


# SUPER-IVIM-DC-BOOT: A Bootstrapped Physics-Informed IVIM Framework for Robust Placental Microstructure Analysis in Uncontrolled Maternal Diabetes

Naama Gavrielov<sup>1,2</sup> 

NAAMAGAV@CAMPUS.TECHNION.AC.IL

Moran Gawie-Rotman<sup>3,4</sup>

MORAN.GAWIE7@GMAIL.COM

Abdel-Rauf Zeina<sup>4,5</sup>

RAUFZEINA@GMAIL.COM

Roni Shreter<sup>5</sup>

RONISH@HYMC.GOV.IL

Daphna Link-Sourani<sup>1,2</sup>

LDAPHNA@BM.TECHNION.AC.IL

Esther Maor-Sagie<sup>3,4</sup>

ESTIMAORSAGIE@GMAIL.COM

Rinat Gabbay-Benziv<sup>3,4</sup>

GABBAYRINAT@GMAIL.COM

Moti Freiman<sup>1,2</sup> 

MOTI.FREIMAN@TECHNION.AC.IL

<sup>1</sup> Faculty of Biomedical Engineering, Technion – Israel Institute of Technology, Haifa, Israel

<sup>2</sup> The May-Blum-Dahl Human MRI Research Center, Faculty of Biomedical Engineering, Technion – Israel Institute of Technology, Haifa, Israel

<sup>3</sup> Department of Obstetrics and Gynecology, Hillel Yaffe Medical Center, Hadera, Israel

<sup>4</sup> The Ruth and Bruce Rappaport Faculty of Medicine, Technion - Israel Institute of Technology, Haifa, Israel

<sup>5</sup> Department of Radiology, Hillel Yaffe Medical Center, Hadera, Israel

**Editors:** Under Review for MIDL 2026

## Abstract

Intravoxel Incoherent Motion (IVIM) modeling decomposes diffusion-weighted MRI (DWI) signals into diffusion- and perfusion-related components, enabling non-invasive characterization of microvascular structure in highly vascularized tissues such as the placenta. However, accurate recovery of IVIM parameters remains an ill-posed inverse problem, particularly under the low signal-to-noise ratio (SNR) and sparse b-value sampling common in fetal and placental imaging. We introduce SUPER-IVIM-DC-BOOT, an acquisition-aware, physics-informed neural network that integrates explicit b-value encoding with inference-time bootstrap resampling to stabilize bi-exponential parameter estimation. By aggregating predictions from stratified subsampled inputs and enforcing a data-consistent forward-model constraint, the method improves robustness to noise and protocol sparsity. Numerical simulations demonstrate that bootstrap aggregation substantially reduces Normalized Root Mean Squared Error (NRMSE) in low-SNR regimes (SNR 7–10) relative to non-bootstrapped baselines. In a healthy volunteer study, SUPER-IVIM-DC-BOOT recovered IVIM parameters from a sparse 9 b-value protocol with accuracy comparable to reference estimates obtained from a dense 16 b-value acquisition. Applied to placental DWI in pregnancies with uncontrolled maternal diabetes ( $N = 6$ ) versus matched controls ( $N = 8$ ), the method detected significant alterations in perfusion fraction ( $f$ ) and tissue diffusivity ( $D$ ) ( $p < 0.05$ ), whereas conventional Trust-Region Reflective fitting (SLS-TRF) underestimated perfusion changes and failed to identify restricted diffusivity. These findings demonstrate that combining protocol-aware physics constraints with inference-time bootstrapping yields robust and clinically meaningful IVIM quantification under realistic, sparse acquisition conditions. Code will be released upon acceptance.

**Keywords:** Physics-Informed Deep Learning, Diffusion-Weighted MRI, IVIM (Intravoxel Incoherent Motion), bootstrapping, Maternal Diabetes, Placenta Microstructure

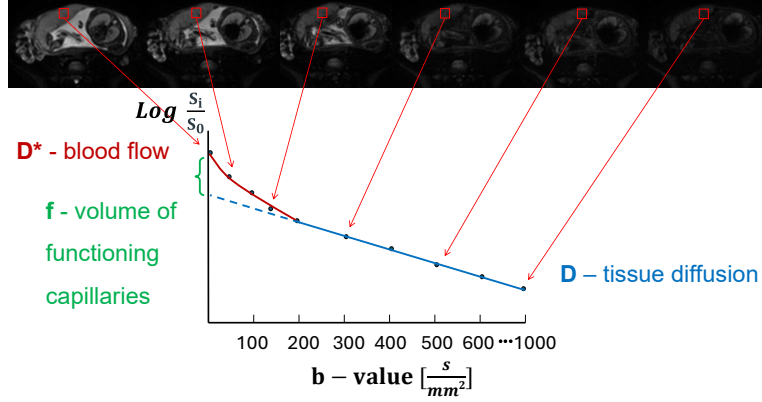


Figure 1: Schematic illustration of the IVIM model in the placenta. The diffusion-weighted signal exhibits a rapid low-  $b$  decay driven by pseudo-diffusion ( $D^*$ , red), modulated by the perfusion fraction  $f$  (green), followed by a slower decay dominated by tissue diffusion  $D$  (blue). Example DW-MRI images of the placenta at increasing  $b$  - values illustrate the corresponding signal attenuation.

