at the cost of potential bias, making it preferable for low-SNR data dominated by noise or for datasets with distributions closely aligned with the pretraining domain. Furthermore, transfer learning offers a marked advantage in computational efficiency, substantially reducing training time (e.g., from 9 h and 30 min to 35 min). A promising future direction lies in pretraining KAN models on larger and more diverse datasets such as HCP, UK Biobank, and ADNI, where exposure to heterogeneous populations and acquisition protocols is expected to substantially enhance the generalizability and robustness of transfer learning, thereby enabling reliable adaptation across scanners, sites, and clinical cohorts with only minimal fine-tuning.

Beyond methodological innovations, MicroKAN has the potential to significantly influence both clinical practice and large-scale population studies. By reducing acquisition requirements from dozens of diffusion directions to only six carefully optimized directions, MicroKAN enables high-quality diffusion tensor imaging in patient groups who cannot tolerate prolonged scanning, such as critically ill individuals or children, thereby improving diagnostic accessibility for neurological disorders and focal lesions. At the same time, the substantial acceleration in model fitting, reducing computation from several hours to only minutes for tensor and NODDI reconstruction, allows advanced microstructural modeling to be applied to large datasets that were previously impractical due to excessive computational demands. In the era of big data, such efficiency gains are particularly valuable, as they make it possible to extract diffusion-derived phenotypes at scale with markedly lower computational costs. This scalability not only supports population-level neuroimaging research but also facilitates the integration of microstructural biomarkers into real-time precision medicine workflows.

Despite these promising results, several limitations warrant attention. First, the current architecture employs similar design choices across different diffusion models, which may not be optimal. For instance, tensor models might require less expressive architectures than NODDI. Second, hyperparameters related to spline activation, such as knot placement and density, also require systematic comprehensive optimization. Automated architecture search and hyperparameter tuning may potentially address these challenges. Third, extending MicroKAN to additional models such as diffusion kurtosis imaging and multi-compartment microstructure models would further demonstrate its versatility. Finally, Connect-KAN requires higher GPU memory than conventional CNN during training (Supplementary Table S5), which would be addressed through efficient operator, such as optimized CUDA kernels.

5. Conclusion

In this work, we present MicroKAN, a unified deep learning framework that redefines diffusion MRI microstructure modeling through convolutional Kolmogorov-Arnold Networks with adaptive spline activations. By seamlessly bridging supervised and self-supervised paradigms, MicroKAN adapts to both richly annotated and fully unlabeled scenarios. Extensive experiments on the HCP dataset demonstrate that Connect-KAN achieves state-of-the-art accuracy, reducing DTI errors by up to 83% and NODDI errors by nearly 50% compared with conventional fitting and CNN baselines. Meanwhile, DIMOND-KAN introduces physics-informed self-supervision to recover microstructural parameters

without ground-truth labels, cutting FA and V1 errors by >76% and 56%, respectively. Cross-dataset validation on an in-house Tsinghua dataset further highlights MicroKAN's robustness, revealing clear trade-offs between bias and variance under zero-shot transfer, few-shot fine-tuning, and large-scale full supervision. Collectively, these results establish MicroKAN as a powerful, scalable, and clinically practical paradigm for accelerated and accurate diffusion MRI analysis, with broad potential to advance population-scale neuroimaging studies and enable next-generation microstructural biomarkers in clinical research.

Ethics statement

This study involves human participants and was approved by the Institution Review Board of Tsinghua University (Project No. 20230002). Given the retrospective nature of the study, the requirement for informed consent is waived. All procedures are conducted in accordance with the Declaration of Helsinki and relevant ethical guidelines for medical research.

Data and code availability

The diffusion MRI data from 93 subjects used in the DTI and NODDI studies were provided by the Human Connectome Project WU-Minn-Ox Consortium and are publicly available via public database at https://www.humanconnectome.org. The source code of supervised learning Connect-KAN and self-supervised learning DIMOND-KAN architectures was implemented using PyTorch and MATLAB, and is available on GitHub at https://github.com/JustlfC03/MicroKAN.

