

Evidential Perfusion Physics-Informed Neural Networks with Residual Uncertainty Quantification

Junhyeok Lee¹, Minseo Choi², Han Jang³, Young Hun Jeon², Heeseong Eum¹,
Joon Jang⁴, Chul-Ho Sohn², and Kyu Sung Choi^{2,5,6}

¹ Cancer Biology, Seoul National University College of Medicine, Seoul, Korea

² Radiology, Seoul National University Hospital, Seoul, Korea

³ Bioengineering, Seoul National University, Seoul, Korea

⁴ Biomedical Sciences, Seoul National University, Seoul, Korea

⁵ Radiology, Seoul National University College of Medicine, Seoul, Korea

⁶ Healthcare AI Research Institute, Seoul National University Hospital, Seoul, Korea
`ent1127@snu.ac.kr`

Abstract. Physics-informed neural networks (PINNs) have shown promise in addressing the ill-posed deconvolution problem in computed tomography perfusion (CTP) imaging for acute ischemic stroke assessment. However, existing PINN-based approaches remain deterministic and do not quantify uncertainty associated with violations of physics constraints, limiting reliability assessment. We propose Evidential Perfusion Physics-Informed Neural Networks (EPPINN), a framework that integrates evidential deep learning with physics-informed modeling to enable uncertainty-aware perfusion parameter estimation. EPPINN models arterial input, tissue concentration, and perfusion parameters using coordinate-based networks, and places a Normal-Inverse-Gamma distribution over the physics residual to characterize voxel-wise aleatoric and epistemic uncertainty in physics consistency without requiring Bayesian sampling or ensemble inference. The framework further incorporates physiologically constrained parameterization and stabilization strategies to promote robust per-case optimization. We evaluate EPPINN on digital phantom data, the ISLES 2018 benchmark, and a clinical cohort. On the evaluated datasets, EPPINN achieves lower normalized mean absolute error than classical deconvolution and PINN baselines, particularly under sparse temporal sampling and low signal-to-noise conditions, while providing conservative uncertainty estimates with high empirical coverage. On clinical data, EPPINN attains the highest voxel-level and case-level infarct-core detection sensitivity. These results suggest that evidential physics-informed learning can improve both accuracy and reliability of CTP analysis for time-critical stroke assessment. Source code is available at <https://github.com/jhlee0619/EPPINN>.

Keywords: CT Perfusion · Physics-Informed Neural Networks · Evidential Deep Learning · Uncertainty Quantification

1 Introduction

Acute ischemic stroke management is time-critical, with reperfusion therapies indicated within 4.5 hours for intravenous thrombolysis and within 6 hours for mechanical thrombectomy, extending up to 24 hours in selected patients [19, 2]. Rapid differentiation between the irreversible ischemic core and the salvageable penumbra is therefore essential for triage [9]. Computed Tomography Perfusion (CTP) is widely used to quantify hemodynamic parameters such as cerebral blood flow (CBF), cerebral blood volume (CBV), and mean transit time (MTT) [3]. This quantification relies on deconvolution, an ill-posed inverse problem sensitive to noise and arterial input function (AIF) variability [8]. Traditional singular value decomposition (SVD) pipelines are efficient but unstable under low signal-to-noise ratios or sparse temporal sampling [16, 23]. Although deep learning models improve robustness by learning direct mappings from time-attenuation curves [11], purely data-driven approaches may violate tracer-kinetic constraints and yield physiologically implausible or poorly generalizable estimates across acquisition protocols, potentially compromising time-critical treatment decisions.

Physics-informed neural networks (PINNs) incorporate tracer-kinetic equations as soft constraints and have improved numerical stability in CTP through spatio-temporal formulations that explicitly enforce the convolution model [17, 6, 7]. Nevertheless, existing PINN-based approaches remain deterministic and yield only point estimates, without quantifying uncertainty arising from imperfect physics constraint satisfaction—an important limitation in acute stroke triage, where treatment decisions depend on reliable parameter confidence [2]. This limitation is further amplified in clinical CTP, which is inherently noisy and often sparsely sampled in time, increasing residual ambiguity and destabilizing optimization. Although Bayesian PINNs and ensemble methods can provide uncertainty estimates [24, 10], their computational overhead limits practicality for rapid, per-case clinical deployment.