## 1. Introduction

Diffusion-weighted MRI (DWI) provides a non-invasive method to probe tissue microstructure by sensitizing the MR signal to the random displacement of water molecules. In highly vascularized organs such as the placenta (Turco and Moffett, 2019), the diffusion-weighted signal reflects not only thermal Brownian motion but also microcirculatory flow within the capillary network. The Intravoxel Incoherent Motion (IVIM) framework separates these contributions into true diffusion and perfusion-related pseudo-diffusion within a single voxel (Le Bihan et al., 1986, 1988), using a bi-exponential model:

$$S(b) = S_0 \left( f \cdot e^{-b(D^*+D)} + (1-f) \cdot e^{-bD} \right) \quad (1)$$

where  $b$  is the diffusion weighting;  $D$  is the tissue diffusion coefficient;  $D^*$  represents pseudo-diffusion associated with microvascular flow; and  $f$  is the perfusion fraction (Fig. 1).

Despite its physiological interpretability, IVIM parameter estimation is notoriously challenging. The bi-exponential decay in Eq. 1 is fundamentally ill-posed and highly sensitive to measurement noise, especially for perfusion-related parameters ( $f$  and  $D^*$ ), which are governed by the rapid signal drop at low  $b$ -values (Koopman et al., 2021). These challenges are amplified in placental DWI, where fetal and maternal motion (Zhang et al., 2019) impose strict limits on scan duration, leading to sparse sampling and low signal-to-noise ratio (SNR). Conventional optimization-based fits under these conditions frequently yield unstable, noisy, or physiologically implausible parameter maps (Park et al., 2017).

Several strategies have been proposed to stabilize this inverse problem. The standard Segmented Least Squares (SLS) formulation, often paired with Trust Region Reflective op-

timization (Coleman and Li, 1996), restricts the solution space but remains fragile under sparse and heterogeneous sampling. Deep learning (DL) methods offer improved robustness, typically incorporating the bi-exponential signal model through a data-consistency term (Bertleff et al., 2017; Barbieri et al., 2020; Kaandorp et al., 2021). SUPER-IVIM-DC (Korngut et al., 2022) further augments this with *supervised* training on simulated parameters, combining physical constraints with explicit ground-truth guidance.

However, existing IVIM networks—SUPER-IVIM-DC included—*do not explicitly condition on the acquisition b-values*. This omission limits their ability to generalize across the diverse diffusion protocols used in clinical and research settings, particularly in placental imaging, where inherently *sparse and variable b-value* sampling is common. Because IVIM estimation depends critically on the relative distribution of low- and high-*b* measurements, ignoring protocol information can lead to systematic bias and unstable parameter recovery.

To address this limitation, we introduce **SUPER-IVIM-DC-BOOT**, a physics-informed neural network (Raissi et al., 2019) with two key technical innovations. First, the model explicitly conditions on the acquired b-values, enabling acquisition-aware parameter estimation that adapts to heterogeneous clinical protocols. Second, we introduce a stratified subsampling strategy during training that supports inference-time bootstrap resampling, substantially enhancing stability under the extreme noise and sparsity characteristic of placental DWI. This bootstrap mechanism is tailored to the IVIM setting by aggregating parameter predictions arising from different low-/high-*b* subsets, improving perfusion and diffusion estimate reliability.

We demonstrate the clinical added value of our approach in assessing pregnancies with uncontrolled maternal diabetes, a condition linked to placental malperfusion and villous structural abnormalities (Scifres et al., 2017; Jirkovská et al., 2012). Such impairments contribute to adverse fetal outcomes, including elevated risk for Large for Gestational Age (LGA), delivery complications, and, in severe cases, stillbirth (Dodd et al., 2007; Kong et al., 2019). Our aim is to enable non-invasive, *in utero* characterization of placental microstructural alterations in pregnancies with uncontrolled maternal diabetes using sparse and clinically practical diffusion protocols.

## Contributions

- **Methodological Innovation: SUPER-IVIM-DC-BOOT**, an acquisition-aware, physics-informed neural network that explicitly incorporates b-values and leverages stratified subsampling and inference-time bootstrapping to stabilize IVIM parameter estimation under sparse, low-SNR conditions.
- **Comprehensive Validation:** Through numerical simulations, healthy volunteer data, and placenta-specific acquisition configurations, we demonstrate improved stability and lower estimation error relative to conventional fitting and prior deep learning baselines.
- **Clinical Impact:** In a cohort of LGA pregnancies with uncontrolled maternal diabetes, our method detects significant ( $p < 0.05$ ) alterations in placental perfusion and diffusion in comparison to control subjects that standard SLS-TRF fitting does not capture, highlighting its potential for a sensitive assessment of placental dysfunction.

## 2. Methods

We propose **SUPER-IVIM-DC-BOOT**, a physics-informed deep learning framework for robust IVIM parameter estimation ( $\theta = \{f, D^*, D\}$ ). The method introduces two components: (1) explicit b-value encoding for protocol generalization and (2) intrinsic bootstrap resampling for variance reduction.

