

# Deep Decomposition for Stochastic Normal-Abnormal Transport

Peirong Liu<sup>1</sup> Yueh Lee<sup>2</sup> Stephen Aylward<sup>3</sup> Marc Niethammer<sup>1</sup>

<sup>1</sup>Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, USA

<sup>2</sup>Department of Radiology, University of North Carolina at Chapel Hill, Chapel Hill, USA

<sup>3</sup>Kitware Inc., New York, USA

{peirong, mn}@cs.unc.edu stephen.aylward@kitware.com yueh\_lee@med.unc.edu

## Abstract

*Advection-diffusion equations describe a large family of natural transport processes, e.g., fluid flow, heat transfer, and wind transport. They are also used for optical flow and perfusion imaging computations. We develop a machine learning model,  $D^2$ -SONATA, built upon a stochastic advection-diffusion equation, which predicts the velocity and diffusion fields that drive 2D/3D image time-series of transport. In particular, our proposed model incorporates a model of transport atypicality, which isolates abnormal differences between expected normal transport behavior and the observed transport. In a medical context such a normal-abnormal decomposition can be used, for example, to quantify pathologies. Specifically, our model identifies the advection and diffusion contributions from the transport time-series and simultaneously predicts an anomaly value field to provide a decomposition into normal and abnormal advection and diffusion behavior. To achieve improved estimation performance for the velocity and diffusion-tensor fields underlying the advection-diffusion process and for the estimation of the anomaly fields, we create a 2D/3D anomaly-encoded advection-diffusion simulator, which allows for supervised learning. We further apply our model on a brain perfusion dataset from ischemic stroke patients via transfer learning. Extensive comparisons demonstrate that our model successfully distinguishes stroke lesions (abnormal) from normal brain regions, while reconstructing the underlying velocity and diffusion tensor fields.*

## 1. Introduction

Partial differential equations (PDEs) are used to describe many transport phenomena, e.g., fluid dynamics, heat conduction, and wind dynamics [12]. However, it is expensive to numerically solve PDEs especially in high spatial dimensions and across large timescales [6]. Recent deep learning

approaches have made data-driven solutions to PDEs possible [8, 16, 33, 38, 46, 51, 55, 56].

While the aforementioned deep learning approaches help to speed-up PDE solutions, we are instead interested in *inverse* PDE problems [34, 59]. Specifically, our goal is to estimate the spatially varying velocity and diffusion tensor fields (referred to as advection-diffusion parameters throughout this manuscript) for general advection-diffusion PDEs. Limited work on estimating the parameter fields of advection-diffusion equations exists. Tartakovsky et al. [59] use a deep neural network (DNN) to estimate 2D diffusion fields from diffusion PDEs. Bézenac et al. [12] learn velocity and diffusion fields of advection-diffusion PDEs by DNNs in 2D. Koundal et al. [29] use optimal mass transport combined with a spatially constant diffusion. Optimization approaches were proposed [35, 36, 63] to estimate the advection-diffusion parameters of an advection-diffusion equation in 3D. Though promising, the numerical optimization approach is time-consuming, especially when dealing with large datasets. Moreover, the methods above assume isotropic diffusion (i.e., they do not estimate more general diffusion tensors), which may be insufficient to accurately model complex materials (e.g., anisotropic porous media, brain tissue) where diffusion is mostly anisotropic.

The YETI approach by Liu et al. [37] is a deep learning framework to estimate advection-diffusion parameters from transport image time-series. It addresses some key challenges regarding identifiability (i.e., whether the resulting transport should be attributed to advection or diffusion) and physical constraints (e.g., vector fields should be divergence-free for fluid flow, diffusion tensors should be positive semi-definite (PSD)). Applied on brain perfusion images for stroke lesion detection, YETI achieved clear improvements over existing approaches. However, there is still room for improvement. First, while YETI estimates the advection-diffusion parameters well, it would be useful to quantify the inherent uncertainty in the resulting solutions instead of interpreting the results deterministically. This

This work was supported by the National Institutes of Health (NIH) under award number NIH 2R42NS086295 and NIH R21NS125369.

could be especially important when estimating advection-diffusion parameters from abnormal transport processes, e.g., brain perfusion processes for a stroke patient. Second, YETI predicts advection-diffusion parameters without knowledge of regional anomalies. Hence, additional post-processing is required for lesion detection in real perfusion analysis applications. Furthermore, YETI uses pre-training based on a simulation dataset which only considers normal transport; fine-tuning is then performed using real perfusion data including stroke lesions. This simulation and fine-tuning approach may result in pre-training biased toward normal parameterizations, as abnormal transport data is not simulated. (See the comparisons in Sec. 4.1-4.2.)

We introduce,  $D^2$ -SONATA (**Deep Decomposition for Stochastic Normal-Abnormal Transport**), a new deep-learning-based stochastic model designed to predict the underlying advection-diffusion parameters from observed transport time-series with potential anomalies, in both 2D and 3D.<sup>1</sup> Our main contributions are threefold:

- 1) *A learning-based stochastic advection-diffusion model.*  $D^2$ -SONATA models advection-diffusion processes as a stochastic system. Given a transport process, this stochastic model permits not only reconstructing the transport dynamics with its underlying advection-diffusion parameters, but also captures the epistemic uncertainties with Brownian motion.
- 2) *Representation theorems for anomaly-decomposed divergence-free vector fields and symmetric PSD tensor fields.* Our estimates are grounded in theorems ensuring realistic constraints on the learned advection-diffusion parameter fields *by construction*. The estimation automatically provides a decomposition based on an anomaly value field and the “anomaly-free” advection-diffusion parameters. This provides insights into both the anomaly patterns and what should locally be considered normal advection-diffusion parameters.
- 3) *2D/3D normal-abnormal advection-diffusion dataset.* We develop a simulator for quasi-realistic advection-diffusion that can be used for supervised model pre-training based on the advection-diffusion parameters. Importantly, the simulator is able to generate velocity vector and diffusion tensor fields and applies artificial anomalies. We show that such simulated data boosts model performance for advection-diffusion parameter prediction, in particular for abnormal transport.

## 2. Related work

**Neural Differential Equations and Optical Flow** Significant developments in deep learning have recently led to an explosive growth of deep-learning-based solutions for

<sup>1</sup>Perfusion imaging is the motivating application behind our approach. However, our approach is generally applicable to parameter estimation for any process governed by an advection-diffusion equation.

PDEs. This stream of work either directly models the solution via DNNs [16, 51, 56] or learns mesh-free, infinite-dimensional operators using DNNs [8, 33, 38, 46, 55].

While the above methods focus on solving PDEs, i.e., the *forward* problem, we are interested in solving *inverse* PDE problems [34, 59]. In the context of advection equations such inverse problems have been extensively studied for optical flow and general image registration [25, 58], where the parameter to be estimated is a deformation or a velocity vector field [7, 9, 23, 24, 43]. DNN solutions for the fast prediction of these vector fields have also been studied [5, 14, 54, 62]. In contrast, our goal is the estimation of the parameters for more general advection-diffusion PDEs: specifically, their associated velocity vector and diffusion tensor fields. Further, while optical flow and registration approaches most commonly deal with image pairs, our parameter estimation for advection-diffusion PDEs is based on image *time-series* acquired over multiple time points.

**Perfusion Imaging** Perfusion imaging measures the blood flow through parenchyma by serial imaging [19]. Common perfusion measurement techniques include injecting an intravascular tracer, e.g., Dynamic Susceptibility Contrast-enhanced (DSC) and Dynamic Contrast-Enhanced (DCE) magnetic resonance (MR) perfusion [13, 19, 36], using magnetically-labeled arterial blood water protons as an endogenous tracer (Arterial spin labeling (ASL)) [50], or using positron emission tomography (PET) [21]. Derived quantitative measures help clinical diagnosis and decision-making. So far, the mainstream approach for quantitative perfusion measures is to use tracer kinetic models to estimate hemodynamic parameters and to obtain 3D perfusion parameter maps [19, 44]. However, substantial differences exist in perfusion parameter maps across institutions, mainly caused by different arterial input function (AIF) selection procedures, deconvolution techniques, and interpretations of perfusion parameters [44, 52, 53]. Moreover, these approaches are performed on individual voxels, disregarding spatial dependencies of the injected tracer’s dynamics.