To this end, we propose Evidential Perfusion Physics-Informed Neural Networks (EPPINN), an uncertainty-aware framework that integrates evidential learning [20, 4] with physics-informed modeling for CTP analysis. Unlike deterministic PINNs, EPPINN introduces a probabilistic formulation of physics constraint deviations to enable explicit uncertainty estimation while retaining a single optimization procedure suitable for per-case deployment. The framework further incorporates optimization stabilization strategies to ensure robust performance under noisy and sparsely sampled clinical conditions. We validate the proposed approach on synthetic and clinical datasets.

Our contributions are summarized as follows:

- We introduce an uncertainty-aware physics-informed framework for CTP that models tracer-kinetic constraint deviations probabilistically, yielding voxel-wise aleatoric and epistemic uncertainty estimates without additional computational overhead.




Fig. 1. Overview of EPPINN. Coordinate-based networks model AIF and tissue curves, and a parameter network predicts perfusion parameters and evidential residual uncertainty.

- We propose a stable per-case optimization scheme that integrates evidential residual learning with a structured CBV–MTT parameterization and progressive annealing, ensuring robust convergence in noisy clinical settings.
- Through extensive evaluation on digital phantoms, the ISLES 2018 benchmark, and clinical data, we demonstrate that EPPINN improves perfusion parameter accuracy and infarct-core detection sensitivity compared to classical deconvolution and state-of-the-art PINN baselines.

2 Method

2.1 CTP Forward Model and Physics Constraints

Let $t \in \mathbb{R}$ denote continuous time and $\mathbf{x} \in \mathbb{R}^3$ the voxel coordinate. In CTP, the tissue concentration curve $C(t, \mathbf{x})$ at voxel $\mathbf{x}$ is related to the AIF $C_a(t)$ through the convolution model [16]

$$C(t, \mathbf{x}) = \text{CBF}(\mathbf{x}) \int_0^t C_a(\tau - \Delta t(\mathbf{x})) R(t - \tau, \mathbf{x}) d\tau, \quad (1)$$

where $\text{CBF}(\mathbf{x})$ denotes cerebral blood flow, $\Delta t(\mathbf{x})$ the transit delay, and $R(\cdot, \mathbf{x})$ the residue function, so that the impulse response function (IRF) is given by $\text{IRF}(t, \mathbf{x}) = \text{CBF}(\mathbf{x}) R(t, \mathbf{x})$. Following prior spatio-temporal PINN-based CTP formulations [6, 7], we adopt a box residue function, i.e., $R(s, \mathbf{x}) = 1$ for $0 \leq s \leq \text{MTT}(\mathbf{x})$ and 0 otherwise. Substituting this into Eq. (1) and applying the Leibniz rule yields the endpoint-difference constraint and physics residual:

$$\frac{\partial C(t, \mathbf{x})}{\partial t} = \text{CBF}(\mathbf{x}) [C_a(t - \Delta t(\mathbf{x})) - C_a(t - \Delta t(\mathbf{x}) - \text{MTT}(\mathbf{x}))], \quad (2)$$

$$r(t, \mathbf{x}) := \frac{\partial C(t, \mathbf{x})}{\partial t} - \text{CBF}(\mathbf{x}) [C_a(t - \Delta t(\mathbf{x})) - C_a(t - \Delta t(\mathbf{x}) - \text{MTT}(\mathbf{x}))]. \quad (3)$$

Under the governing constraint, the residual $r(t, \mathbf{x})$ should vanish. To enforce this constraint within a neural approximation framework, let $f_{\text{AIF}}(t)$ and $f_{\text{tissue}}(t, \mathbf{x})$ denote neural approximations of $C_a(t)$ and $C(t, \mathbf{x})$, respectively (see Sec. 2.2). In practice, we evaluate $r(t, \mathbf{x})$ by substituting C with f_{tissue} and computing the temporal derivative via automatic differentiation.