### 2.1. Protocol-Aware Network Architecture

To enable generalization across acquisition protocols, the model uses an **acquisition-aware** input. Given a vector of observed signals  $\mathbf{S}$  and their corresponding b-values  $\mathbf{b}$ , the input vector is constructed as:

$$\mathbf{x} = \text{concat}(\mathbf{S}, \mathbf{b}), \quad (2)$$

allowing the MLP to learn how IVIM signal decay depends on diffusion weighting:

$$\hat{\theta} = \mathcal{F}_{\Phi}(\mathbf{x}), \quad (3)$$

where  $\mathcal{F}$  represents the mapping function parameterized by the learnable weights and biases  $\Phi$ . The final layer applies a **scaled Sigmoid** activation, constraining predicted parameters to physiologically valid ranges and preventing non-physical values such as negative diffusivities.

### 2.2. Training Strategy: Stratified Subsampling and Hybrid Loss

To enable bootstrap inference, we use **stratified random subsampling** of the simulated decay curves (originally comprising  $N$  measurement points). Non-zero b-values are partitioned into low and high-diffusion regimes,  $\mathcal{I}_{low}$  and  $\mathcal{I}_{high}$ . For each training sample, a sampling index set (with cardinality  $M < N$ )

$$\mathcal{K} = \{0\} \cup \mathcal{K}_{low} \cup \mathcal{K}_{high} \quad (4)$$

is generated. This set always includes the  $b = 0$  measurement and randomly selects points from each regime ( $\mathcal{K}_{low} \subset \mathcal{I}_{low}$  and  $\mathcal{K}_{high} \subset \mathcal{I}_{high}$ ). The mask  $\mathcal{K}$  is applied to both the signal intensities and their corresponding b-values. These subsets are concatenated to form input vectors of size  $2M$ . The resulting inputs are fed through a shuffled data loader, exposing the network to diverse b-value combinations.

The network is trained by minimizing a hybrid loss function (Fig. 2), adapted from the SUPER-IVIM-DC framework (Korngut et al., 2022):

$$\mathcal{L}_{total} = \mathcal{L}_{supervised} + \alpha_{dc} \mathcal{L}_{dc}. \quad (5)$$

**Supervised loss:**

$$\mathcal{L}_{supervised} = \alpha_f \|\hat{f} - f_{ref}\|^2 + \alpha_{D^*} \|\hat{D}^* - D_{ref}^*\|^2 + \alpha_D \|\hat{D} - D_{ref}\|^2, \quad (6)$$

with weights  $\alpha_{\{f, D^*, D\}}$  chosen to balance gradients across parameters.

**Data Consistency loss:**

$$\mathcal{L}_{dc} = \|\mathbf{R} - \mathbf{S}_{\mathcal{K}}\|^2, \quad (7)$$

where  $\mathbf{R} \in \mathbb{R}^M$  is the forward-model reconstruction from  $\hat{\theta}$  and  $\mathbf{S}_{\mathcal{K}} \in \mathbb{R}^M$  is the vector of subsampled signal intensities. This term enforces biophysical plausibility even when labels are noisy, with the hyperparameter  $\alpha_{dc}$  controlling its relative contribution to the total loss.

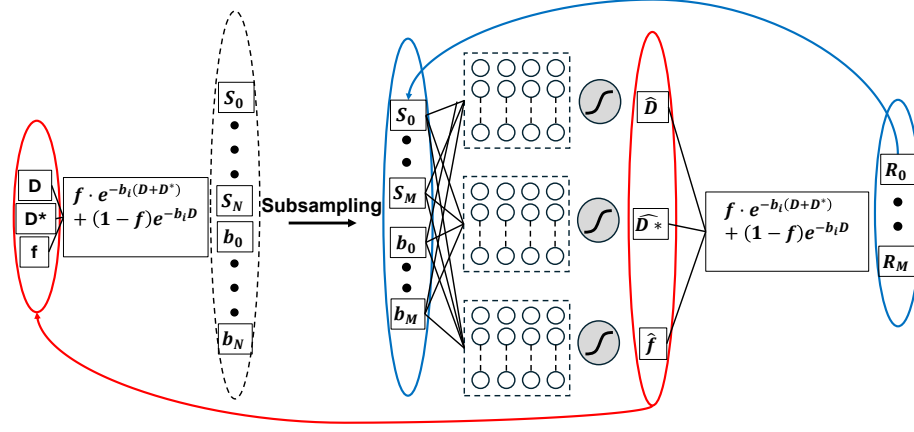


Figure 2: **Training Scheme of SUPER-IVIM-DC-BOOT.** The network receives paired inputs of subsampled signals  $\{S_i\}$  and b-values  $\{b_i\}$ . The model is optimized via a dual objective: a supervised loss against ground truth parameters (red) and a Data Consistency term (blue) that enforces the physical reconstruction of the input signal.

### 2.3. Inference: Bootstrap Resampling

At inference time, we apply a **bootstrap resampling** pipeline (Fig. 3) inspired by classical statistical resampling (Efron, 1979). Given a full clinical protocol of size  $N$ , the signal is subsampled  $P$  times to a size of  $M$ :

1. **Subsampling:** Generate  $P$  random index sets of size  $M$  from the acquired signal.
2. **Prediction:** Construct input vectors by concatenating the subsampled signals with their corresponding b-values, and process them via the trained network.
3. **Aggregation:** Average the  $P$  predicted parameter sets to obtain the final estimate.