Works to fit tracer transport via PDEs exist, where the observed tracer concentration time-series is assumed to reflect the blood flow within the vessels (advection) while diffusion captures the movements of freely-diffusive tracer within capillaries and the macroscopic effect of capillary transport [11, 11, 22, 36, 57]. However, these works assume that both velocity and diffusion are *constant* over the entire domain, which is unrealistic in real tissue. Zhou et al. [64, 65] and Liu et al. [35, 36] proposed to model the perfusion process via a transport model using numerical optimization. However, Zhou et al. [64, 65] assume the diffusion process is negligible; therefore, only the velocity field is estimated similar to optical flow [9, 25, 48]. Liu et al. [35, 36] and Zhang et al. [63] estimate both the spatial-varying velocity and diffusion fields, yet modeling the diffusion as a

scalar field, which cannot express the diffusion anisotropy.

Liu et al. [37] propose to predict the perfusion process via an advection-diffusion model with *intrinsically constructed* divergence-free velocity vector fields and symmetric positive-semi-definite (PSD) diffusion *tensor* fields, using a deep-learning approach. This work allows for modeling diffusion anisotropy via tensors, and significantly reduces the inference time by resorting to deep learning approaches. However, regional anomalies are not explicitly modeled. Hence, additional post-processing steps, such as lesion detection, are required in perfusion analysis applications. In contrast, D<sup>2</sup>-SONATA not only predicts the spatial dependency of advection-diffusion parameters in a stochastic manner, but it also decomposes transport into its normal and abnormal parts and can thereby provide additional insights into existing patterns of anomaly.

### 3. Method

#### 3.1. Problem Setup

Let  $C = C(\mathbf{x}, t)$  denote the mass concentration at location  $\mathbf{x}$  in a bounded domain  $\Omega \subset \mathbb{R}^d$  ( $d = 2, 3$ ), at time  $t \in [0, T]$ . The local mass concentration changes of an advection-diffusion process can be modeled as:

$$\begin{aligned} \frac{\partial C}{\partial t} &= -\mathbf{V} \cdot \nabla C + \nabla \cdot (\mathbf{D} \nabla C) + \sigma \partial W_t \quad (1) \\ &= \underbrace{-(A \diamond \bar{\mathbf{V}}) \cdot \nabla C}_{\text{Anomaly-decomposed Incompressible flow}} + \underbrace{\nabla \cdot ((A \circ \bar{\mathbf{D}}) \nabla C)}_{\text{Anomaly-decomposed PSD Diffusion}} + \underbrace{\sigma \partial W_t}_{\text{Model Uncertainty}}, \end{aligned}$$

with specified boundary conditions (B.C.). The advection term captures the transport related to the flow of the fluid and the diffusion is driven by the gradient of mass concentration. The spatially-varying velocity field  $\mathbf{V}$  ( $\mathbf{V}(\mathbf{x}) \in \mathbb{R}^d$ ) and diffusion tensor field  $\mathbf{D}$  ( $\mathbf{D}(\mathbf{x}) \in \mathbb{R}^{d \times d}$ ) describe the advection and diffusion. Eq. (1) contains the common assumption that fluids show negligible density variation in practice [26, 37], which means  $\mathbf{V}$  has zero divergence (i.e., is divergence-free).<sup>2</sup> As in [37], we model diffusion tensors  $\mathbf{D}(x)$  as symmetric PSD, to capture the predominant diffusion direction and diffusion anisotropy [47].

In contrast to existing approaches we explicitly model abnormalities via an anomaly value field  $A$  ( $A(x) \in \mathbb{R}_{(0,1]}$ , higher  $A(x)$  means closer to normal), which modulates both the velocity and diffusion fields. Specifically, we assume that the velocity field can be decomposed as  $\mathbf{V} = A \diamond \bar{\mathbf{V}}$  and the diffusion tensor field as  $\mathbf{D} = A \circ \bar{\mathbf{D}}$ , where  $\bar{\mathbf{V}}$  and  $\bar{\mathbf{D}}$  denote the ‘‘anomaly-free’’ velocity and diffusion tensor fields respectively. Here,  $\diamond$  and  $\circ$  denote the chosen interactions between the ‘‘anomaly-free’’ velocity and diffusion

<sup>2</sup>Note when  $\mathbf{D} \rightarrow 0$ , Eq. (1) is an advection equation, which is the basis for many variational optical flow or image registration methods [5, 7, 9, 14, 23, 24, 43, 54, 61, 62]. If density variations should be modeled one can simply replace  $\mathbf{V} \cdot \nabla C$  with  $\nabla \cdot (C \mathbf{V})$ . (See more details in Supp. C.)

tensor fields and the anomaly value field respectively (see the definition in Eq. (4)). This interaction can be chosen in different ways. See Sec. 3.2 (Theorems 1 & 2 and Definition 1) for details on the anomaly-decomposed representations (and how  $\diamond$  and  $\circ$  are operationalized) which assure that the decompositions can express divergence-free velocity fields and PSD diffusion tensor fields.

In addition, we model the advection-diffusion process via a stochastic PDE (SPDE), where  $\sigma$  ( $\sigma(\mathbf{x}) \in \mathbb{R}$ ) denotes the variance of the Brownian motion  $W_t$  ( $W(\mathbf{x}, t) \in \mathbb{R}$ ) [4] and represents the epistemic uncertainty for the dynamical system. With this additional stochastic term, the existence and uniqueness of the solution to Eq. (1) still holds [4, 49, 60].

#### 3.2. Anomaly-decomposed Constraint-free Representation

As discussed in Sec. 3.1, incompressibility and symmetric PSD-ness are commonly used assumptions for fluid flow and diffusion. In YETI [37], divergence-free vectors and PSD tensors are obtained by construction via a suitable parameterization. Hence, when integrated into a deep-learning model no extra losses need to be imposed during training to obtain these properties. However, YETI does not provide an explicit abnormality model for the predicted  $\mathbf{V}$  and  $\mathbf{D}$  fields. D<sup>2</sup>-SONATA, therefore introduces two representation theorems which allow constructing divergence-free velocity fields and symmetric PSD diffusion tensor fields in such a way that a decomposition into normal and abnormal value fields is obtained.

##### Anomaly-decomposed Divergence-free Velocity Vectors

Bézenac et al. [12] penalize deviations of predicted velocity fields from zero divergence during training, with the assumption that test time predictions will be *approximately* divergence free. Kim et al. [26] parametrize velocity vectors via the curl of vector fields, yet do not account for the boundary conditions that need to be imposed in bounded domain scenarios [2, 3, 15, 40]. Liu et al. [37] propose to predict a parameterization of velocity fields  $\mathbf{V}$  which is by construction divergence-free; however, abnormal regions are not explicitly modeled for  $\mathbf{V}$  and extra post-processing steps are required for applications targeted at anomaly detection. Therefore, our goal is a representation strategy for  $\mathbf{V}$  on a domain  $\Omega \subset \mathbb{R}^d$  ( $d = 2, 3$ ) with a smooth boundary such that (1)  $\mathbf{V}$  is divergence-free *by construction*; (2) Any divergence-free  $\mathbf{V}$  can be represented; (3)  $\mathbf{V}$  can be decomposed into an anomaly value field  $A$  which characterizes the abnormal fields, and  $\bar{\mathbf{V}}$  which refers to the corresponding normal value field without anomaly.

##### Theorem 1 (Anomaly-decomposed Divergence-free Vector Representation).

*For any vector field  $\mathbf{V} \in L^p(\Omega)^d$  and scalar field  $A$  in  $\mathbb{R}_{(0,1]}(\Omega)$  on a bounded domain  $\Omega \subset \mathbb{R}^d$  with smooth boundary  $\partial\Omega$ . If  $\mathbf{V}$  satisfies  $\nabla \cdot \mathbf{V} = 0$ , there*

exists a potential  $\Psi$  in  $L^p(\Omega)^\alpha$  ( $\alpha = 1(3)$  when  $d = 2(3)$ ):

$$\mathbf{V} = \nabla \times (A\Psi), \quad (A\Psi) \cdot \mathbf{n}|_{\partial\Omega} = 0. \quad (2)$$

Conversely, for any  $A \in \mathbb{R}_{(0,1]}(\Omega)$ ,  $\Psi \in L^p(\Omega)^\alpha$ ,  $\nabla \cdot \mathbf{V} = \nabla \cdot (\nabla \times (A\Psi)) = 0$ . (See complete proof in Supp. A.)