In addition to the physics residual, we define an L_1 data-fidelity loss on the arterial input function and tissue concentration curves, evaluated at the acquired time points and sampled voxel locations:

$$\mathcal{L}_{\text{data}} = \|C_a - f_{\text{AIF}}\|_1 + \|C - f_{\text{tissue}}\|_1. \quad (4)$$

2.2 Architecture and Evidential Formulation

EPPINN consists of three coordinate-based networks (Fig. 1): the AIF and tissue networks f_{AIF} and f_{tissue} , and a parameter network f_{par} that predicts CBV, MTT, and Δt . The flow parameter CBF is derived from (CBV, MTT) following the CBV–MTT parameterization (Sec. 2.3). Spatial coordinates $\mathbf{x}$ are encoded by a multi-resolution hash grid [13] $h(\mathbf{x})$. The tissue network takes $[t, h(\mathbf{x})]$ as input, and the parameter network takes $h(\mathbf{x})$ as input. For brevity, we write $f_{\text{tissue}}(t, \mathbf{x})$ and $f_{\text{par}}(\mathbf{x})$.

To enable uncertainty estimation, f_{par} additionally outputs evidential parameters that parameterize a Normal–Inverse–Gamma (NIG) distribution over the physics residual $r(t, \mathbf{x})$. The NIG formulation provides computationally efficient, closed-form uncertainty decomposition under the prior $\sigma^2 \sim \text{Inv}\Gamma(\alpha, \beta)$ and $\mu \sim \mathcal{N}(0, \sigma^2/\nu)$ [4]. Specifically, the network predicts $(\tilde{\alpha}, \tilde{\beta}, \tilde{\nu})$, which are transformed to valid domains by: $\alpha = 1 + \text{softplus}(\tilde{\alpha}) + \epsilon_{\text{edl}}$, $\beta = \text{softplus}(\tilde{\beta}) + \epsilon_{\text{edl}}$, $\nu = \text{softplus}(\tilde{\nu}) + \epsilon_{\text{edl}}$.

Evidential residual learning and uncertainty. We treat the physics residual as an evidential regression target with $y_r := 0$ under this NIG prior [4]. Since the governing constraint requires $r(t, \mathbf{x}) = 0$, fixing the predictive mean ($\gamma = 0$) enforces physics consistency while allowing (α, β, ν) to encode uncertainty in constraint satisfaction. We optimize the standard NIG negative log-likelihood with an anti-degeneracy regularizer:

$$\mathcal{L}_{\text{NLL}} = -\alpha \log \beta + \left(\alpha + \frac{1}{2}\right) \log\left(\beta + \frac{\nu}{2} r^2\right) + C(\alpha, \nu), \quad (5)$$

where $C(\alpha, \nu)$ collects terms independent of r , and

$$\mathcal{L}_{\text{reg}} = |r| (2\nu + \alpha). \quad (6)$$

The NIG prior promotes small residuals while quantifying remaining discrepancies as aleatoric and epistemic components. Voxel-wise uncertainty maps are obtained via the closed-form decomposition [4]:

$$\hat{\sigma}_{\text{ale}}^2 = \frac{\beta}{\alpha - 1}, \quad \hat{\sigma}_{\text{epi}}^2 = \frac{\hat{\sigma}_{\text{ale}}^2}{\nu}, \quad \hat{\sigma}_{\text{total}}^2 = \hat{\sigma}_{\text{ale}}^2 + \hat{\sigma}_{\text{epi}}^2,$$

providing physics-consistency uncertainty in a single forward pass.