## 3. Experimental Setup

### 3.1. Numerical Simulations

We generated 600,000 synthetic IVIM signals using the forward model (Eq. 1). Parameters were uniformly sampled from physiological ranges ( $f \in [0.05, 0.55]$ ,  $D \in [0.5, 3.0] \times 10^{-3} \text{ mm}^2/\text{s}$ ,  $D^* \in [10, 100] \times 10^{-3} \text{ mm}^2/\text{s}$ ). Rician noise was added across six SNR levels (7, 10, 15, 20, 30, 50). The simulation utilized a protocol of  $N = 9$  b-values (0, 10, 20, 40, 80, 200, 400, 600, 1000 s/mm<sup>2</sup>). The dataset was partitioned into 90% for training and 10% for validation.

**Optimization:** Grid search identified an optimal subsampling mask of  $M = 7$  points:

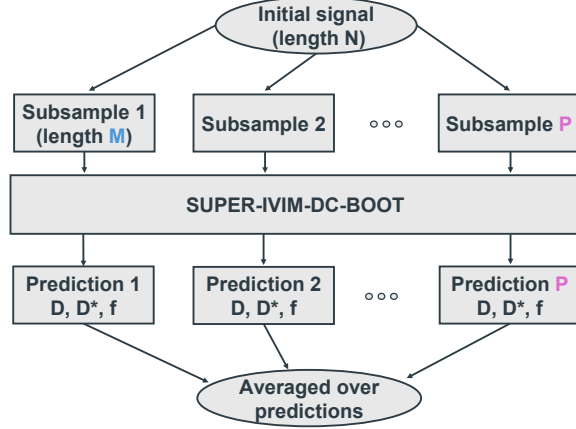


Figure 3: **Inference Pipeline: Bootstrap Resampling.** To reduce estimation uncertainty, the initial signal (length  $N$ ) is randomly subsampled  $P$  times. These subsets are processed by the shared weights of the **SUPER-IVIM-DC-BOOT** network, and the resulting parameter estimates are averaged to yield the final prediction.

the  $b = 0$  signal, three low  $b$ -values ( $< 200$ ), and three high  $b$ -values ( $\geq 200$ ). Training used PyTorch with Adam optimizer ( $\text{lr} = 10^{-4}$ ), batch size 128, and early stopping with a patience of 10 epochs.

A separate test set of 3,000 signals at low SNR (7, 10, 15) was generated to compare SUPER-IVIM-DC-BOOT with SUPER-IVIM-DC. For inference, SUPER-IVIM-DC-BOOT used  $P = 16$  bootstrap samples. Accuracy was reported using Normalized RMSE (NRMSE).

### 3.2. Healthy Volunteer Validation

To assess performance on sparse clinical data, we acquired dense DWI from a healthy volunteer using 16  $b$ -values (0–1000  $\text{s/mm}^2$ ).

**ROI Analysis:** Four ROIs were defined (Liver, Right Kidney, Left Kidney, Spleen) with estimated SNRs of 22.8, 53.9, 49.8, and 43.4, respectively.

**Evaluation Protocol:** Reference parameters were obtained from the full dataset using Segmented Least Squares (SLS). The data were then subsampled to the 9-point clinical protocol, and parameters were estimated using SUPER-IVIM-DC-BOOT, SUPER-IVIM-DC, and SLS-TRF. NRMSE was computed relative to the dense SLS reference.

### 3.3. Clinical Study: Pregnancies with Uncontrolled Maternal Diabetes

We evaluated clinical utility in 14 pregnant subjects scanned at Hillel-Yaffe Medical Center: 6 pregnancies with uncontrolled maternal diabetes and LGA fetuses and 8 matched controls. Placental DWI was acquired using the same 9  $b$ -value protocol. The placenta was manually segmented with ITK-SNAP, and mean diffusion signals were extracted per subject.

**Data Quality:** The data exhibited low SNR (2.6–8.3; mean  $4.4 \pm 1.9$ ), a regime where




Figure 4: **Numerical Simulation Results.** The NRMSE of the IVIM model parameter estimates obtained with SUPER-IVIM-DC-BOOT (blue) and the baseline SUPER-IVIM-DC (orange) as a function of SNR. Our proposed method demonstrates superior stability in the clinically relevant low-SNR regime (SNR 7-10).

conventional IVIM fitting is unstable.

**Statistical Analysis:** Parameters ( $D, f, D^*$ ) were estimated using SLS-TRF, SUPER-IVIM-DC, and SUPER-IVIM-DC-BOOT. Group differences between LGA and control pregnancies were assessed using Welch’s t-tests.

**Qualitative Analysis:** Parametric maps ( $f, D, D^*$ ) were generated for representative subjects to compare noise suppression and spatial coherence across methods.

## 4. Results

### 4.1. Numerical Simulations

The simulation study evaluated the robustness of the proposed method against noise. Figure 4 summarizes the NRMSE for parameter estimation across varying noise levels.

SUPER-IVIM-DC-BOOT consistently showed the lowest error for all IVIM parameters in the low-SNR regime (SNR 7–10), confirming that bootstrap aggregation improves robustness to noise. At higher SNRs, the performance gap narrows, with the non-bootstrapped baseline performing marginally better, likely due to the reduced variance in the input signal.