**Anomaly-decomposed Symmetric PSD Diffusion Tensors** We seek a representation for  $\mathbf{D}$  (for a  $\mathbf{D}(x)$ ) such that (1)  $\mathbf{D}$  is a symmetric PSD tensor by construction; (2) any symmetric PSD  $\mathbf{D}$  can be represented; (3)  $\mathbf{D}$  can be decomposed into an anomaly value field  $A$  indicating the abnormal patterns, and  $\overline{\mathbf{D}}$  which is the corresponding normal value field. We resort to the spectral decomposition theorem and the surjective Lie exponential mapping on  $SO(n)$  ( $\exp: \mathfrak{so}(n) \mapsto SO(n)$ ,  $\mathfrak{so}(n)$  is the group of skew-symmetric matrices,  $SO(n)$  is the real orthogonal group) [31, 32, 37].

**Theorem 2 (Anomaly-decomposed Symmetric PSD Tensor Representation).** For any  $n \times n$  symmetric PSD tensor  $\mathbf{D}$  and  $A \in \mathbb{R}_{(0,1]}(\Omega)$ , there exist an upper triangular matrix with zero diagonal entries,  $\mathbf{B} \in \mathbb{R}^{\frac{n(n-1)}{2}}$ , and a non-negative diagonal matrix,  $\Lambda \in SD(n)$ , satisfying:

$$\mathbf{D} = \mathbf{U}(A\Lambda)\mathbf{U}^T, \quad \mathbf{U} = \exp(\mathbf{B} - \mathbf{B}^T) \in SO(n). \quad (3)$$

Conversely, for  $\forall A \in \mathbb{R}_{(0,1]}(\Omega)$ ,  $\forall \mathbf{B} \in \mathbb{R}^{\frac{n(n-1)}{2}}$ , and  $\forall \Lambda \in SD(n)$ , Eq. (3) results in a symmetric PSD tensor,  $\mathbf{D}$ . (See complete proof in Supp. B.)

**Definition 1 (“Anomaly-free” Fields).** According to Eq. (2-3), when  $A$  equals 1 over the entire domain ( $\Omega$ ), we denote the corresponding parameters as “anomaly-free”. For the sake of convenience, we write the relation:

$$\mathbf{V} = A \diamond \overline{\mathbf{V}}, \quad \mathbf{D} = A \circ \overline{\mathbf{D}}, \quad (4)$$

where the overline is used to denote “anomaly-free” fields. Obviously, when  $A \rightarrow 1$ ,  $\mathbf{V}, \mathbf{D}$  are closer to normal. (See explicit expressions of the  $\diamond, \circ$  operations in Supp. A & B.)

In this way, we not only obtain the underlying advection-diffusion parameters ( $\mathbf{V}, \mathbf{D}$ ) satisfying the divergence-free and PSD constraints from the observed transport processes by construction, but we also gain knowledge as to where the anomalies exist and what the normal patterns of the advection-diffusion parameters are. Note that in our proposed approach we will directly predict  $A, \overline{\mathbf{V}}$ , and  $\overline{\mathbf{D}}$ . I.e., we will directly work in the decomposed domain rather than predicting  $\mathbf{V}$  and  $\mathbf{D}$  followed by a decomposition.

### 3.3. $D^2$ -SONATA: Deep Decomposition for Stochastic Normal-Abnormal Transport

Sec. 3.2 described the representation theorems of  $D^2$ -SONATA for anomaly-encoded divergence-free velocity fields and symmetric PSD diffusion fields. This enables expressing the advection-diffusion parameters with realistic constraints along with the corresponding anomaly patterns by construction. In particular, using these representations

within a deep network will assure these constraints during training and testing. This section introduces  $D^2$ -SONATA’s two-stage learning framework, to predict the velocity vector and diffusion tensor fields driving a potentially anomalous observed advection-diffusion process:

(1) *Physics-informed learning*: reconstructs the anomaly value field ( $A$ ) and “anomaly-free” advection-diffusion parameters ( $\overline{\mathbf{V}}, \overline{\mathbf{D}}$ ), from observed transport time-series. The anomaly-encoded parameters ( $\mathbf{V}, \mathbf{D}$ ) can then be obtained (Eq. (4)). In this stage, the model is trained on a simulated dataset (Sec. 4.1-4.2) under the supervision of the ground truth parameters and a corresponding anomaly value field; (2) *Transport-informed learning*: reconstructs the transport concentration time-series, where the ground truth advection-diffusion parameters can be unknown. During this stage, the supervision is imposed by the observed mass transport dynamics, via the integration of the stochastic advection-diffusion PDE in Eq. (1).

**Patch-based Input Time-series** Given a time-series  $C = \{C^i \in \mathbb{R}(\Omega) | i = 1, 2, \dots, N_T\}$ , we first extract  $32^3$  ( $32^2$  for 2D domains) patches randomly (uniform over the space) over the same spatial domain ( $\Omega_p \subset \Omega$ ) at  $N_{\text{in}}$  sequential time points, starting from a random  $t_i$  ( $i \in \{1, 2, \dots, N_T - N_{\text{in}} + 1\}$ ). Each training sample ( $C_p$ ) is denoted as  $C_p = \{C_p^j \in \mathbb{R}(\Omega_p) | j = i, \dots, i + N_{\text{in}} - 1\}$  (Fig. 1 (top left)).

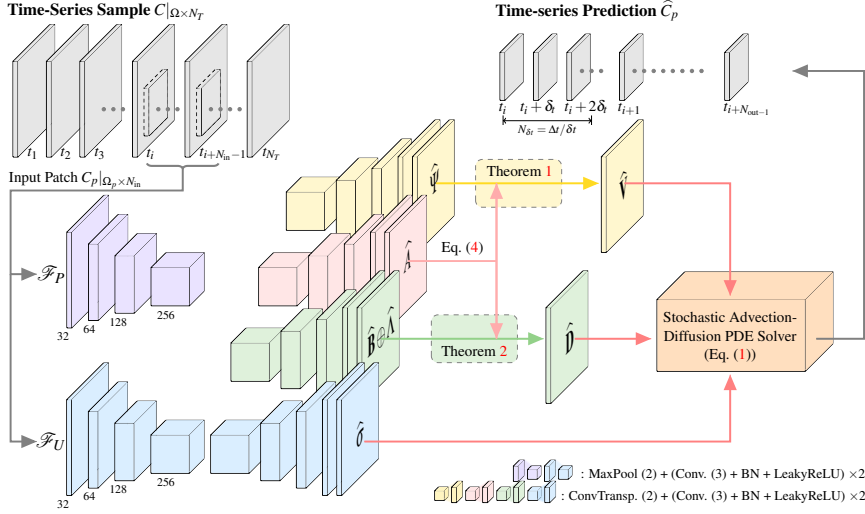
**Physics-informed Learning** We use a U-Net [1, 10]) as the backbone for the parameter prediction network ( $\mathcal{F}_p$ ) in this stage (Fig. 1 (middle)). A transport encoder is first used to encode the input  $C_p$  into latent features. The number of the input time points determines the input channels. Two decoders then learn the mappings to the representations of the “anomaly-free” parameters ( $\overline{\mathbf{V}}, \overline{\mathbf{D}}$ ). Specifically, the  $\overline{\mathbf{V}}$ -decoder predicts the potential  $\widehat{\Psi}$  to represent a divergence-free velocity field  $\widehat{\mathbf{V}}$  via Theorem 1, and the  $\overline{\mathbf{D}}$ -decoder predicts  $\widehat{\mathbf{B}}$  along with  $\widehat{\Lambda}$  to represent a symmetric PSD diffusion tensor field  $\widehat{\mathbf{D}}$  via Theorem 2. Another decoder outputs the predicted anomaly value field ( $A$ ). The resulting  $\widehat{A}, \widehat{\mathbf{V}}$  and  $\widehat{\mathbf{D}}$  then allow computing the final prediction of the estimated velocity ( $\widehat{\mathbf{V}}$ ) and diffusion ( $\widehat{\mathbf{D}}$ ) fields (Eq. (4)).

The reconstruction loss in this stage is three-fold. It is a supervised loss based on the actual and “anomaly-free” advection-diffusion parameters and the anomaly value field, which are all known from simulation:

$$\mathcal{L}_{\text{VD}} = \frac{1}{|\Omega_p|} \int_{\Omega_p} \left\{ \|\overline{\mathbf{V}} - \widehat{\mathbf{V}}\|_2 + \|\mathbf{V} - \widehat{\mathbf{V}}\|_2 + \|\overline{\mathbf{D}} - \widehat{\mathbf{D}}\|_F + \|\mathbf{D} - \widehat{\mathbf{D}}\|_F + |A - \widehat{A}| \right\} dx, \quad (5)$$

where the hatted symbols denote the predictions,  $\|\cdot\|_2, \|\cdot\|_F$  the vector 2-norm and tensor Frobenius norm respectively.

