SUPER-IVIM-DC-SUB: A Physics-Informed Subset Ensembling Framework for Robust Placental IVIM Analysis in Uncontrolled Maternal Diabetes

Naama Gavrielov[1,2], Moran Gawie-Rotman[3,4], Abdel-Rauf Zeina[4,5], Roni Shreter[5], Daphna Link-Sourani[1,2], Esther Maor-Sagie[3,4], Rinat Gabbay-Benziv[3,4], and Moti Freiman[1,2]

[1] Faculty of Biomedical Engineering, Technion - Israel Institute of Technology, Haifa, Israel
[2] The May-Blum-Dahl Human MRI Research Center, Faculty of Biomedical Engineering, Technion - Israel Institute of Technology, Haifa, Israel
[3] Department of Obstetrics and Gynecology, Hillel Yaffe Medical Center, Hadera, Israel
[4] The Ruth and Bruce Rappaport Faculty of Medicine, Technion - Israel Institute of Technology, Haifa, Israel
[5] Department of Radiology, Hillel Yaffe Medical Center, Hadera, Israel
naamagav@campus.technion.ac.il

Abstract. Intravoxel Incoherent Motion (IVIM) modeling decomposes Diffusion-Weighted Magnetic Resonance Imaging (DWI) signals into diffusion and pseudo-diffusion components, enabling non-invasive characterization of highly vascularized tissues like the placenta. However, accurate recovery of IVIM parameters remains an ill-posed inverse problem, particularly at low signal-to-noise ratio (SNR) and sparse b-value sampling, common in fetal and placental imaging. We introduce SUPER-IVIM-DC-SUB, an acquisition-aware, physics-informed neural network combining explicit b-value encoding with inference-time subset ensembling. By aggregating predictions from stratified subsamples, the method stabilizes estimation against noise and sparsity. Simulations demonstrate reduced estimation error in low-SNR regimes (3–10) compared to conventional and deep learning baselines. In a healthy volunteer, the method recovered stable parameters from a sparse 9 b-value protocol. In a proof-of-concept clinical application comparing pregnancies complicated by uncontrolled maternal diabetes with large-for-gestational-age (LGA) fetuses ($N = 6$) to matched controls ($N = 8$), the proposed method detected preliminary significant alterations ($p < 0.05$) in pseudo-diffusion fraction (f) and tissue diffusion (D). Notably, conventional fitting (SLS-TRF) detected the alteration in f but missed the decline in D. These results demonstrate that protocol-aware subset ensembling yields robust IVIM quantification under sparse acquisition. Code is available at: https://github.com/TechnionComputationalMRILab/SUPER-IVIM-DC-SUB.

Keywords: Physics-Informed Deep Learning · Diffusion-Weighted MRI · IVIM · Subset Ensembling · Maternal Diabetes · Placenta.

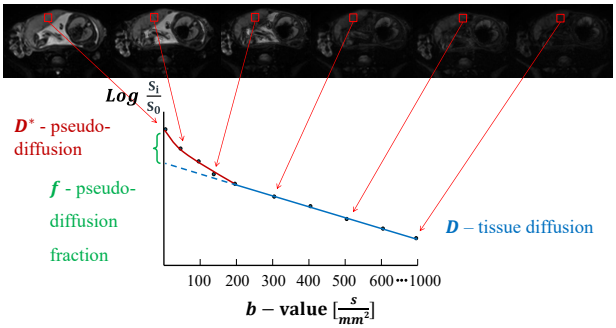

Fig. 1. Placental IVIM schematic. The signal exhibits rapid low-b decay driven mainly by pseudo-diffusion (D^*, red) and pseudo-diffusion fraction (f, green), followed by slower tissue diffusion (D, blue). Example DWI images illustrate corresponding signal attenuation at increasing b-values.

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 like the placenta [20], the Intravoxel Incoherent Motion (IVIM) framework distinguishes thermal diffusion from microcirculatory pseudo-diffusion [15,16] using a bi-exponential model [14]:

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

where b is the diffusion weighting, D is the tissue diffusion coefficient, D^* represents pseudo-diffusion, and f is the pseudo-diffusion fraction (Fig. 1). In this work, the fast pseudo-diffusion exponential term is explicitly parameterized as $(D^* + D)$; all evaluated baselines utilize this equivalent parameterization to ensure fair comparison of D^*.