### 4.2. Healthy Volunteer Study

Table 1 compares the estimation error (NRMSE) of the sparse fitting methods against the high-quality reference parameters (derived via SLS on dense, 16 b-value data) using the healthy volunteer data.

**SUPER-IVIM-DC-BOOT** demonstrated competitive or superior performance in recovering parameters from the subsampled data. Notably, it achieved the lowest error for all parameters in the Right Kidney, as well as for  $f$  and  $D^*$  in the Left Kidney and  $D$  in the Liver. This suggests that even in high-SNR regimes, the physics-informed bootstrapping strategy offers advantages in complex multi-compartment tissues compared to standard sparse fitting.



Table 1: **Healthy Volunteer Study (Reference: High-Quality SLS)**. NRMSE comparison for parameter recovery from sparse data (9 b-values). The reference parameters were derived using the Segmented Least Squares (SLS) method on the full 16 b-value dataset.

Param.	Method	Liver	R. Kidney	L. Kidney	Spleen
<b>D</b>	SLS-TRF (Sparse)	0.117	0.152	<b>0.149</b>	0.069
	SUPER-IVIM-DC	0.102	0.127	0.182	<b>0.032</b>
	<b>SUPER-IVIM-DC-BOOT</b>	<b>0.065</b>	<b>0.105</b>	0.168	0.181
<b>f</b>	SLS-TRF (Sparse)	0.273	0.513	0.378	<b>0.259</b>
	SUPER-IVIM-DC	<b>0.214</b>	0.300	0.350	0.380
	<b>SUPER-IVIM-DC-BOOT</b>	0.345	<b>0.194</b>	<b>0.271</b>	0.483
<b>D*</b>	SLS-TRF (Sparse)	0.396	6.266	4.258	<b>1.132</b>
	SUPER-IVIM-DC	<b>0.386</b>	3.772	3.440	1.209
	<b>SUPER-IVIM-DC-BOOT</b>	0.410	<b>3.288</b>	<b>3.325</b>	1.331

We attribute the varied performance in these ROIs to their relatively high SNR (22–54), a regime where SUPER-IVIM-DC-BOOT did not uniformly yield the lowest NRMSE. This is consistent with our simulation results, which indicate that the stabilizing benefits of bootstrapping are most pronounced under noisier conditions.

### 4.3. Clinical Study: Diabetic Pregnancies

The most significant findings emerged from the clinical cohort ( $N = 14$ ), where data quality was severely compromised (Mean SNR  $\approx 4.4$ ). Figure 5 illustrates the distribution of placental parameters between the Diabetic (Study) and Healthy (Control) groups.

Table 2 details the statistical significance of these differences. **SUPER-IVIM-DC-BOOT** demonstrated superior sensitivity in this low-SNR regime. It identified a reduction in perfusion fraction ( $f$ ) with higher statistical significance ( $p = 0.016$ ) compared to standard SLS-TRF ( $p = 0.027$ ) and SUPER-IVIM-DC ( $p = 0.019$ ). Crucially, it also successfully detected a significant alteration in tissue diffusivity ( $D$ ) ( $p = 0.022$ ), whereas standard fitting failed to do so ( $p = 0.190$ ). Furthermore, the proposed method showed a strong trend toward significance for pseudo-diffusion ( $D^*$ ) ( $p = 0.065$ ), whereas other methods showed no such trend.

Qualitatively (Fig. 6), SLS-TRF produced noisy maps with visible fitting artifacts, whereas SUPER-IVIM-DC-BOOT generated smoother and more coherent parameter maps that better preserved placental structure. Similar improvements were observed for  $D$  and  $D^*$  (Supplementary Figs. A1–A2).

Together, these results indicate that the proposed bootstrapped physics-informed model enhances parameter stability and improves sensitivity to subtle placental microstructural differences in low-SNR clinical data.



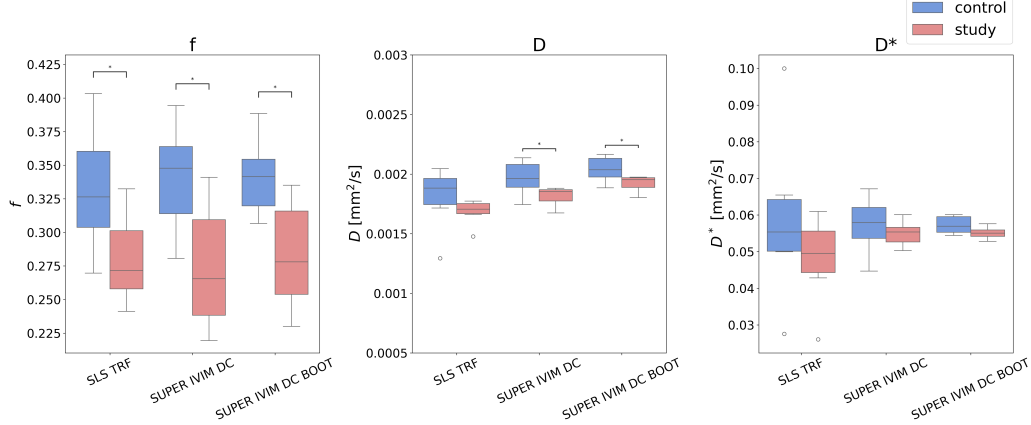


Figure 5: **Clinical Group Differences.** Comparison of placental IVIM parameters between Healthy Controls and LGA Pregnancies estimated by the three methods. \* indicates statistical significance ( $p < 0.05$ ).