CRediT authorship contribution statement

Yifei Chen: Writing – original draft, Methodology, Conceptualization, Writing – review & editing. **Zihan Li:** Writing – original draft, Methodology. **Yuanhan Wang:** Visualization, Validation, Methodology. **Yijin Li:** Visualization, Validation. **Jialan Zheng:** Visualization, Validation. **Hongjia Yang:** Visualization, Validation. **Mingxuan Liu:** Visualization, Validation. **Tolga Cukur:** Writing – review & editing. **Qiuyun Fan:** Writing – review & editing. **Ziyu Li:** Writing – review & editing, Supervision. **Jie Lu:** Writing – review & editing. **Qiyuan Tian:** Writing – review & editing, Project administration, Conceptualization, Funding acquisition, Methodology, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (grant number 82302166), the Tsinghua University Startup Fund, and the Dushi Program (grant number 20241080026). Preliminary results from this project were presented in a poster abstract at the annual scientific meeting of the International Society of Magnetic Resonance in Medicine (ISMRM) in May 2025 and in a poster abstract at the ISMRM Workshop on 40 Years of Diffusion: Past, Present & Future Perspectives in February 2025.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.neuroimage.2026.122032.

References

Alfaro-Almagro, F., Jenkinson, M., Bangerter, N.K., Andersson, J.L.R., Griffanti, L., Douaud, G., Sotiropoulos, S.N., Jbabdi, S., Hernandez-Fernandez, M., Vallee, E., Vidaurre, D., Webster, M., McCarthy, P., Rorden, C., Daducci, A., Alexander, D.C., Zhang, H., Dragonu, I., Matthews, P.M., Miller, K.L., Smith, S.M., 2018. Image processing and quality control for the first 10,000 brain imaging datasets from UK Biobank. NeuroImage 166, 400–424. https://doi.org/10.1016/j.neuroimage.2017.10.034.

Aliotta, E., Nourzadeh, H., Sanders, J., Muller, D., Ennis, D.B., 2019. Highly accelerated, model-free diffusion tensor MRI reconstruction using neural networks. Med. Phys. 46, 1581–1591. https://doi.org/10.1002/mp.13400.

Al-Sharif, N.B., Zavaliangos-Petropulu, A., Narr, K.L., 2025. A review of diffusion MRI in mood disorders: mechanisms and predictors of treatment response. Neuropsychopharmacology 50, 211–229. https://doi.org/10.1038/s41386-024-01894-3.

Andersson, J.L.R., Skare, S., Ashburner, J., 2003. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. NeuroImage 20, 870–888. https://doi.org/10.1016/S1053-8119(03)00336-7.

Bertleff, M., Domsch, S., Weingärtner, S., Zapp, J., O'Brien, K., Barth, M., Schad, L.R., 2017. Diffusion parameter mapping with the combined intravoxel incoherent motion and kurtosis model using artificial neural networks at 3 T. NMR Biomed 30, e3833. https://doi.org/10.1002/nbm.3833.

Chen, G., Hong, Y., Huynh, K.M., Yap, P.-T., 2023. Deep learning prediction of diffusion MRI data with microstructure-sensitive loss functions. Med. Image Anal. 85, 102742. https://doi.org/10.1016/j.media.2023.102742.

Chen, Y., Zhu, Z., Zhu, S., Qiu, L., Zou, B., Jia, F., Zhu, Y., Zhang, C., Fang, Z., Qin, F., Fan, J., Wang, C., Yu, G., Gao, Y., 2025. SCKansformer: fine-grained classification of bone marrow cells via Kansformer backbone and hierarchical attention mechanisms. IEEE J. Biomed. Health Inform. 29, 558–571. https://doi.org/10.1109/JBHI.2024.3471928.

Chiari-Correia, R.D., Tumas, V., Santos, A.C., Salmon, C.E.G., 2023. Structural and functional differences in the brains of patients with MCI with and without depressive symptoms and their relations with Alzheimer's disease: an MRI study. Psychoradiology 3, kkad008. https://doi.org/10.1093/psyrad/kkad008.

Daducci, A., Canales-Rodríguez, E.J., Zhang, H., Dyrby, T.B., Alexander, D.C., Thiran, J.-P., 2015. Accelerated microstructure imaging via Convex optimization (AMICO) from diffusion MRI data. NeuroImage 105, 32–44. https://doi.org/10.1016/j.neuroimage.2014.10.026.