Parameter estimation for Eq. 1 is fundamentally ill-posed and noise-sensitive [11]. These challenges are amplified in placental DWI, where fetal and maternal motion impose strict limits on scan duration, leading to sparse sampling and low signal-to-noise ratio (SNR) [22]. Conventional optimization-based fits under these conditions frequently yield unstable, noisy, or physiologically implausible parameter maps [17]. While Segmented Least Squares (SLS) with Trust Region Reflective (TRF) optimization [3] improves robustness, recent deep learning (DL) methods [1,2,7,21], such as SUPER-IVIM-DC [12], have demonstrated improved stability. However, standard DL networks typically lack explicit conditioning on acquisition b-values, preventing inference-time subset ensembling (which relies on model awareness of specific diffusion encodings).

To address this, we propose **SUPER-IVIM-DC-SUB**, a physics-informed neural network [18] that explicitly conditions on acquired b-values. This de-

sign enables inference-time stratified subset ensembling (aggregating predictions across distinct subsamples without replacement), substantially improving estimation stability under sparse sampling. We demonstrate this approach in a preliminary clinical cohort of pregnancies complicated by maternal diabetes with large-for-gestational-age (LGA) fetuses, a condition linked to adverse fetal outcomes including delivery complications, pre-eclampsia, and, in severe cases, stillbirth [4,6,10,19]. This complex condition requires sensitive, non-invasive *in utero* assessment.

Contributions: We propose **SUPER-IVIM-DC-SUB**, an acquisition-aware network leveraging explicit b-value conditioning and subset ensembling. We demonstrate its improved stability under low-SNR conditions over baselines in both simulations and *in vivo* data. In a proof-of-concept clinical cohort, our method detects preliminary but significant reductions in placental f and D in diabetic pregnancies with LGA fetuses, whereas conventional fitting misses the alteration in D.

2 Methods

We propose a physics-informed deep learning framework estimating $\theta = \{f, D^*, D\}$ via (1) explicit b-value encoding and (2) intrinsic subset ensembling for variance reduction.

2.1 Protocol-Aware Network Architecture

To enable robustness across varying subset combinations, the network conditions on the specific acquisition parameters. We construct the input $\mathbf{x}$ by concatenating observed signals $\mathbf{S}$ and their corresponding b-values $\mathbf{b}$: $\mathbf{x} = \text{concat}(\mathbf{S}, \mathbf{b})$. A multilayer perceptron (MLP) maps this input to parameters $\hat{\theta} = \mathcal{F}_{\Phi}(\mathbf{x})$.

2.2 Training Strategy: Stratified Subsampling and a Hybrid Loss

Training utilizes **stratified random subsampling** to simulate diverse subprotocols. From the full decay curve (N points), we generate a sampling mask $\mathcal{K}$ of size $M < N$: $\mathcal{K} = \{0\} \cup \mathcal{K}_{low} \cup \mathcal{K}_{high}$, where $\mathcal{K}_{low}$ and $\mathcal{K}_{high}$ are random subsets from the pseudo- and pure-diffusion regimes, respectively. This ensures the network learns from varied but structurally complete b-value combinations. The network minimizes a hybrid loss (Fig. 2) adapted from [12], combining supervised regression on simulated parameters and data consistency:

$$\mathcal{L}_{total} = \mathcal{L}_{supervised} + \alpha_{dc}\,\mathcal{L}_{dc} \tag{2}$$

Supervised Parameter Loss:

$$\mathcal{L}_{supervised} = \alpha_f \|\hat{f} - f_{ref}\|^2 + \alpha_{D^*} \|\hat{D}^* - D^*_{ref}\|^2 + \alpha_D \|\hat{D} - D_{ref}\|^2 \tag{3}$$