Table 2: **Statistical Evaluation (Welch’s t-test).** P-values for pairwise comparisons between Diabetic (Study) and Healthy (Control) groups. \* indicates statistical significance ( $p < 0.05$ ).

Method	$f$ (Perfusion)	$D$ (Diffusion)	$D^*$ (Pseudo-diff.)
SLS-TRF	0.02746*	0.19017	0.24697
SUPER-IVIM-DC	0.01899*	<b>0.0215*</b>	0.38686
<b>SUPER-IVIM-DC-BOOT</b>	<b>0.01622*</b>	0.02192*	<b>0.06492</b>

## 5. Discussion and Conclusion

SUPER-IVIM-DC-BOOT combines explicit b-value encoding with bootstrap resampling to address two central challenges in clinical diffusion MRI: low SNR and heterogeneous acquisition protocols. In simulations, it substantially reduced estimation errors in low-SNR regimes (SNR 7–10), and in a high-SNR healthy volunteer (SNR 22.8–53.9) it improved characterization of multi-compartment organs, demonstrating that acquisition-aware physics constraints benefit both noisy and clean conditions.

A major strength of the framework is its treatment of the acquisition as a distribution rather than a fixed set of measurements. Whereas conventional fits are highly sensitive to individual motion-corrupted points, randomized subsampling and aggregation effectively suppress transient noise and encourage the model to learn the underlying bi-exponential signal behavior. This explains the observed robustness under sparse, clinically constrained protocols.

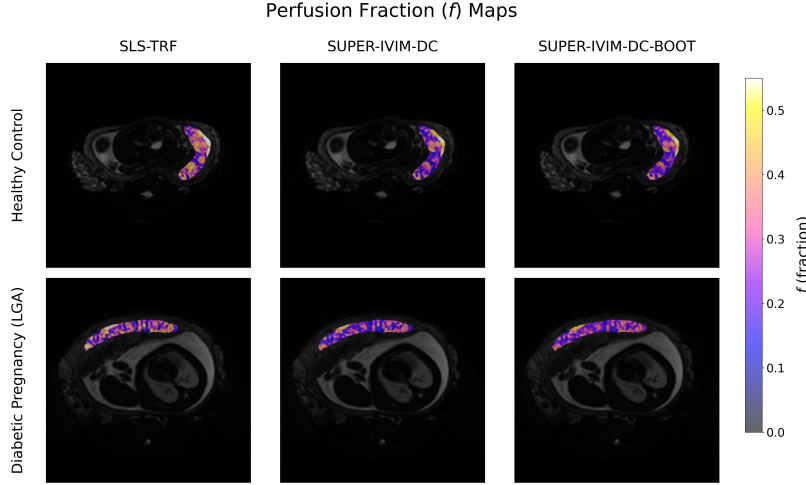


Figure 6: **Representative perfusion fraction ( $f$ ) maps of the placenta.** Comparison of a healthy control (top row) and a pregnancy with uncontrolled maternal diabetes (bottom row) using SLS-TRF, SUPER-IVIM-DC, and SUPER-IVIM-DC-BOOT (left to right). The proposed method improves spatial coherence and anatomical definition, significantly reducing the fitting artifacts and noise observed in the standard SLS-TRF estimation.

Clinically, SUPER-IVIM-DC-BOOT detected significant reductions in perfusion fraction ( $f$ ) and tissue diffusivity ( $D$ ) in pregnancies with uncontrolled maternal diabetes, differences that SLS-TRF failed to capture. These findings highlight its potential for non-invasive, *in utero* assessment of placental health and fetal risk. The resulting parameter maps were smoother and more anatomically coherent than those obtained with conventional fitting, reflecting the stabilizing effect of combining physics constraints with bootstrapping.

In addition to improved accuracy, the method offers practical benefits: while traditional segmented least-squares fitting is iterative and initialization-sensitive, neural network inference remains fast and parallelizable, even when multiple bootstrap passes are used.

**Limitations and Future Work:** The clinical cohort was small ( $N = 14$ ), warranting larger multi-center studies. Motion correction is not yet integrated and will be addressed via the IVIM-Morph framework (Kertes et al., 2025). Our current analysis used whole-placenta averages; future work will examine regional heterogeneity and compartment-specific function.

In summary, SUPER-IVIM-DC-BOOT provides a robust, acquisition-aware solution for IVIM quantification under sparse and noisy diffusion protocols. By stabilizing parameter estimation and detecting subtle placental abnormalities in pregnancies with uncontrolled maternal diabetes, it represents a meaningful step toward non-invasive, *in utero* screening and earlier identification of at-risk fetuses.

## Acknowledgments

This research was supported by the Israeli Ministry of Innovation, Science, and Technology (#1001577479). We extend our sincere gratitude to the clinical staff at Hillel-Yaffe Medical Center for their assistance in data acquisition, and most importantly, to the study participants for their invaluable contribution.