Objective and annealing. The overall objective combines data fidelity (Eq. 4), residual minimization, and evidential regularization:

$$\mathcal{L} = \lambda_{\text{data}} \mathcal{L}_{\text{data}} + \lambda_{\text{res}} (\|r\|_1 + \omega(i) \lambda_{\text{EDL}} [\mathcal{L}_{\text{NLL}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}]), \quad (7)$$

where $\omega(i) \in [0, 1]$ is a progressive annealing schedule that gradually increases the influence of the evidential terms, enabling initial curve fitting before enforcing physics consistency and thereby improving optimization stability.

2.3 Parameterization and Optimization Stability

To improve numerical stability and identifiability in the ill-posed inverse problem, we incorporate structured parameterization and training strategies. **(1) CBV–MTT parameterization.** The perfusion parameters satisfy the Central Volume Principle [16] $\text{CBV}(\mathbf{x}) = \text{CBF}(\mathbf{x}) \cdot \text{MTT}(\mathbf{x})$. We therefore predict CBV and MTT as primary outputs and derive $\text{CBF}(\mathbf{x}) = \text{CBV}(\mathbf{x}) / (\text{MTT}(\mathbf{x}) + \epsilon_{\text{cbf}})$ by construction. This enforces physiological consistency and reduces ambiguity between flow and transit time parameters. The time-to-maximum (T_{max}) is derived as $T_{\text{max}}(\mathbf{x}) = \Delta t(\mathbf{x}) + \frac{1}{2} \text{MTT}(\mathbf{x})$. **(2) FOV-aware adaptive encoding.** Spatial coordinates are normalized by the physical field-of-view prior to hash encoding. This maintains consistent effective frequency representation across voxel resolutions and scanning protocols, mitigating scale-dependent instability during optimization. **(3) AIF pre-training.** The AIF network f_{AIF} is first fitted to the global arterial signal. This reduces gradient variance during early training and stabilizes subsequent joint optimization. **(4) Physics-informed initialization.** The parameter network f_{par} is initialized to physiologically plausible values (e.g., $\text{CBF} \approx 20 \text{ ml}/100\text{g}/\text{min}$), anchoring optimization in a realistic regime and reducing convergence to implausible local minima. Together, these strategies improve robustness under sparse temporal sampling and heterogeneous acquisition protocols.

3 Experiments

3.1 Datasets

We evaluated our method on three datasets: (i) a digital brain perfusion phantom dataset [1, 18] with known ground-truth perfusion and timing parameters across controlled noise and temporal sampling regimes, (ii) the ISLES 2018 benchmark [12] (n=93) with 4D CTP and diffusion-weighted imaging (DWI) core masks (ROI evaluation), and (iii) a retrospective clinical cohort (n=42) with 4D CTP, DWI core masks, and AIF signals. AIF ground truth is from manual ROIs (clinical), RAPID-derived curves [6] (ISLES), and the analytic curve (phantom); cohorts are partitioned at the patient level. The clinical scans were acquired with in-plane resolution 0.4 mm, slice thickness 10 mm, temporal sampling interval 2.5 s (range 1.0–2.8 s), and 36 time points.




Fig. 2. Digital phantom results under controlled noise and temporal sampling. Left: NMAE for CBF/CBV across PSNR and δt . Right: empirical coverage of uncertainty intervals. Spearman $\rho(\sigma_r, |r|) = 0.66$ with $\pm 1\sigma/\pm 2\sigma$ coverage 87.0%/95.6% (conservative, monotone).