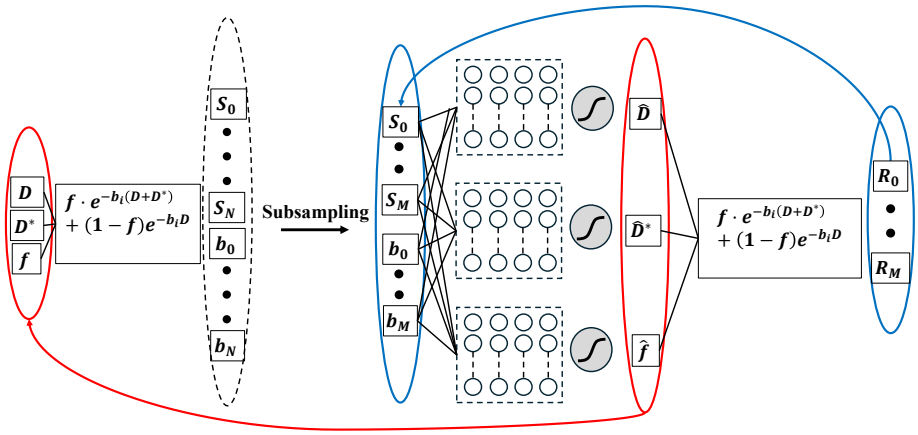

Fig. 2. Training Scheme of SUPER-IVIM-DC-SUB. Paired subsampled signals $\{S_i\}$ and b-values $\{b_i\}$ are optimized via a dual objective: supervised parameter loss (red) and data consistency (blue), enforcing physical reconstruction.

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

$$\mathcal{L}_{dc} = \|\mathbf{R} - \mathbf{S}_{\mathcal{K}}\|^2 \tag{4}$$

where $\mathbf{R}$ is the forward-model reconstruction and $\mathbf{S}_{\mathcal{K}}$ is the subsampled input. This term enforces biophysical plausibility, weighted by α_{dc}.

2.3 Inference: Subset Ensembling (SUB)

To reduce uncertainty during inference, we apply subset ensembling (Fig. 3). The acquired clinical signal (size N) is subsampled P times without replacement to a fixed size M, utilizing combinations of available low and high b-values. The network predicts parameters for each subset independently, and the results are averaged.

2.4 Implementation and Architecture Details

For SUPER-IVIM-DC-SUB, the MLP consists of 4 hidden layers with 14 neurons each, utilizing ELU activations and Batch Normalization. The final output layer uses a scaled Sigmoid activation to constrain predictions to physiologically valid ranges, incorporating a margin to avoid gradient saturation near the boundaries. The network is trained exclusively on simulated data due to the lack of ground truth for *in vivo* data; consequently, no *in vivo* fine-tuning was performed.

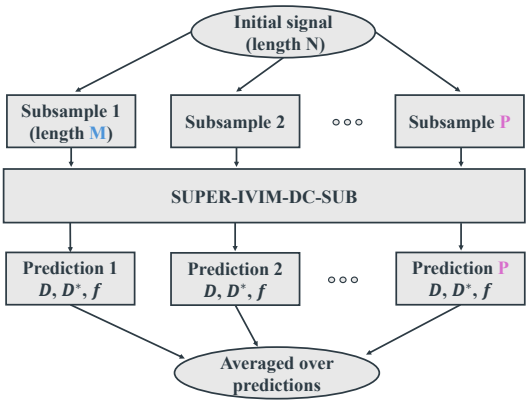

Fig. 3. Subset Ensembling Inference Pipeline. To reduce uncertainty, the initial signal (length N) is subsampled P times. Network predictions for each subset are averaged to yield the final estimate.

3 Experimental Setup

3.1 Baseline Methods and Ablation Strategy

We compared against **SLS-TRF** and the physics-informed **SUPER-IVIM-DC** [12]. To isolate contributions, we evaluated two ablations: **SLS-TRF-SUB** applies subset ensembling to the standard solver; **SUPER-IVIM-DC-BVAL** extends the baseline network with explicit b-value inputs but performs standard single-pass inference using the full 9-point acquisition. In contrast, the full SUB variant averages predictions across $P = 16$ distinct 7-point subsets drawn from the 9-point protocol, cleanly isolating the architectural gain of acquisition awareness (BVAL) from the statistical variance reduction of ensembling (SUB).

3.2 Numerical Simulations

We generated 600,000 synthetic IVIM signals using realistic physiological ranges ($f \in [0.05, 0.55]$, $D \in [0.5, 3.0] \times 10^{-3}$ mm^2/s, $D^* \in [10, 100] \times 10^{-3}$ mm^2/s). Rician noise was added at SNR $\in \{7, 10, 15, 20, 30, 50\}$. The protocol utilized nine b-values (0, 10, 20, 40, 80, 200, 400, 600, 1000 s/mm^2), and the dataset was partitioned into 90% training and 10% validation sets. Models were trained in PyTorch (Adam, lr $= 10^{-4}$) with a subsampling mask of $M = 7$ points (including $b = 0$ s/mm^2, three low $b < 200$ s/mm^2, and three high $b \geq 200$ s/mm^2). Testing evaluated 7,000 low-SNR signals (SNR $\in \{3, 4, 5, 6, 7, 10, 15\}$). Accuracy was evaluated using the Normalized Root Mean Square Error (NRMSE).