## References

- Sebastiano Barbieri, Oliver J Gurney-Champion, Remy Klaassen, and Harriet C Thoeny. Deep learning how to fit an intravoxel incoherent motion model to diffusion-weighted mri. *Magnetic resonance in medicine*, 83(1):312–321, 2020.
- Marco Bertleff, Sebastian Domsch, Sebastian Weingärtner, Jascha Zapp, Kieran O’Brien, Markus Barth, and Lothar R Schad. Diffusion parameter mapping with the combined intravoxel incoherent motion and kurtosis model using artificial neural networks at 3 t. *NMR in Biomedicine*, 30(12):e3833, 2017.
- Thomas F Coleman and Yuying Li. An interior trust region approach for nonlinear minimization subject to bounds. *SIAM Journal on optimization*, 6(2):418–445, 1996.
- Jodie M Dodd, Caroline A Crowther, Georgia Antoniou, Peter Baghurst, and Jeffrey S Robinson. Screening for gestational diabetes: the effect of varying blood glucose definitions in the prediction of adverse maternal and infant health outcomes. *Australian and New Zealand Journal of Obstetrics and Gynaecology*, 47(4):307–312, 2007.
- B. Efron. Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics*, 7(1):1 – 26, 1979. doi: 10.1214/aos/1176344552. URL <https://doi.org/10.1214/aos/1176344552>.
- M Jirkovská, T Kučera, J Kaláb, M Jadrníček, V Niedobová, J Janáček, L Kubínová, M Moravcová, Z Žizka, and V Krejčí. The branching pattern of villous capillaries and structural changes of placental terminal villi in type 1 diabetes mellitus. *Placenta*, 33(5): 343–351, 2012.
- Misha PT Kaandorp, Sebastiano Barbieri, Remy Klaassen, Hanneke WM van Laarhoven, Hans Crezee, Peter T While, Aart J Nederveen, and Oliver J Gurney-Champion. Improved unsupervised physics-informed deep learning for intravoxel incoherent motion modeling and evaluation in pancreatic cancer patients. *Magnetic resonance in medicine*, 86(4):2250–2265, 2021.
- Noga Kertes, Yael Zaffrani-Reznikov, Onur Afacan, Sila Kurugol, Simon K Warfield, and Moti Freiman. Ivim-morph: Motion-compensated quantitative intra-voxel incoherent motion (ivim) analysis for functional fetal lung maturity assessment from diffusion-weighted mri data. *Medical Image Analysis*, 101:103445, 2025.
- Linghua Kong, Ida AK Nilsson, Mika Gissler, and Catharina Lavebratt. Associations of maternal diabetes and body mass index with offspring birth weight and prematurity. *JAMA pediatrics*, 173(4):371–378, 2019.

- Thomas Koopman, Roland Martens, Oliver J Gurney-Champion, Maqsood Yaqub, Cristina Lavini, Pim de Graaf, Jonas Castelijns, Ronald Boellaard, and J Tim Marcus. Repeatability of ivim biomarkers from diffusion-weighted mri in head and neck: Bayesian probability versus neural network. *Magnetic resonance in medicine*, 85(6):3394–3402, 2021.
- Noam Korngut, Elad Rotman, Onur Afacan, Sila Kurugol, Yael Zaffrani-Reznikov, Shira Nemirovsky-Rotman, Simon Warfield, and Moti Freiman. Super-ivim-dc: Intra-voxel incoherent motion based fetal lung maturity assessment from limited dwi data using supervised learning coupled with data-consistency. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 743–752. Springer, 2022.
- Denis Le Bihan, Eric Breton, Denis Lallemand, Philippe Grenier, Emmanuel Cabanis, and Maurice Laval-Jeantet. Mr imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*, 161(2):401–407, 1986.
- Denis Le Bihan, Eric Breton, Denis Lallemand, Marie-Louise Aubin, Jean Vignaud, and Maurice Laval-Jeantet. Separation of diffusion and perfusion in intravoxel incoherent motion mr imaging. *Radiology*, 168(2):497–505, 1988.
- Hyo Jung Park, Yu Sub Sung, Seung Soo Lee, Yedaun Lee, Hyunhee Cheong, Yeong Jae Kim, and Moon-gyu Lee. Intravoxel incoherent motion diffusion-weighted mri of the abdomen: the effect of fitting algorithms on the accuracy and reliability of the parameters. *Journal of Magnetic Resonance Imaging*, 45(6):1637–1647, 2017.
- Maziar Raissi, Paris Perdikaris, and George E 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.
- Christina M Scifres, W Tony Parks, Maisa Feghali, Steve N Caritis, and Janet M Catov. Placental maternal vascular malperfusion and adverse pregnancy outcomes in gestational diabetes mellitus. *Placenta*, 49:10–15, 2017.
- Margherita Y Turco and Ashley Moffett. Development of the human placenta. *Development*, 146(22):dev163428, 2019.
- Lin Zhang, Valery Vishnevskiy, Andras Jakab, and Orcun Goksel. Implicit modeling with uncertainty estimation for intravoxel incoherent motion imaging. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 1003–1007. IEEE, 2019.