Similar to YETI [37], we additionally impose supervision on the eigenvectors ( $\widehat{\mathbf{U}}$ ) and eigenvalues ( $\widehat{\Lambda}$ ) of  $\widehat{\mathbf{D}}$  to improve the network’s ability to capture the anisotropic struc-



Alg. 1: Pseudocode for  $D^2$ -SONATA

---

**Input:** Time series of mass concentration  
 $C_p|_{\Omega \times N_T} = \{C_p^i \in \mathbb{R}(\Omega) | i = 1, 2, \dots, N_T\}$

**Output:**  $\hat{\mathbf{V}}, \hat{\mathbf{D}}, \hat{\mathbf{U}}, \hat{\Lambda}, \hat{\sigma}, \{\hat{C}_p^j | j = i, \dots, i + N_{out} - 1\}$

- 1 while  $\mathcal{L}_{\text{Phy}}$  not converged do // Physics-informed Learning
- 2 Randomly pick time-course as training sample  
 $C_p|_{\Omega_p \times N_{in}} = \{C_p^j \in \mathbb{R}(\Omega_p) | j = i, \dots, i + N_{in} - 1\}$
- 3 Reconstruct potential  $\hat{\Psi}, \hat{\mathbf{B}}, \hat{\Lambda}$  and  $\hat{\sigma}$
- 4 Represent divergence-free  $\hat{\mathbf{V}}$  by  $\hat{\Psi}, \hat{\Lambda}$  (Eq. (2.4))
- 5 Represent symmetric PSD  $\hat{\mathbf{D}}$  by  $\hat{\mathbf{B}}, \hat{\Lambda}$  (Eq. (3.4))
- 6 Compute  $\mathcal{L}_{\text{Phy}}$  and backpropagate
- 7 while  $\mathcal{L}_{\text{Tm}}$  not converged do // Transport-informed Learning
- 8 Process forward line 2-6
- 9 for  $t = t_i + \delta t, \dots, t_{i+1}, t_{i+1} + \delta t, \dots, t_{i+N_{out}-1}$  do
- 10 Discretize in space (Sec. 3.3)
- 11 Compute advection-diffusion SPDE via Eq. (1)
- 12 Impose patch-based Cauchy B.C.
- 13 Integrate in time, obtain  $\hat{C}_p^{t+\delta t}$
- 14 Compute  $\mathcal{L}_{\text{Tm}}$  and backpropagate

---

Legend:  
█ : MaxPool (2) + (Conv. (3) + BN + LeakyReLU)  $\times 2$   
█ : ConvTransp. (2) + (Conv. (3) + BN + LeakyReLU)  $\times 2$   
█ : Stochastic Advection-Diffusion PDE Solver (Eq. (1))

Figure 1. The two-stage learning framework for  $D^2$ -SONATA. During the physics-informed stage, the model is trained under the supervision of the ground truth anomaly value field ( $A$ ), and the “anomaly-free” advection-diffusion parameters ( $\hat{\mathbf{V}}, \hat{\mathbf{D}}$ ). In the transport-informed stage, the model is trained based on the observed concentration time-series, which allows recovering the underlying  $\mathbf{V}, \mathbf{D}$  that best fit the input time-series. The transport-informed stage is therefore applicable to real data where only the concentration time-series are observed.

ture of diffusion tensors. Note  $\hat{\mathbf{U}} = [\hat{\mathbf{u}}_1, \hat{\mathbf{u}}_2, \hat{\mathbf{u}}_3]$  is an intermediate output from  $\hat{\mathbf{B}}$  via Eq. (3).

$$\mathcal{L}_{\text{UA}} = \frac{1}{|\Omega_p|} \int_{\Omega_p} \sum_{i=1}^3 \min\{\|\mathbf{u}_i \pm \hat{\mathbf{u}}_i\|_2\} + \|\Lambda - \hat{\Lambda}\|_{\text{F}} dx, \quad (6)$$

where  $\min$  addresses the eigenvector sign ambiguities by selecting the minimum between  $\|\mathbf{u}_i + \hat{\mathbf{u}}_i\|_2$  and  $\|\mathbf{u}_i - \hat{\mathbf{u}}_i\|_2$ .

The loss for the physics-informed learning stage is

$$\mathcal{L}_{\text{Phy}} = \mathcal{L}_{\text{VD}} + w_{\text{UA}} \mathcal{L}_{\text{UA}}, \quad w_{\text{UA}} > 0. \quad (7)$$

**Transport-informed Learning** In this stage (Fig. 1), the model is trained under the supervision of the transport dynamics. It contains two parts: (1) an uncertainty prediction network ( $\mathcal{F}_U$ ); (2) a stochastic advection-diffusion PDE integrator. Specifically,  $\mathcal{F}_U$  is used to represent the model’s epistemic uncertainty, namely  $\sigma$  in Eq. (1). Intuitively, the model’s epistemic uncertainty should capture the following: (1) For values within the normal range, the variance of the Brownian motion should be small (low uncertainty). In this case the system state is dominated by the advection-diffusion term. (2) For anomaly regions with parameter values outside the normal range, the variance of the Brownian motion should be large and the system should exhibit higher uncertainty [28]. Based on this desired property, we design our training loss for model uncertainty as

$$\mathcal{L}_{\sigma} = \frac{1}{|\Omega_p|} \int_{\Omega_p} ((1 - A) - \hat{\sigma})^2 dx, \quad (8)$$

with the prior assumption that uncertainties should be small when the ground truth  $A$  is close to the normal value, 1.

Next, we implement a stochastic advection-diffusion PDE solver to integrate the initial state ( $C_p^i$ ) from time to  $t_{i+N_{out}-1}$  via Eq. (1) (Fig. 1 (right)), and to train the

model by minimizing the differences between the predicted ( $\hat{C}_p$ ) and the input ( $C_p$ ) time-series. We follow the patch-based Cauchy B.C. and the spatial discretization scheme described in Liu et al. [30, 37, 59]. We also use RK45 for time-integration ( $\delta t$ ) for the concentration prediction  $\hat{C}_p^{t+\delta t}$ . For the stochastic term ( $\sigma$ ) in Eq. (1), we use the Euler-Maruyama scheme [27, 28], and the discretized version of the stochastic term in Eq. (1) can therefore be written as  $\sigma W / \sqrt{\Delta t}$ , where  $W \sim \mathcal{N}(0, 1)$ .<sup>3</sup> (See Supp. C for details.)

Given a training patch sample  $\{C_p^j \in \mathbb{R}(\Omega_p) | j = i, \dots, i + N_{in} - 1\}$ , we compute the collocation concentration loss ( $\mathcal{L}_{\text{CC}}$ ) [37, 41], the mean squared error of the predicted time series at output collocation time points, to encourage the predicted values to be close to the observed ones. We also use a regularizer ( $\mathcal{L}_{\text{SS}}$ ) [37] on the gradient fields of each component of  $\hat{\mathbf{V}}, \hat{\mathbf{D}}$  to encourage the predicted parameter fields to be spatially smooth.

The complete loss for this stage is thus ( $w_{\text{SS}}, w_{\sigma} > 0$ )

$$\mathcal{L}_{\text{Tm}} = \mathcal{L}_{\text{CC}} + w_{\text{SS}} \mathcal{L}_{\text{SS}} + w_{\sigma} \mathcal{L}_{\sigma}. \quad (9)$$

## 4. Experiments

Secs. 4.1-4.2 present results on simulated time-series in 2D/3D. We analyze the individual contribution of  $D^2$ -SONATA’s components, and compare the improvements achieved to the state of the art approaches. In Sec. 4.3, we further transfer our model pre-trained on simulated data to real time-series of magnetic resonance (MR) perfusion images from ischemic stroke patients. We demonstrate  $D^2$ -SONATA’s ability to distinguish stroke lesions

<sup>3</sup>Note when the temporal resolution ( $\Delta t$ ) is coarse,  $\delta t$  should be set smaller than  $\Delta t$  (Fig. 1 (top right)) to satisfy the Courant-Friedrichs-Lewy (CFL) condition [20, 30] and to allow for stable numerical integration.

Models	Training Data	V, D	$\bar{V}, \bar{D}$	A	$\sigma$	Type
PIANO [35]	N/A	✓	✗	✗	✗	Optimization-based
YETI [37]	Normal	✓	✗	✗	✗	Learning-based
Anom.-YETI	(Ab-)Normal	✓	✗	✗	✗	Learning-based
w/o- $\sigma$	(Ab-)Normal	✓	✓	✓	✗	Learning-based
w/o-A	(Ab-)Normal	✓	✓	✗	✓	Learning-based
D <sup>2</sup> -SONATA	(Ab-)Normal	✓	✓	✓	✓	Learning-based

Table 1. Model properties comparison.