Hyperparameter Optimization: The optimal $M = 7$ configuration was determined via grid search. The value of $P = 16$ was specifically chosen to ensure exhaustive coverage of the unique $\binom{4}{3} \times \binom{4}{3}$ subsampling combinations permitted by our acquisition protocol.

Table 1. Demographic and clinical characteristics of the study cohort. Data are presented as median (interquartile range, IQR). Group comparisons were evaluated using the Mann–Whitney U test; *ns*: not significant ($p \geq 0.05$).

Characteristic	Healthy Controls	Diabetic + LGA	p-value
Gestational Age (weeks)	33.2 (27.3-35.3)	34.4 (29.4-37.1)	ns
Maternal BMI	26 (24-27.7)	37.7 (29.7-44.5)	0.008
Estimated Fetal Weight Percentiles	41 (29-71.5)	99 (97.5-99)	0.002
Placental SNR	4.91 (3.30-7.04)	2.98 (2.82-3.55)	ns

3.3 Healthy Volunteer Validation

To establish performance against a stable, high-SNR dense reference estimate, which is ethically and practically challenging to acquire via dense *in utero* scanning, dense DWI was acquired from a healthy adult volunteer ($N = 1$) using 16 b-values (0, 10, 20, 30, 40, 50, 60, 80, 100, 150, 200, 300, 400, 600, 800, 1000 s/mm^2). Four regions of interest (ROIs) provided high-SNR regimes: the liver (SNR 22.8), right kidney (SNR 53.9), left kidney (SNR 49.8), and spleen (SNR 43.4). Reference parameters were estimated from this dense dataset using SLS. We then acquired 10 additional datasets of the same subject using the sparse 9-point protocol (Sec. 3.2). IVIM parameters were calculated for each of the 10 sparse datasets using all evaluated methods, and the resulting estimates were compared against the dense reference via NRMSE.

3.4 Clinical Study: Pregnancies with Uncontrolled Maternal Diabetes and LGA fetuses

This proof-of-concept study included 14 pregnancies (Table 1). Institutional review board (Helsinki Committee) approval was obtained. The study group consisted of six patients with a confirmed clinical diagnosis of maternal diabetes. To specifically study the effects of uncontrolled diabetes with a uniform phenotype, inclusion in this cohort required the presence of fetal overgrowth, characterized as a large-for-gestational-age (LGA) fetus with an estimated fetal weight above the 90th percentile. Eight healthy pregnancies with appropriately grown fetuses served as matched controls. Gestational age at the time of scan and placental SNR showed no significant differences between groups. As expected, significant differences were observed in maternal Body Mass Index (BMI) and estimated fetal weight percentiles.

MRI Acquisition: Placental MRI was performed using a 1.5 T Philips Ingenia scanner (Philips Medical Systems, Best, The Netherlands) equipped with a multi-channel abdominal coil. Parameters included: TR = 4215 ms, TE = 103 ms, voxel size = $1.01 \times 1.01 \times 5.0$ mm^3. The sparse protocol used 9 b-values (0, 10, 20, 40, 80, 200, 400, 600, 1000 s/mm^2). Manual segmentation of the whole placenta was used to extract mean ROI signals; explicit spatial motion registration across b-values was not applied, representing a known limitation. The

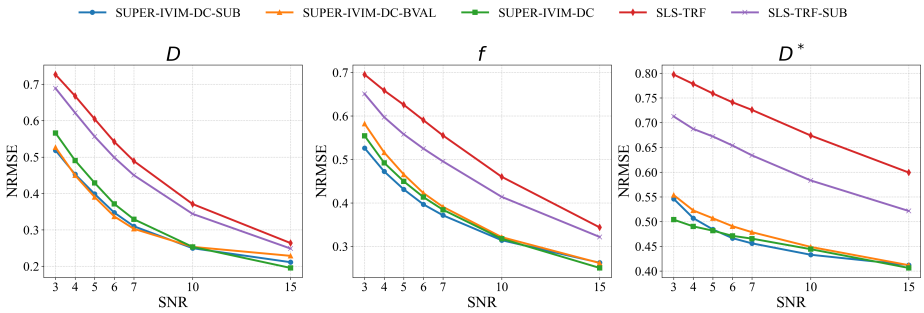

Fig. 4. Numerical Simulation Results. The NRMSE of the IVIM model parameter estimates obtained with the different methods as a function of SNR.

in vivo data exhibited low SNR (4.4 ± 1.9), representing a challenging regime where conventional fitting approaches often prove unstable. We evaluated group differences in the estimated parameters (D, f, D^*) using Welch's t-tests and generated parametric maps to qualitatively assess both noise suppression and spatial coherence.