## Appendix A. Supplementary Figures

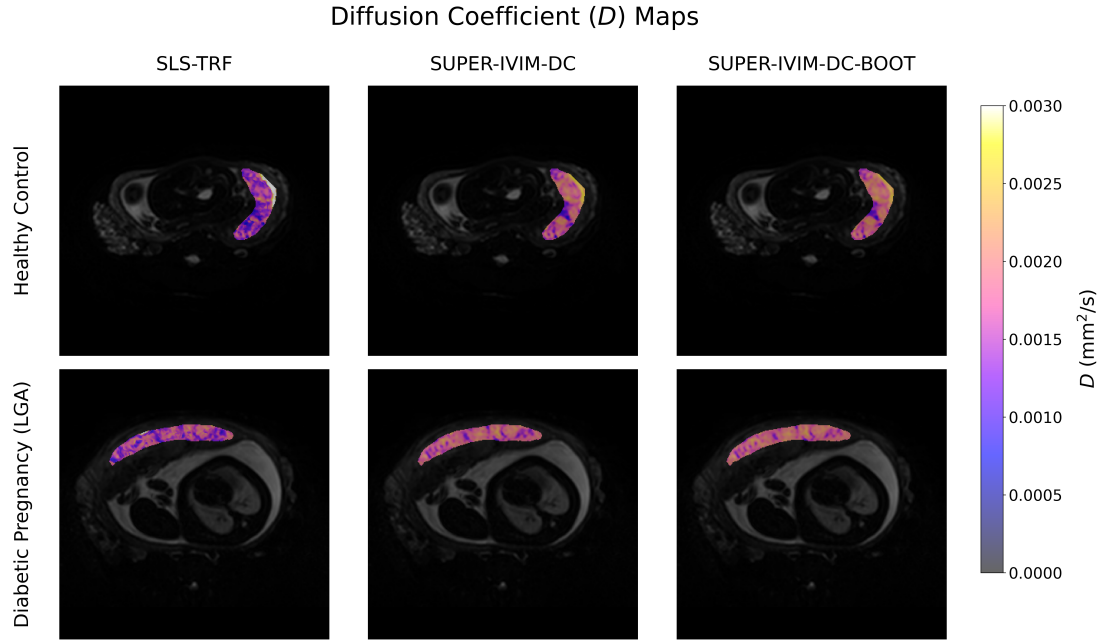


Figure A1: **Visual comparison of Tissue Diffusion ( $D$ ) maps.** Representative maps for a Large for Gestational Age (LGA) pregnancy with uncontrolled maternal diabetes compared to a healthy control. Parameters were estimated using SLS-TRF, SUPER-IVIM-DC, and the proposed SUPER-IVIM-DC-BOOT method.

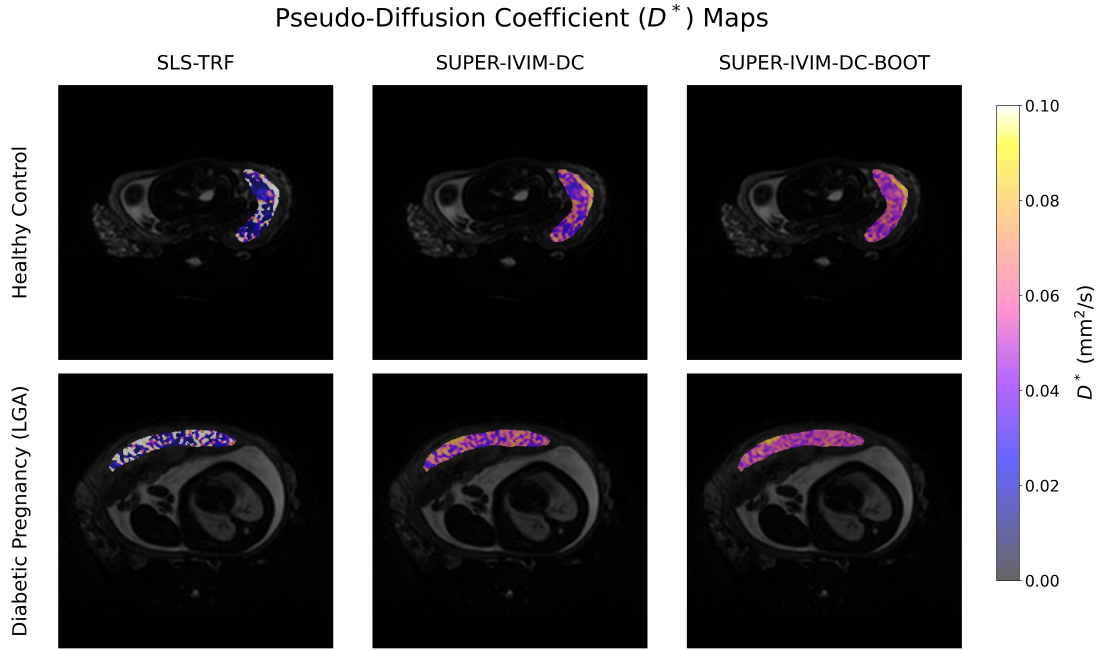


Figure A2: **Visual comparison of Tissue Pseudo-Diffusion ( $D^*$ ) maps.** Representative maps for a Large for Gestational Age (LGA) pregnancy with uncontrolled maternal diabetes compared to a healthy control. Parameters were estimated using SLS-TRF, SUPER-IVIM-DC, and the proposed SUPER-IVIM-DC-BOOT method.