De Almeida Martins, J.P., Nilsson, M., Lampinen, B., Palombo, M., While, P.T., Westin, C.-F., Szczepankiewicz, F., 2021. Neural networks for parameter estimation in microstructural MRI: application to a diffusion-relaxation model of white matter. NeuroImage 244, 118601. https://doi.org/10.1016/j.neuroimage.2021.118601.

Drokin, I., 2024. Kolmogorov-Arnold convolutions: design principles and empirical studies. https://doi.org/10.48550/arXiv.2407.01092.

Fujiyoshi, K., Hikishima, K., Nakahara, J., Tsuji, O., Hata, J., Konomi, T., Nagai, T., Shibata, S., Kaneko, S., Iwanami, A., Momoshima, S., Takahashi, S., Jinzaki, M., Suzuki, N., Toyama, Y., Nakamura, M., Okano, H., 2016. Application of q -space diffusion MRI for the visualization of white matter. J. Neurosci. 36, 2796–2808. https://doi.org/10.1523/JNEUROSCI.1770-15.2016.

Ge, R., Yu, X., Chen, Y., Jia, F., Zhu, S., Zeng, D., Wang, C., Liu, Q., Niu, S., 2026. TC-KANRecon: High-quality and accelerated MRI reconstruction via adaptive KAN mechanisms and intelligent feature scaling. IEEE J. Biomed. Health Inform. 30, 4200–4213. https://doi.org/10.1109/JBHI.2025.3626343.

Gibbons, E.K., Hodgson, K.K., Chaudhari, A.S., Richards, L.G., Majersik, J.J., Adluru, G., DiBella, E.V.R., 2019. Simultaneous NODDI and GFA parameter map generation from subsampled q-space imaging using deep learning. Magn. Reson. Med. 81, 2399–2411. https://doi.org/10.1002/mrm.27568.

Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M., 2013. The minimal preprocessing pipelines for the human connectome project. NeuroImage 80, 105–124. https://doi.org/10.1016/j.neuroimage.2013.04.127.

Golkov, V., Dosovitskiy, A., Sperl, J.I., Menzel, M.I., Czisch, M., Samann, P., Brox, T., Cremers, D., 2016. Q-space deep learning: twelve-fold shorter and model-free diffusion MRI scans. IEEE Trans. Med. Imaging 35, 1344–1351. https://doi.org/10.1109/TMI.2016.2551324.

Hernández, M., Guerrero, G.D., Cecilia, J.M., García, J.M., Inuggi, A., Jbabdi, S., Behrens, T.E.J., Sotiropoulos, S.N., 2013. Accelerating fibre orientation estimation from diffusion weighted magnetic resonance imaging using GPUs. PLoS ONE 8, e61892. https://doi.org/10.1371/journal.pone.0061892.

Jack Jr., C.R., Bernstein, M.A., Fox, N.C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P.J., Whitwell, J.L., Ward, C., Dale, A.M., Felmlee, J.P., Gunter, J.L., Hill, D.L.G., Killiany, R., Schuff, N., Fox-Bosetti, S., Lin, C., Studholme, C., DeCarli, C.S., Krueger, G., Ward, H.A., Metzger, G.J., Scott, K.T., Mallozzi, R., Blezek, D., Levy, J., Debbins, J.P., Fleisher, A.S., Albert, M., Green, R., Bartzokis, G., Glover, G., Mugler, J., Weiner, M.W., 2008. The Alzheimer's disease neuroimaging initiative (ADNI): MRI methods. J. Magn. Reson. Imaging 27, 685–691. https://doi.org/10.1002/jmri.21049.

Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M., 2012. FSL. NeuroImage 62, 782–790. https://doi.org/10.1016/j.neuroimage.2011.09.015.

Jiang, J., Ferraro, S., Zhao, Y., Wu, B., Lin, J., Chen, T., Gao, J., Li, L., 2024. Common and divergent neuroimaging features in major depression, posttraumatic stress disorder, and their comorbidity. Psychoradiology 4, kkae022. https://doi.org/10.1093/psyrad/kkae022.