4 Results

4.1 Numerical Simulations: Error and the Noise Floor

Fig. 4 summarizes NRMSE across noise levels, where all deep learning variants substantially outperformed standard fitting. Among deep learning models in the critical low-SNR regime (3–10), both SUPER-IVIM-DC-SUB and SUPER-IVIM-DC-BVAL achieved lower errors for D than the baseline SUPER-IVIM-DC. Furthermore, SUPER-IVIM-DC-SUB uniquely minimized error for f in this regime, emphasizing the added value of ensembling for noise robustness. Finally, performance for D^* was inconsistent among network variants; however, this is expected given the notoriously difficult and low-sensitivity nature of estimating this parameter.

Beyond the specific benefits of subset ensembling, the overall superiority of the deep learning variants in these low-SNR regimes stems from their handling of the Rician noise floor. The supervised loss mitigates this inherent bias by explicitly training the network to map noisy signals directly to their ground-truth parameters. By learning this mapping from data, the network captures the true biophysical decay shape, preventing it from merely fitting the artificially flattened, noise-biased tail at high b-values, a common pitfall of conventional optimization solvers.

4.2 Healthy Volunteer Study

Table 2 compares sparse reconstruction against a dense 16-point SLS reference. In the high-SNR adult organs (SNR 22–54), SUPER-IVIM-DC-SUB demon-

Table 2. Healthy Volunteer Study (NRMSE). Comparison of sparse (9 b-value) recovery against a dense (16 b-value) SLS reference. Values are mean ± standard deviation.

Param.	Method	Liver	R. Kidney	L. Kidney	Spleen
D	SLS-TRF	0.117 ± 0.109	0.152 ± 0.069	0.149 ± 0.086	0.069 ± 0.069
	SLS-TRF-SUB	0.113 ± 0.096	0.138 ± 0.037	**0.144 ± 0.051**	0.066 ± 0.065
	SUPER-IVIM-DC	0.102 ± 0.061	0.127 ± 0.016	0.182 ± 0.026	**0.032 ± 0.030**
	SUPER-IVIM-DC-BVAL	0.105 ± 0.091	0.203 ± 0.011	0.263 ± 0.017	0.137 ± 0.034
	SUPER-IVIM-DC-SUB	**0.061 ± 0.060**	**0.112 ± 0.013**	0.175 ± 0.021	0.187 ± 0.024
f	SLS-TRF	0.273 ± 0.265	0.513 ± 0.180	0.378 ± 0.224	0.259 ± 0.251
	SLS-TRF-SUB	0.290 ± 0.236	0.448 ± 0.100	0.338 ± 0.157	**0.196 ± 0.184**
	SUPER-IVIM-DC	**0.214 ± 0.070**	0.300 ± 0.014	0.350 ± 0.039	0.380 ± 0.042
	SUPER-IVIM-DC-BVAL	0.317 ± 0.101	0.263 ± 0.017	0.325 ± 0.033	0.506 ± 0.034
	SUPER-IVIM-DC-SUB	0.338 ± 0.063	**0.192 ± 0.017**	**0.270 ± 0.030**	0.485 ± 0.022
D^*	SLS-TRF	0.396 ± 0.279	6.266 ± 3.310	4.258 ± 3.475	1.132 ± 0.897
	SLS-TRF-SUB	0.407 ± 0.265	5.325 ± 1.916	**3.228 ± 2.089**	**1.098 ± 0.699**
	SUPER-IVIM-DC	0.386 ± 0.082	3.772 ± 0.372	3.440 ± 0.488	1.209 ± 0.242
	SUPER-IVIM-DC-BVAL	**0.386 ± 0.075**	3.511 ± 0.197	3.419 ± 0.347	1.393 ± 0.186
	SUPER-IVIM-DC-SUB	0.414 ± 0.093	**3.256 ± 0.097**	3.290 ± 0.140	1.315 ± 0.106

strated improved stability, achieving the lowest NRMSE in the right kidney and for specific parameters in the left kidney (f) and liver (D), with consistently low standard deviations. While the method did not uniformly yield the lowest error across all metrics in this high-SNR regime, this aligns with our simulations, which demonstrate that subset ensembling offers the greatest gain under noisier conditions typical of placental imaging. The ablation study confirms the benefit of ensembling: both subset variants (SLS-TRF-SUB and SUPER-IVIM-DC-SUB) generally outperformed their non-ensembled counterparts. Finally, as expected, D^* estimation exhibited substantial inaccuracies (with NRMSE values exceeding 3.0, corresponding to relative errors of over 300%) across all methods, reinforcing that pseudo-diffusion remains fundamentally difficult to pinpoint even in ideal conditions.