Figure 2. Anisotropic moving Gaussians (Sec. 4.1), examples for normal (1<sup>st</sup> row) and abnormal (2<sup>nd</sup> row) transport time-series.




Figure 3. Mean relative absolute error (RAE) of YETI, Anom.-YETI, w/o- $\sigma$ , w/o-A and D<sup>2</sup>-SONATA for 2D ab-(normal) anisotropic moving Gaussian (Sec. 4.1). X-axis: training iterations; Y-axis: the average RAE (log scale) of 100 test samples.

from normal brain regions via its predicted anomaly value field  $\hat{A}$  and reconstructed advection-diffusion parameters ( $\hat{V}, \hat{D}$ ).

#### 4.1. 2D (Ab-)Normal Anisotropic Moving Gaussian Dataset

We simulate 2D advection-diffusion with random anomalies on-the-fly (Fig. 2). Each sample is a 2D image time-series ( $N_T = 40, \Delta t = 0.01 s$ ) on a  $64^2$  domain with  $1 mm$  spacing. Every time-series is an advection-diffusion process initialized by a Gaussian ( $\mathcal{N}(0, 2)$ ) at a uniformly sampled center. The “anomaly-free” advection-diffusion parameters ( $\bar{V}, \bar{D}$ ) are obtained from randomly generated potentials  $\Psi, \mathbf{B}, \Lambda$  with uniformly sampled dominant diffusion directions. All components of  $\Lambda$  are randomly sampled from  $\mathcal{U}_{[0,1]}$ , and  $\Psi$  is randomly sampled from  $\mathcal{U}_{[-10,10]}$ . Anomaly value fields are applied on the original “anomaly-free” parameters with a probability of 50%, where the value ( $\in [0, 1]$ ) and area of anomaly fields are determined by a multivariate Gaussian with its center and standard deviation uniformly sampled in the domain. The initial Gaussian is then transported by  $\mathbf{V}, \mathbf{D}$ . (See Supp. C for details.)

**Comparisons** To examine the contributions of different components in D<sup>2</sup>-SONATA, we compare five vari-

Models	$\hat{C}$	$\hat{V}$	$\hat{D}$	$\hat{U}$	$\hat{\Lambda}$	$\hat{A}$
YETI [37]	8.53	18.21	6.48	4.43	4.71	N/A
Anom.-YETI	2.52	12.23	5.15	4.12	3.54	N/A
w/o- $\sigma$	2.41	10.91	4.67	4.07	3.08	0.51
w/o-A	2.29	12.10	4.61	4.11	3.17	N/A
D <sup>2</sup> -SONATA	<b>2.12</b>	<b>10.68</b>	<b>4.21</b>	<b>3.08</b>	<b>2.94</b>	<b>0.47</b>

Table 2. Reconstruction comparisons on mean RAE (%) (Eq. (10)) for 3D (ab-)normal brain advection-diffusion dataset (Sec. 4.2).

ants (Tab. 1): (1) YETI: the original setting of YETI by Liu et al. [37], where the model is trained with anisotropic Gaussians with no anomaly value field applied, and the advection-diffusion parameters are deterministically predicted without predicting an anomaly field or uncertainty; (2) Anom.-YETI: here the YETI model remains unchanged, but is trained with anisotropic Gaussians of which 50% are affected by the anomaly fields introduced in Sec. 4.1; (3) w/o- $\sigma$ : D<sup>2</sup>-SONATA without the stochastic term and the model uncertainty network  $\mathcal{F}_U$ ; (4) w/o-A: D<sup>2</sup>-SONATA without the decoder branch for the anomaly value field (A) in  $\mathcal{F}_P$ . In this case, the  $\bar{V}, \bar{D}$ -decoders directly output the overall value for  $\mathbf{V}, \mathbf{D}$ , instead of predicting the anomaly value field and corresponding anomaly-free parameters separately; (5) The full D<sup>2</sup>-SONATA model.

We follow the training schedule in [37], and re-train YETI with its original settings. The time-series sample length for all models is  $N_{in} = 10$ . We set  $N_{out} = 10, w_{U\Lambda} = 0.5, w_{SS} = 0.1, w_{\sigma} = 0.5$ . We test on 100 samples after every 500 training epochs and use the mean relative absolute error (RAE) [37] for the predicted  $\hat{V}, \hat{D}, \hat{U}, \hat{\Lambda}, \hat{A}$ :

$$Err(\mathbf{F}) = \frac{1}{|\Omega|} \int_{\Omega} \|\mathbf{F} - \hat{\mathbf{F}}\| / \|\mathbf{F}\| dx, \quad (10)$$

where  $\mathbf{F}(\hat{\mathbf{F}})$  denotes the ground truth (prediction),  $\|\cdot\|$  is the absolute, 2-norm and Frobenius norm for scalars, vectors and tensors. The time-series error ( $Err(C)$ ) computes the RAE averaged over all predicted collocation time points.

Fig. 3 shows the testing reconstruction errors of the five models throughout training. Being exposed to abnormal data (Anom.-YETI) during training significantly boosts YETI’s reconstruction performance. The anomaly value field (A) prediction helps  $\mathbf{V}, \mathbf{D}$ ’s reconstruction, and modeling as a stochastic system ( $\sigma$ ) further improves the time-series predictions. Overall, with anomaly-prediction and stochastic modeling, D<sup>2</sup>-SONATA outperforms all other variants with respect to both (ab-)normal time-series prediction and advection-diffusion parameter reconstruction.

#### 4.2. 3D (Ab-)Normal Brain Advection-Diffusion Dataset

**Dataset** We develop a normal-abnormal brain advection-diffusion simulator based on the IXI brain dataset<sup>4</sup>. Following the simulation procedure in [37], we use 200 pa-

<sup>4</sup><http://brain-development.org/ixi-dataset/>




Figure 4. Reconstruction comparisons (Tab. 1) of one slice from a test case for 3D (ab-)normal brain advection-diffusion dataset (Sec. 4.2). ( $\|\mathbf{V}\|_2$  shown in maximum intensity projection.)

tients that have T1/T2-weighted images, a magnetic resonance angiography (MRA) image, and diffusion-weighted images (DWI) with 15 directions, to simulate the normal 3D velocity and diffusion tensor fields<sup>5</sup>. For each case, we simulate both normal samples and anomaly-encoded samples, i.e., with simulated anomaly value fields applied to both  $\bar{\mathbf{V}}$  and  $\bar{\mathbf{D}}$ , where the value (within  $[0, 1]$ ) and area of the anomaly fields are calculated by a multivariate Gaussian with its center and standard deviation uniformly sampled over the spatial domain. Each resulting image has isotropic spacing ( $1\text{ mm}$ ) and has been rigidly registered intra-subject. In summary, we simulated 2,400 brain advection-diffusion time-series, where 50% contain simulated anomalies. We randomly select 40 time-series for validation and testing respectively. (See Supp. D for more simulation details.)

**Comparisons** We compare the same five models as in Sec. 4.1 (Tab. 1). We follow the training strategy in YETI [37] and re-train YETI with its original settings. Predicted parameters ( $\hat{\mathbf{V}}, \hat{\mathbf{D}}, \hat{\mathbf{U}}, \hat{\Lambda}, \hat{A}$ ) on the entire domain are obtained by splicing the output patches together. Integrating the stochastic advection-diffusion PDE forward with  $\hat{\mathbf{V}}, \hat{\mathbf{D}}$ , we obtain the predicted time-series  $\hat{C}$  on the original domain. For better visualization, we use the 2-norm ( $\|\mathbf{V}\|_2$ ) map for the velocity fields, and the commonly used tensor feature maps [37, 45] for the diffusion fields: (1) Trace ( $tr_{\mathbf{D}}$ ), the sum of tensor eigenvalues ( $\Lambda$ ): the overall diffusion strength; and (2) Fractional anisotropy (FA): the amount of

<sup>5</sup>The goal is to obtain nontrivial advection-diffusion parameters for 3D simulation to boost a network’s pre-training with supervision. This will likely not mimick real perfusion across the entire brain, but is expected to provide quasi-realistic local patterns which would appear in real data.

anisotropy across diffusing directions.