Jokinen, H., Schmidt, R., Ropele, S., Fazekas, F., Gouw, A.A., Barkhof, F., Scheltens, P., Madureira, S., Verdelho, A., Ferro, J.M., Wallin, A., Poggesi, A., Inzitari, D.,

Pantoni, L., Erkinjuntti, T., LADIS Study Group, 2013. Diffusion changes predict cognitive and functional outcome: the LADIS study. Ann. Neurol. 73, 576–583. https://doi.org/10.1002/ana.23802.

Karimi, D., Jaimes, C., Machado-Rivas, F., Vasung, L., Khan, S., Warfield, S.K., Gholipour, A., 2021a. Deep learning-based parameter estimation in fetal diffusion-weighted MRI. NeuroImage 243, 118482. https://doi.org/10.1016/j.neuroimage.2021.118482.

Karimi, D., Vasung, L., Jaimes, C., Machado-Rivas, F., Warfield, S.K., Gholipour, A., 2021b. Learning to estimate the fiber orientation distribution function from diffusion-weighted MRI. NeuroImage 239, 118316. https://doi.org/10.1016/j.neuroimage.2021.118316.

Kingsley, P.B., 2006. Introduction to diffusion tensor imaging mathematics: part III. Tensor calculation, noise, simulations, and optimization. Concepts Magn. Reson. Part A 28A, 155–179. https://doi.org/10.1002/cmr.a.20050.

Kubicki, M., Park, H., Westin, C.F., Nestor, P.G., Mulkern, R.V., Maier, S.E., Niznikiewicz, M., Connor, E.E., Levitt, J.J., Frumin, M., Kikinis, R., Jolesz, F.A., McCarley, R.W., Shenton, M.E., 2005. DTI and MTR abnormalities in schizophrenia: analysis of white matter integrity. NeuroImage 26, 1109–1118. https://doi.org/10.1016/j.neuroimage.2005.03.026.

Le Bihan, D., Breton, E., Lallemand, D., Grenier, P., Cabanis, E., Laval-Jeantet, M., 1986. MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. Radiology 161, 401–407. https://doi.org/10.1148/radiology.161.2.3763909.

Le Bihan, D., Mangin, J., Poupon, C., Clark, C.A., Pappata, S., Molko, N., Chabriat, H., 2001a. Diffusion tensor imaging: concepts and applications. J. Magn. Reson. Imaging 13, 534–546. https://doi.org/10.1002/jmri.1076.

Le Bihan, D., Mangin, J., Poupon, C., Clark, C.A., Pappata, S., Molko, N., Chabriat, H., 2001b. Diffusion tensor imaging: concepts and applications. J. Magn. Reson. Imaging 13, 534–546. https://doi.org/10.1002/jmri.1076.

Li, Z., Gong, T., Lin, Z., He, H., Tong, Q., Li, C., Sun, Y., Yu, F., Zhong, J., 2019. Fast and robust diffusion kurtosis parametric mapping using a three-dimensional convolutional neural network. IEEE Access 7, 71398–71411. https://doi.org/10.1109/ACCESS.2019.2919241.

Li, Z., Li, Z., Bilgic, B., Lee, H., Ying, K., Huang, S.Y., Liao, H., Tian, Q., 2024. DIMOND: diffusion model optimization with deep learning. Adv. Sci. 11, 2307965. https://doi.org/10.1002/advs.202307965.

Liang, Z., Arefin, T.M., Lee, C.H., Zhang, J., 2023. Using mesoscopic tract-tracing data to guide the estimation of fiber orientation distributions in the mouse brain from diffusion MRI. NeuroImage 270, 119999. https://doi.org/10.1016/j.neuroimage.2023.119999.

Lin, Z., Gong, T., Wang, K., Li, Z., He, H., Tong, Q., Yu, F., Zhong, J., 2019. Fast learning of fiber orientation distribution function for MR tractography using convolutional neural network. Med. Phys. 46, 3101–3116. https://doi.org/10.1002/mp.13555.

Liu, Z., Wang, Y., Vaidya, S., Ruehle, F., Halverson, J., Soljačić, M., Hou, T.Y., Tegmark, M., 2025. KAN: Kolmogorov–Arnold Networks. The Thirteenth International Conference on Learning Representations.https://openreview.net/forum?id=Ozo7qJ5vZi.