4.3 Clinical Study: Diabetic + LGA Pregnancies

Fig. 5 illustrates the distribution of placental parameters between the Diabetic + LGA (Study) and Healthy (Control) groups. Computationally, inference requires ≈85 ms per subject (averaged signal) on a standard CPU, facilitating real-time clinical workflows.

Table 3 details the statistical significance, effect sizes, and 95% confidence intervals (CI) between groups. SUPER-IVIM-DC-SUB identified a reduced pseudo-diffusion fraction (f, $p = 0.016$, mean difference: -0.060, 95% CI: $[-0.105, -0.014]$). Crucially, it also detected significant alterations in tissue diffusion (D, $p = 0.022$, mean difference: -0.00012, 95% CI: $[-0.00022, -0.00002]$), which standard fitting (SLS-TRF) entirely missed ($p = 0.190$, 95% CI crossing zero: $[-0.00034, 0.00008]$).

Analysis of the statistical results highlights the quantitative advantages of our acquisition-aware ensembling framework. While all methods report reductions in f, SUPER-IVIM-DC-SUB identifies this alteration with the smallest

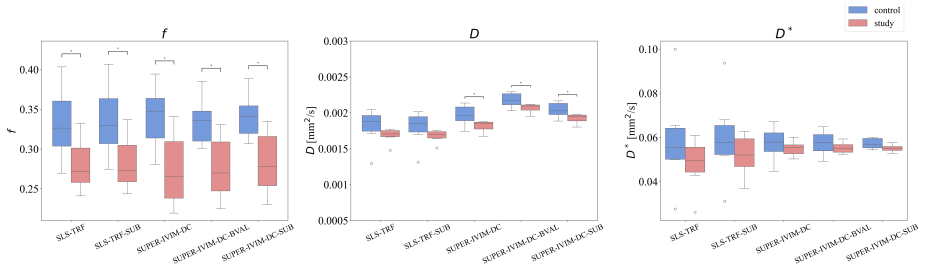

Fig. 5. Clinical Group Differences. Comparison of placental IVIM parameters between Healthy Controls and Diabetic + LGA Pregnancies estimated by the five methods. ***** indicates statistical significance ($p < 0.05$).

Table 3. Clinical Group Differences. P-values, mean differences (Effect Size), and 95% Confidence Intervals comparing Diabetic + LGA to Healthy groups. ***** indicates $p < 0.05$.

	p-value (Welch's t-test)			Effect Size (Mean Diff) [95% CI]		
Method	f	D	D^*	f	D	D^*
SLS-TRF	0.027*	0.190	0.247	-0.052 [-0.096, -0.007]	-0.00013 [-0.00034, 0.00008]	-0.01075 [-0.03002, 0.00853]
SLS-TRF-SUB	0.027*	0.192	0.312	-0.053 [-0.098, -0.007]	-0.00012 [-0.00032, 0.00007]	-0.00786 [-0.02414, 0.00843]
SUPER-IVIM-DC	0.019*	**0.022***	0.387	-0.068 [-0.121, -0.014]	-0.00016 [-0.00028, -0.00003]	-0.00261 [-0.00901, 0.00378]
SUPER-IVIM-DC-BVAL	0.017*	0.025*	0.312	-0.060 [-0.106, -0.014]	-0.00011 [-0.00021, -0.00002]	-0.00222 [-0.00685, 0.00240]
SUPER-IVIM-DC-SUB	**0.016***	**0.022***	**0.065**	-0.060 [-0.105, -0.014]	-0.00012 [-0.00022, -0.00002]	-0.00217 [-0.00450, 0.00016]

p-value, achieving a narrower confidence interval than SUPER-IVIM-DC while maintaining a comparable effect magnitude. Furthermore, for the notoriously noisy pseudo-diffusion parameter D^*, the ensembling framework dramatically narrows the confidence interval width by nearly an order of magnitude compared to traditional fitting (SLS-TRF span: 0.039 vs. SUPER-IVIM-DC-SUB span: 0.0047), directly validating its variance-reduction efficacy under sparse clinical acquisition. Qualitatively, the DL methods also yielded physically smoother, more anatomically coherent parametric maps (Fig. 6).