Tab. 2 compares the reconstruction errors of the predicted  $\hat{\mathbf{V}}, \hat{\mathbf{D}}, \hat{\mathbf{U}}, \hat{\Lambda},$  and  $\hat{A}$  from the five models, where  $D^2$ -SONATA consistently outperforms *all* other models. Specifically, training with abnormal samples and the stochastic term ( $\sigma$ ) to model uncertainties helps achieve better overall performance, particularly in predicting the transport time-series. Predicting with the proposed decomposition based on the anomaly value field ( $A$ ) further boosts the reconstruction of the advection-diffusion parameters. The advantages of  $D^2$ -SONATA can be observed more clearly in Fig. 4, where it successfully captures the anomaly value field ( $A$ ), and accurately reconstructs the overall magnitudes of the advection-diffusion parameters ( $\mathbf{V}, \mathbf{D}$ ). Without seeing abnormal samples during training, YETI struggles to identify the anomalies and to differentiate between advection ( $\mathbf{V}$ ) and diffusion ( $\mathbf{D}$ ) effectively (2<sup>nd</sup> column), where the resulting  $tr_{\mathbf{D}}$  reflects insufficient differences between normal and abnormal regions and the reconstructed  $\|\mathbf{V}\|_2$  is noisiest among all models. Training with abnormal samples generally helps the models to locate anomalies. Interestingly, decomposing into an anomaly value field and the “anomaly-free” parameters improves the reconstruction of the diffusion’s fractional anisotropy (FA), especially within the abnormal regions.

### 4.3. ISLES2017: Brain Perfusion Dataset from Ischemic Stroke Patients

**Dataset** We test on the Ischemic Stroke Lesion Segmentation (ISLES) 2017 dataset [39], which contains perfusion data from 75 (43 training, 32 testing) ischemic stroke patients. Each patient has a dynamic susceptibility contrast (DSC) MR perfusion image (4D, with 40 to 80 available time points, time interval  $\approx 1\text{ s}$ ), along with the corresponding gold-standard lesion segmentation maps [17]. All images are resampled to isotropic spacing ( $1\text{ mm}$ ) and rigidly registered intra-subject via `ITK` [42] following [37]. Perfusion images are converted to the tracer concentration time-series,  $\{C^i \in \mathbb{R}(\Omega) | i = 1, 2, \dots, N_T\}$  ( $t_1$  is the time-to-peak for total concentration over the entire brain, at which the injected tracer is assumed to have been fully transported into the brain [35, 37]), via the relation between MR signal and tracer intensity [19]. 10 patients with lesion segmentation maps are randomly chosen for testing, while the remaining 65 concentration time-series are left-right flipped (i.e., the brain hemispheres are flipped) for data augmentation. In total, we obtain 130 time-series training samples, from which 10 samples are randomly picked for validation.

**Results** We transfer the pre-trained model from our simulated 3D normal-abnormal brain advection-diffusion dataset (Sec. 4.2) to the ISLES tracer concentration time-series dataset, using the Adam optimizer and learning rate  $10^{-4}$ . We set  $N_{in} = N_{out} = 5$ ,  $w_{\sigma} = 0.5$ ,  $w_{SS} = 0.1$ .

Following the feature maps for comparison proposed




Figure 5. Lesion segmentation and corresponding  $D^2$ -SONATA feature maps of the same slice for four test stroke patients.

Metrics	$D^2$ -SONATA			YETI [37]		PIANO [35]		ISLES [39]			
	A	$\ \mathbf{V}\ _2$	$tr_{\mathbf{D}}$	$\ \mathbf{V}\ _2$	$tr_{\mathbf{D}}$	$\ \mathbf{V}\ _2$	$D$	CBF	CBV	MTT	
$\mu^r$	<i>Me.</i>	0.47	<b>0.29</b>	0.42	0.30	0.59	0.55	0.58	0.67	0.78	0.57
	<i>Med.</i>	0.49	<b>0.30</b>	0.48	0.32	0.59	0.54	0.55	0.59	0.79	0.58
	( $\downarrow$ ) <i>(STD)</i>	(0.13)	(0.17)	(0.15)	(0.11)	(0.19)	(0.15)	(0.16)	(0.12)	(0.23)	(0.13)
$ t $	<i>Me.</i>	<b>280</b>	165	166	155	49	108	52	34	16	31
	<i>Med.</i>	<b>286</b>	164	158	134	42	89	48	28	11	32
	( $\uparrow$ ) <i>(STD)</i>	(58)	(37)	(60)	(62)	(22)	(35)	(26)	(22)	(12)	(37)
AUC	<i>Me.</i>	<b>0.79</b>	0.70	0.64	0.73	0.51	0.74	0.68	0.72	0.65	0.65
	<i>Med.</i>	<b>0.76</b>	0.71	0.65	0.73	0.50	0.74	0.69	0.73	0.68	0.66
	( $\uparrow$ ) <i>(STD)</i>	(0.05)	(0.04)	(0.07)	(0.06)	(0.03)	(0.04)	(0.03)	(0.07)	(0.06)	(0.06)

\*  $\downarrow$  ( $\uparrow$ ) indicates the lower (higher) values are better.

Table 3. Quantitative comparison between  $D^2$ -SONATA, YETI, PIANO and ISLES maps across 10 test subjects from ISLES2017 dataset (Sec. 4.3), using *Mean (Me.)*, *Median (Med.)*, *Standard Deviation (STD)* of relative mean  $\mu^r$ , absolute  $|t|$ , and area under the curve (AUC).

in [35,37], we compute (1) the 2-norm ( $\|\mathbf{V}\|_2$ ) for the velocity field, and (2) the trace ( $tr_{\mathbf{D}}$ ) for diffusion tensors (both introduced in Sec. 4.2). In  $D^2$ -SONATA, we also show (3) the output model uncertainty ( $\sigma$ ), (4) the predicted anomaly value field ( $A$ ), and (5)  $Seg(A)$ , the segmentation map obtained by thresholding<sup>6</sup>  $A$ .

Fig. 5 shows the feature maps for four test patients from our model. All show features consistent with the manually segmented lesion maps. Importantly, besides the feature maps of the predicted advection-diffusion parameters ( $\|\mathbf{V}\|_2$ ,  $tr_{\mathbf{D}}$ ), we additionally obtain (1) the anomaly value field ( $A$ ) from which one can directly obtain the lesion segmentation ( $Seg(A)$ ), and (2) the ‘‘anomaly-free’’ parameters ( $\|\mathbf{V}\|_2$ ,  $tr_{\mathbf{D}}$ ) which provide insights into the expected normal brain perfusion of a patient.

**Comparisons** We compare the performance of feature maps from  $D^2$ -SONATA ( $\|\mathbf{V}\|_2$ ,  $\mathbf{D}$ ,  $A$ ) with (1) YETI [37]:  $\|\mathbf{V}\|_2$ ,  $\mathbf{D}$ ; (2) PIANO [35]:  $\|\mathbf{V}\|_2$ ,  $D$ <sup>7</sup>; and (2) ISLES [39]

<sup>6</sup>The chosen threshold is the optimal threshold for the segmentation Dice score, averaged from 6 randomly-selected test patients.

<sup>7</sup>We directly use  $D$  for PIANO as it models diffusion as scalar fields.

(perfusion summary maps): cerebral blood flow (CBF), cerebral blood volume (CBV), mean transit time (MTT). We report the two measures proposed by Liu et al. [37] that focus on the feature map differences between lesions and normal regions: (1) Relative mean ( $\mu^r \in [0, 1]$ ):

$$\mu^r = \min \left\{ \frac{\text{mean in lesion}}{\text{mean in c-lesion}}, \frac{\text{mean in c-lesion}}{\text{mean in lesion}} \right\}, \quad (11)$$

where c-lesion refers to the contralateral region of a lesion obtained by mirroring via the midline of the cerebral hemispheres, and  $\min$  accounts for typically larger MTT values (while other metrics are typically smaller) in the lesion rather than the c-lesion region; (2) Absolute t-value ( $|t|$ ): the absolute value of the unpaired t-statistic between voxels in the lesion and the c-lesion regions.

The above two measures are defined given a known segmented lesion region. In this work, we additionally consider how well the feature maps can distinguish lesion from non-lesion regions. I.e., we want to directly use them for lesion segmentation. To this end we threshold all feature maps and additionally compare the resulting areas under the curve (AUC) of the receiver operating characteristic (ROC) curves computed from all feature maps.

Tab. 3 compares feature maps from  $D^2$ -SONATA, YETI, PIANO, and ISLES based on the above metrics for the 10 test patients. Despite predicting the advection-diffusion parameters in the form of an anomaly-encoded decomposition,  $D^2$ -SONATA still results in the lowest relative means ( $\mu^r$ ) between lesion and c-lesion for  $\|\mathbf{V}\|_2$  and  $tr_{\mathbf{D}}$ . This indicates  $D^2$ -SONATA’s stronger ability to differentiate between normal and abnormal fields. Moreover, the predicted anomaly value field ( $A$ ) obtains a much higher absolute t-value ( $|t|$ ) which indicates great potential for lesion segmentation. This is confirmed by  $A$  achieving the highest AUC score among all feature maps.