Miller, K.L., Alfaro-Almagro, F., Bangerter, N.K., Thomas, D.L., Yacoub, E., Xu, J., Bartsch, A.J., Jbabdi, S., Sotiropoulos, S.N., Andersson, J.L.R., Griffanti, L., Douaud, G., Okell, T.W., Weale, P., Dragonu, I., Garratt, S., Hudson, S., Collins, R., Jenkinson, M., Matthews, P.M., Smith, S.M., 2016. Multimodal population brain imaging in the UK biobank prospective epidemiological study. Nat. Neurosci. 19, 1523–1536. https://doi.org/10.1038/nn.4393.

Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J.M., Matthews, P.M., 2004. Advances in functional and structural MR image analysis and implementation as FSL. NeuroImage 23, S208–S219. https://doi.org/10.1016/j.neuroimage.2004.07.051.

Sotiropoulos, S.N., Jbabdi, S., Xu, J., Andersson, J.L., Moeller, S., Auerbach, E.J., Glasser, M.F., Hernandez, M., Sapiro, G., Jenkinson, M., Feinberg, D.A., Yacoub, E., Lenglet, C., Van Essen, D.C., Ugurbil, K., Behrens, T.E.J., 2013. Advances in diffusion MRI acquisition and processing in the human connectome project. NeuroImage 80, 125–143. https://doi.org/10.1016/j.neuroimage.2013.05.057.

Tian, Q., Bilgic, B., Fan, Q., Liao, C., Ngamsombat, C., Hu, Y., Witzel, T., Setsompop, K., Polimeni, J.R., Huang, S.Y., 2020. DeepDTI: high-fidelity six-direction diffusion tensor imaging using deep learning. NeuroImage 219, 117017. https://doi.org/10.1016/j.neuroimage.2020.117017.

Uğurbil, K., Xu, J., Auerbach, E.J., Moeller, S., Vu, A.T., Duarte-Carvajalino, J.M., Lenglet, C., Wu, X., Schmitter, S., Van De Moortele, P.F., Strupp, J., Sapiro, G., De Martino, F., Wang, D., Harel, N., Garwood, M., Chen, L., Feinberg, D.A., Smith, S.M., Miller, K.L., Sotiropoulos, S.N., Jbabdi, S., Andersson, J.L.R., Behrens, T.E.J., Glasser, M.F., Van Essen, D.C., Yacoub, E., 2013. Pushing spatial and temporal resolution for functional and diffusion MRI in the human connectome project. NeuroImage 80, 80–104. https://doi.org/10.1016/j.neuroimage.2013.05.012.

Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K., 2013. The WU-Minn human connectome project: an overview. NeuroImage 80, 62–79. https://doi.org/10.1016/j.neuroimage.2013.05.041.

Van Everdingen, K.J., Van Der Grond, J., Kappelle, L.J., Ramos, L.M.P., Mali, W.P.T.M., 1998. Diffusion-weighted magnetic resonance imaging in acute stroke. Stroke 29, 1783–1790. https://doi.org/10.1161/01.STR.29.9.1783.

Xu, M., Zhang, X., Wang, Y., Ren, L., Wen, Z., Xu, Y., Gong, G., Xu, N., Yang, H., 2012. Probabilistic brain Fiber tractography on GPUs. In: 2012 IEEE 26th International Parallel and Distributed Processing Symposium Workshops & PhD Forum. Presented at the 2012 26th IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW). Shanghai, China. IEEE, pp. 742–751. https://doi.org/10.1109/IPDPSW.2012.92.

Ye, C., 2017. Tissue microstructure estimation using a deep network inspired by a dictionary-based framework. Med. Image Anal. 42, 288–299. https://doi.org/10.1016/j.media.2017.09.001.

Yeatman, J.D., Wandell, B.A., Mezer, A.A., 2014. Lifespan maturation and degeneration of human brain white matter. Nat. Commun. 5, 4932. https://doi.org/10.1038/ncomms5932.

Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C., 2012. NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. NeuroImage 61, 1000–1016. https://doi.org/10.1016/j.neuroimage.2012.03.072.