5 Discussion and Conclusion

SUPER-IVIM-DC-SUB addresses low SNR and protocol heterogeneity in sparse placental DWI by combining explicit b-value encoding with subset ensembling. By treating the acquisition as a distribution of sub-protocols, the model suppresses transient noise, yielding stability superior to standard fitting and non-ensembled deep learning. Ablation confirms this robustness relies on the synergy of b-value encoding (flexibility) and subset ensembling (noise mitigation). Clinically, in a proof-of-concept cohort, the method identified significant reductions in placental f and D. While conventional fittings successfully captured the decline in f, they completely missed the alteration in D; our framework captured both, achieving the smallest p-value for f ($p = 0.016$). Additionally, it produces

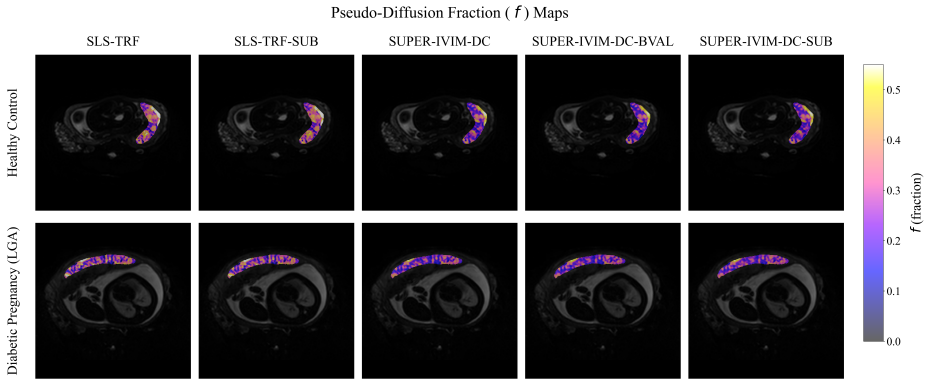

Fig. 6. Representative pseudo-diffusion fraction (f) maps of the placenta. Comparison of healthy control (top) and uncontrolled maternal diabetes with LGA fetuses (bottom) across methods (left to right). The proposed method improves spatial coherence and reduces fitting artifacts relative to standard SLS-TRF.

smoother parameter maps with rapid inference times suitable for clinical workflows.

Our approach mitigates severe Rician noise floor bias at high b-values through the supervised learning of clean target parameters. Estimating D^* remains highly unreliable under sparse, low-SNR conditions, often generating substantial inaccuracies even in high-SNR adult reference tissues. Consequently, we position f and D as the primary reliable clinical biomarkers.

Biophysically, we acknowledge that standard IVIM is a phenomenological model. In the placenta, the bi-exponential decomposition conflates fetal perfusion, maternal blood space, and inherent diffusion heterogeneity into overlapping terms. Improved numerical stability does not strictly imply improved biophysical specificity, highlighting the need for advanced multi-compartment formulations.

Limitations and Future Work: These preliminary clinical findings ($N = 14$) warrant larger validation. Additionally, while our ensembling strategy demonstrates robustness within a predefined sparse grid, it does not inherently generalize to arbitrarily long, unseen acquisition protocols. A vital direction for future work is a direct empirical comparison between our subset ensembling approach and recent sequence-based architectures that process irregular b-value sequences, such as transformer-based estimators [8] and neural controlled differential equations [13]. To thoroughly disentangle systematic bias from precision, a known limitation of relying solely on NRMSE [5], future evaluations will incorporate predicted-versus-ground-truth scatter plots across varying SNR levels, median percentage error (MPE) to quantify directional bias, and median absolute percentage error (MAPE) for robust error magnitude. Furthermore, to verify that our method preserves genuine spatial heterogeneity rather than reverting to a smooth training prior, subsequent studies will evaluate quantitative spatial metrics using spatially correlated synthetic ground truth, such as fractal-noise

parameter phantoms, alongside structural similarity metrics (SSIM). Finally, future iterations will integrate explicit spatial motion correction via the IVIM-Morph framework [9] and leverage ensembling distributions to extract voxel-wise uncertainty intervals for regional placental mapping.