## 5. Conclusions

We presented  $D^2$ -SONATA, a two-stage learning scheme to estimate the velocity vector fields and diffusion tensor fields that drive real-world-observed advection-diffusion processes in the presence of anomalies. The model is built upon a stochastic advection-diffusion PDE and predicts divergence-free velocity fields and symmetric PSD diffusion tensor fields by construction. Further, our model directly predicts an anomaly field and provides measures of uncertainty. Extensive experiments on 2D and 3D simulated data illustrate that our model can successfully reconstruct the advection-diffusion fields (with or without anomalies) from the time-series data of observed transport processes. When applied to real brain perfusion data from ischemic stroke patients, our model outperforms the state of the art across *all* metrics. By predicting the underlying advection-diffusion parameters along with an anomaly value field our model provides additional insights into real stroke data.



## References

- [1] U-Net: Convolutional Networks for Biomedical Image Segmentation, author=Olaf Ronneberger and Philipp Fischer and Thomas Brox, year=2015, journal=arXiv preprint arXiv:1505.04597, . 4
- [2] Chérif Amrouche, Christine Bernardi, Monique Dauge, and Vivette Girault. Vector potentials in three-dimensional non-smooth domains. *Mathematical Methods in the Applied Sciences*, 21(9):823–864, 1998. 3
- [3] Chérif Amrouche and Nour El Houda Seloula. Lp-theory for vector potentials and Sobolev’s inequalities for vector fields: Application to the Stokes equations with pressure boundary conditions. *Mathematical Models and Methods in Applied Sciences*, 23(01):37–92, 2013. 3
- [4] Jo ao Guerra and David Nualart. Stochastic differential equations driven by fractional Brownian motion and standard brownian motion. *Stochastic Analysis and Applications*, 26(5):1053–1075, 2008. 3, 14
- [5] Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging*, 38(8):1788–1800, 2019. 2, 3
- [6] Yohai Bar-Sinai, Stephan Hoyer, Jason Hickey, and Michael P. Brenner. Learning data-driven discretizations for partial differential equations. *Proceedings of the National Academy of Sciences*, 116(31):15344–15349, 2019. 1
- [7] Mirza Faisal Beg, Michael I. Miller, Alain Trouvé, and Laurent Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision*, 61(2):139–157, 2005. 2, 3
- [8] Kaushik Bhattacharya, Bamdad Hosseini, Nikola B. Kovachki, and Andrew M. Stuart. Model reduction and neural networks for parametric PDEs. *arXiv preprint arXiv:2005.03180*, 2020. 1, 2
- [9] Alfio Borzi, Kazufumi Ito, and Karl Kunisch. Optimal control formulation for determining optical flow. *SIAM journal on scientific computing*, 24(3):818–847, 2003. 2, 3
- [10] Özgün Çiçek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. 3D U-Net: Learning dense volumetric segmentation from sparse annotation. In Sebastien Ourselin, Leo Joskowicz, Mert R. Sabuncu, Gozde Unal, and William Wells, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*, pages 424–432, Cham, 2016. Springer International Publishing. 4
- [11] Andrew Cookson, Jack Lee, Christian Michler, Radomir Chabiniok, Eoin R Hyde, David Nordsletten, and Nicolas Smith. A spatially-distributed computational model to quantify behaviour of contrast agents in MR perfusion imaging. *Medical Image Analysis*, 18(7):1200–1216, 2014. 2
- [12] Emmanuel de Bézenac, Arthur Pajot, and Patrick Gallinari. Deep learning for physical processes: Incorporating prior scientific knowledge. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018. 1, 3
- [13] Jelle Demeestere, Anke Wouters, Soren Christensen, Robin Lemmens, and Maarten G. Lansberg. Review of perfusion imaging in acute ischemic stroke. *Stroke*, 51(3):1017–1024, 2020. 2
- [14] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Häusser, Caner Hazirbas, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 2758–2766, 2015. 2, 3
- [15] Francois Dubois. Discrete vector potential representation of a divergence-free vector field in three-dimensional domains: Numerical analysis of a model problem. *SIAM Journal on Numerical Analysis*, 27(5):1103–1141, 1990. 3
- [16] Weinan E and Bing Yu. The deep Ritz method: A deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6(1):e2019WR026731, 2018. 1, 2
- [17] Marco Essig, Mark S Shiroishi, Thanh Binh Nguyen, Marc Saake, James M Provenzale, David Enterline, Nicoletta Anzalone, Arnd Dörfler, Alex Rovira, Max Wintermark, and Meng Law. Perfusion MRI: the five most frequently asked technical questions. *AJR. American journal of roentgenology*, 200(1):24–34, 2013. 7
- [18] Gwynne A. Evans, Jonathan M. Blackledge, and Peter D. Yardley. *Numerical Methods for Partial Differential Equations*. Springer, 2000. 14
- [19] Andreas Fieselmann, Markus Kowarschik, Arundhuti Ganguly, Joachim Hornegger, and Rebecca Fahrig. Deconvolution-based CT and MR brain perfusion measurement: Theoretical model revisited and practical implementation details. *Journal of Biomedical Imaging*, 2011, 2011. 2, 7
- [20] Sigal Gottlieb and Lee-Ad J. Gottlieb. Strong stability preserving properties of Runge-Kutta time discretization methods for linear constant coefficient operators. *Journal of Scientific Computing*, pages 83–109, 2003. 5
- [21] Julie Marie Grøner, Rune Paamand, Liselotte Højgaard, and Ian Law. Brain perfusion CT compared with <sup>15</sup>O-H<sub>2</sub>O-PET in healthy subjects. *EJNMMI research*, 1(1), 2011. 2
- [22] Vratislav Harabis, Radim Kolar, Martin Mezl, and Radovan Jirik. Comparison and evaluation of indicator dilution models for bolus of ultrasound contrast agents. *Physiological measurement*, 34(2):151–162, 2013. 2
- [23] Gabriel L. Hart, Christopher Zach, and Marc Niethammer. An optimal control approach for deformable registration. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 9–16. IEEE, 2009. 2, 3
- [24] Berthold K.P. Horn and Brian G. Rhunck. Determining optical flow. In *Techniques and Applications of Image Understanding*, volume 281, pages 319–331. International Society for Optics and Photonics, 1981. 2, 3
- [25] Berthold K.P. Horn and Brian G. Schunck. Determining optical flow. Technical report, 1980. 2
- [26] Byungsoo Kim, Vinicius C. Azevedo, Nils Thuerey, Theodore Kim, Markus Gross, and Barbara Solenthaler. Deep fluids: A generative network for parameterized fluid simulations. *Computer Graphics Forum (Proc. Eurographics)*, 38(2), 2019. 3

- [27] Peter E. Kloeden and Eckhard Platen. *Numerical Solution of Stochastic Differential Equations*. Springer, 2000. 5, 14
- [28] Lingkai Kong, Jimeng Sun, and Chao Zhang. SDE-Net: Equipping deep neural networks with uncertainty estimates. In *International Conference on Machine Learning (ICML)*, 2020. 5, 14
- [29] Sunil Koundal, Rena Elkin, Saad Nadeem, Yuechuan Xue, Stefan Constantinou, Simon Sanggaard, Xiaodan Liu, Britany Monte, Feng Xu, William Van Nostrand, Maiken Nedergaard, Hedok Lee, Joanna Wardlaw, Helene Benveniste, and Allen Tannenbaum. Optimal mass transport with Lagrangian workflow reveals advective and diffusion driven solute transport in the lymphatic system. *Scientific Reports*, 10(1):1990, 2020. 1
- [30] Randall J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge Texts in Applied Mathematics. Cambridge University Press, 2002. 5
- [31] Mario Lezcano-Casado. Trivializations for gradient-based optimization on manifolds. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 9154–9164, 2019. 4
- [32] Mario Lezcano-Casado and David Martínez-Rubio. Cheap orthogonal constraints in neural networks: A simple parametrization of the orthogonal and unitary group. In *International Conference on Machine Learning (ICML)*, pages 3794–3803, 2019. 4
- [33] Zongyi Li, Nikola Kovachki, Kamyar Aizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020. 1, 2
- [34] Chad Lieberman, Karen Willcox, and Omar Ghattas. Parameter and state model reduction for large-scale statistical inverse problems. *SIAM J. Sci. Comput.*, 32(5):2523–2542, Aug. 2010. 1, 2
- [35] Peirong Liu, Yueh Z. Lee, Stephen R. Aylward, and Marc Niethammer. PIANO: Perfusion imaging via advection-diffusion. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2020. 1, 2, 6, 7, 8
- [36] Peirong Liu, Yueh Z. Lee, Stephen R. Aylward, and Marc Niethammer. Perfusion imaging: An advection diffusion approach. *IEEE Transactions on Medical Imaging*, 2021. 1, 2
- [37] Peirong Liu, Lin Tian, Yubo Zhang, Stephen Aylward, Yueh Lee, and Marc Niethammer. Discovering hidden physics behind transport dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 3, 4, 5, 6, 7, 8, 12, 13, 14
- [38] Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*, 2020. 1, 2
- [39] Oskar Maier, Bjoern H Menze, Janina von der Gablentz, Levin Hani, Mattias P Heinrich, Matthias Liebrand, Stefan Winzeck, Abdul Basit, Paul Bentley, Liang Chen, Daan Christiaens, Francis Dutil, Karl Egger, Chaolu Feng, Ben Glocker, Michael Götz, Tom Haeck, Hanna-Leena Halme, Mohammad Havaei, Khan M Iftekharuddin, Pierre-Marc Jodoin, Konstantinos Kamnitsas, Elias Kellner, Antti Korvenoja, Hugo Larochelle, Christian Ledig, Jia-Hong Lee, Frederik Maes, Qaiser Mahmood, Klaus H Maier-Hein, Richard McKinley, John Muschelli, Chris Pal, Linmin Pei, Janaki Raman Rangarajan, Syed M S Reza, David Robben, Daniel Rueckert, Eero Salli, Paul Suetens, Ching-Wei Wang, Matthias Wilms, Jan S Kirschke, Ulrike M Krämer, Thomas F Münte, Peter Schramm, Roland Wiest, Heinz Handels, and Mauricio Reyes. ISLES 2015 - a public evaluation benchmark for ischemic stroke lesion segmentation from multispectral MRI medical image analysis. *Medical Image Analysis*, 35, 2017. 7, 8
- [40] Filippo Maria Denaro. On the application of the Helmholtz-Hodge decomposition in projection methods for incompressible flows with general boundary conditions. *International Journal for Numerical Methods in Fluids*, 43(1):43–69, 2003. 3
- [41] Michaël Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. 5
- [42] Matthew McCormick, Xiaoxiao Liu, Julien Jomier, Charles Marion, and Luis Ibanez. ITK: enabling reproducible research and open science. *Frontiers in Neuroinformatics*, 8(13), 2014. 7
- [43] Jan Modersitzki. *Numerical methods for image registration*. Oxford University Press on Demand, 2004. 2, 3
- [44] Kim Mouridsen, Søren Christensen, Louise Gyldensted, and Leif Østergaard. Automatic selection of arterial input function using cluster analysis. *Magnetic Resonance in Medicine*, 55(3):524–531, 2006. 2
- [45] Pratik Mukherjee, Jeffrey Berman, Stephen W. Chung, Christopher Hess, and Roland Henry. Diffusion tensor MR imaging and fiber tractography: Theoretic underpinnings. *American Journal of Neuroradiology*, 29(4):632–641, 2008. 7
- [46] Nicholas H. Nelsen and Andrew M. Stuart. The random feature model for input-output maps between banach spaces. *arXiv preprint arXiv:2005.10224*, 2020. 1, 2
- [47] Marc Niethammer, Raul San Jose Estepar, Sylvain Bouix, Martha Shenton, and Carl-Fredrik Westin. On diffusion tensor estimation. In *2006 International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 2622–2625, 2006. 3
- [48] Marc Niethammer, Gabriel L Hart, and Christopher Zach. An optimal control approach for the registration of image time-series. In *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pages 2427–2434. IEEE, 2009. 2
- [49] B. Øksendal. *Stochastic Differential Equations: An Introduction with Applications*. Universitext. Springer Berlin Heidelberg, 2010. 3, 14
- [50] Sasitorn Petcharunpaisan, Joana Ramalho, and Mauricio Castillo. Arterial spin labeling in neuroimaging. *World journal of radiology*, 2(10):384–398, 2010. 2

- [51] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686 – 707, 2019. 1, 2
- [52] K.M. Schmainda, M.A. Prah, L.S. Hu, C.C. Quarles, N. Semmineh, S.D. Rand, J.M. Connelly, B. Anderies, Y. Zhou, Y. Liu, B. Logan, A. Stokes, G. Baird, and J.L. Boxerman. Moving toward a consensus DSC-MRI protocol: Validation of a low-flip angle single-dose option as a reference standard for brain tumors. *American Journal of Neuroradiology*, 2019. 2
- [53] K.M. Schmainda, M.A. Prah, Z. Zhang, B.S. Snyder, S.D. Rand, T.R. Jensen, D.P. Barboriak, and J.L. Boxerman. Quantitative delta T1 (dT1) as a replacement for adjudicated central reader analysis of contrast-enhancing tumor burden: A subanalysis of the american college of radiology imaging network 6677/radiation therapy oncology group 0625 multi-center brain tumor trial. *American Journal of Neuroradiology*, 2019. 2
- [54] Zhengyang Shen, Xu Han, Zhenlin Xu, and Marc Niethammer. Networks for joint affine and non-parametric image registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4224–4233, 2019. 2, 3
- [55] Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *arXiv preprint arXiv:2006.09661*, 2020. 1, 2
- [56] Jonathan D. Smith, Kamyar Azizzadenesheli, and Zachary E. Ross. Eikonet: Solving the eikonal equation with deep neural networks. *arXiv preprint arXiv:2004.00361*, 2020. 1, 2
- [57] Costas Strouthos, Marios Lampaskis, Vassilis Sboros, Alan Mcneilly, and Michalakis Averkiou. Indicator dilution models for the quantification of microvascular blood flow with bolus administration of ultrasound contrast agents. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 57(6):1296–1310, 2010. 2
- [58] Deqing Sun, Stefan Roth, and Michael J. Black. Secrets of optical flow estimation and their principles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2432–2439, 2010. 2
- [59] Alexandre M. Tartakovsky, Carlos Ortiz Marrero, Paris Perdikaris, Guzel D. Tartakovsky, and David Barajas-Solano. Physics-informed deep neural networks for learning parameters and constitutive relationships in subsurface flow problems. *Water Resources Research*, 56(5):e2019WR026731, 2020. 1, 2, 5
- [60] Demian Wassermann, Matthew Toews, Marc Niethammer, and William Wells. Probabilistic diffeomorphic registration: Representing uncertainty. In Sébastien Ourselin and Marc Modat, editors, *Biomedical Image Registration*, pages 72–82, Cham, 2014. Springer International Publishing. 3, 14
- [61] Philippe Weinzaepfel, Jérôme Revaud, Zaid Harchaoui, and Cordelia Schmid. Deepflow: Large displacement optical flow with deep matching. In *2013 IEEE International Conference on Computer Vision*, pages 1385–1392, 2013. 3
- [62] Xiao Yang, Roland Kwitt, Martin Styner, and Marc Niethammer. Quicksilver: Fast predictive image registration—a deep learning approach. *NeuroImage*, 158:378–396, 2017. 2, 3
- [63] Qihao Zhang, Pascal Spincemaille, Michele Drotman, Christine Chen, Sarah Eskreis-Winkler, Weiyuan Huang, Liangdong Zhou, John Morgan, Thanh D. Nguyen, Martin R. Prince, and Yi Wang. Quantitative transport mapping (QTM) for differentiating benign and malignant breast lesion: Comparison with traditional kinetics modeling and semi-quantitative enhancement curve characteristics. *Magnetic Resonance Imaging*, 2021. 1, 2
- [64] Liangdong Zhou, Pascal Spincemaille, Qihao Zhang, Thanh D. Nguyen, and Yi Wang. Vector field perfusion imaging: A validation study by using multiphysics model. *Proc. Intl. Soc. Mag. Recon. Med. (ISMRM)*, 26, 2018. 2
- [65] Liangdong Zhou, Qihao Zhang, Pascal Spincemaille, Thanh D. Nguyen, John Morgan, Weiyang Dai, Yi Li, Ajay Gupta, Martin R. Prince, and Yi Wang. Quantitative transport mapping (QTM) of the kidney with an approximate microvascular network. *Magnetic Resonance in Medicine*, 85(4):2247–2262, 2021. 2