In conclusion, SUPER-IVIM-DC-SUB provides a robust, acquisition-aware framework for IVIM quantification. By stabilizing parameter estimation in sparse and noisy regimes, it shows promising potential to facilitate sensitive, non-invasive *in utero* assessment of placental health, laying the groundwork for future clinical tools aimed at identifying 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.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Barbieri, S., Gurney-Champion, O.J., Klaassen, R., Thoeny, H.C.: Deep learning how to fit an intravoxel incoherent motion model to diffusion-weighted mri. Magnetic resonance in medicine **83**(1), 312–321 (2020)
2. Bertleff, M., Domsch, S., Weingärtner, S., Zapp, J., O'Brien, K., Barth, M., Schad, L.R.: 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)
3. Coleman, T.F., Li, Y.: An interior trust region approach for nonlinear minimization subject to bounds. SIAM Journal on optimization **6**(2), 418–445 (1996)
4. Dodd, J.M., Crowther, C.A., Antoniou, G., Baghurst, P., Robinson, J.S.: 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)
5. Epstein, S.C., Bray, T.J., Hall-Craggs, M., Zhang, H.: Choice of training label matters: how to best use deep learning for quantitative mri parameter estimation. arXiv preprint arXiv:2205.05587 (2022)
6. Jirkovská, M., Kučera, T., Kaláb, J., Jadrníček, M., Niedobová, V., Janáček, J., Kubínová, L., Moravcová, M., Žižka, Z., Krejčí, V.: The branching pattern of villous capillaries and structural changes of placental terminal villi in type 1 diabetes mellitus. Placenta **33**(5), 343–351 (2012)
7. Kaandorp, M.P., Barbieri, S., Klaassen, R., van Laarhoven, H.W., Crezee, H., While, P.T., Nederveen, A.J., Gurney-Champion, O.J.: 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)
8. Kaandorp, M.P., Jakab, A., Federau, C., While, P.T.: A comparative study of ivim-mri fitting techniques in glioma grading: Conventional, bayesian, and voxel-wise and spatially-aware deep learning approaches. Journal of Magnetic Resonance Imaging **64**(1), 184–200 (2026)

9. Kertes, N., Zaffrani-Reznikov, Y., Afacan, O., Kurugol, S., Warfield, S.K., Freiman, M.: 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)
10. Kong, L., Nilsson, I.A., Gissler, M., Lavebratt, C.: Associations of maternal diabetes and body mass index with offspring birth weight and prematurity. JAMA pediatrics **173**(4), 371–378 (2019)
11. Koopman, T., Martens, R., Gurney-Champion, O.J., Yaqub, M., Lavini, C., de Graaf, P., Castelijns, J., Boellaard, R., Marcus, J.T.: 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)
12. Korngut, N., Rotman, E., Afacan, O., Kurugol, S., Zaffrani-Reznikov, Y., Nemirovsky-Rotman, S., Warfield, S., Freiman, M.: 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. pp. 743–752. Springer (2022)
13. Kuppens, D., Barbieri, S., Berg, D.v.d., Schouten, P., Thoeny, H.C., Wennen, M., Gurney-Champion, O.J.: Acquisition-independent deep learning for quantitative mri parameter estimation using neural controlled differential equations. arXiv preprint arXiv:2412.20844 (2024)
14. Le Bihan, D.: What can we see with ivim mri? Neuroimage **187**, 56–67 (2019)
15. Le Bihan, D., Breton, E., Lallemand, D., Aubin, M.L., Vignaud, J., Laval-Jeantet, M.: Separation of diffusion and perfusion in intravoxel incoherent motion mr imaging. Radiology **168**(2), 497–505 (1988)
16. Le Bihan, D., Breton, E., Lallemand, D., Grenier, P., Cabanis, E., Laval-Jeantet, M.: Mr imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. Radiology **161**(2), 401–407 (1986)
17. Park, H.J., Sung, Y.S., Lee, S.S., Lee, Y., Cheong, H., Kim, Y.J., Lee, M.g.: 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)
18. Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational physics **378**, 686–707 (2019)
19. Scifres, C.M., Parks, W.T., Feghali, M., Caritis, S.N., Catov, J.M.: Placental maternal vascular malperfusion and adverse pregnancy outcomes in gestational diabetes mellitus. Placenta **49**, 10–15 (2017)
20. Turco, M.Y., Moffett, A.: Development of the human placenta. Development **146**(22), dev163428 (2019)
21. Vasylechko, S.D., Warfield, S.K., Afacan, O., Kurugol, S.: Self-supervised ivim dwi parameter estimation with a physics based forward model. Magnetic resonance in medicine **87**(2), 904–914 (2022)
22. Zhang, L., Vishnevskiy, V., Jakab, A., Goksel, O.: Implicit modeling with uncertainty estimation for intravoxel incoherent motion imaging. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). pp. 1003–1007. IEEE (2019)