3.2 Baselines and Implementation Details

We compare against standard SVD [16], block-circulant SVD (bcSVD) [23], box-shaped non-linear regression (boxNLR) [5], SPPINN [6], and ReSPPINN [7]. All networks use three hidden layers with sinusoidal representation networks (SIREN; $\omega_0 = 15$) [21] and multi-resolution hash encoding, with hidden dimensions 128 for tissue, 64 for parameters, and 16 for AIF. Optimization uses Adam with a learning rate of 10^{-3} with a OneCycle learning-rate schedule, 25k spatiotemporal samples per iteration, and 5k iterations per case. The AIF network is pre-trained for 5k iterations using only the L_1 data term (Eq. 4), followed by joint optimization under the full objective (Eq. 7). Loss weights are $\lambda_{\text{data}} = \lambda_{\text{res}} = 1$, $\lambda_{\text{EDL}} = 0.5$, and $\lambda_{\text{reg}} = 10^{-3}$ with linear annealing, and $\epsilon_{\text{cbf}} = \epsilon_{\text{edl}} = 10^{-3}$. Experiments were conducted on a single NVIDIA RTX 3090 graphics processing unit (GPU) using tiny-cuda-nn [14], requiring 50s per case (AIF 5–6s, joint optimization 42s). Hyperparameters were fixed on the phantom and reused on ISLES and clinical as held-out tests, with learnable baselines sharing an identical budget.

3.3 Digital Phantom Experiments

Digital phantom experiments enable controlled evaluation under varying noise levels (PSNR 18–27 dB) and temporal sampling intervals ($\delta t \in 1, 2, 3, 4$, s). Accuracy is quantified using normalized mean absolute error (NMAE) across PSNR and δt conditions, and uncertainty calibration is assessed via empirical $\pm 1\sigma$ and $\pm 2\sigma$ coverage (68%/95%). Figure 2 presents NMAE trends over PSNR and temporal resolutions, together with coverage results for two representative settings ($\delta t = 1$ s and $\delta t = 3$ s). EPPINN consistently achieves lower NMAE than competing methods across noise and sampling regimes, with more pronounced gains under sparse temporal sampling ($\delta t = 3$ –4 s) and low-SNR conditions. Empirical coverage generally exceeds nominal levels, particularly outside the lowest PSNR range, indicating conservative uncertainty estimates.

Table 1. Infarct core detection performance on ISLES and a clinical cohort. Det. denotes a case with non-empty overlap between the predicted $\text{CBF} < 25 \text{ ml}/100 \text{ g}/\text{min}$ core and the DWI core.

Method	ISLES (n=93)			Clinical (n=42)		
	Case-level		Voxel-level	Case-level		Voxel-level
	Det. / Total	Sens. $\uparrow$	Sens. $\uparrow$	Det. / Total	Sens. $\uparrow$	Sens. $\uparrow$
SVD	32 / 93	0.344	0.195	12 / 42	0.286	0.125
bcSVD	36 / 93	0.387	0.262	27 / 42	0.643	0.395
boxNLR	38 / 93	0.409	0.243	38 / 42	0.905	0.406
SPPINN	20 / 93	0.215	0.066	29 / 42	0.690	0.304
ReSPPINN	16 / 93	0.172	0.061	28 / 42	0.667	0.304
EPPINN	43 / 93	0.462	0.323	41 / 42	0.976	0.742




Fig. 3. Qualitative perfusion maps on ISLES (left) and a clinical case (right): EPPINN, ReSPPINN, boxNLR, and SVD (CBF/CBV/MTT/ $T_{\max}$).

3.4 Clinical Experiments

We report voxel- and case-level detection sensitivity at a CBF threshold of $25 \text{ ml}/100 \text{ g}/\text{min}$ [15]; case-level sensitivity emphasizes reduction of missed cases, critical for small or heterogeneous cores [22]. EPPINN attains the highest voxel- and case-level sensitivity across both datasets in Table 1, demonstrating superior performance on the ISLES benchmark and correctly identifying 41 of 42 clinical cases. At this threshold, EPPINN’s voxel-level specificity/precision are 0.79/0.04 (ISLES) and 0.54/0.18 (clinical), in a sensitivity-prioritized regime.

Figure 3 shows qualitative comparisons. EPPINN produces coherent perfusion maps on ISLES with less speckle than ReSPPINN/boxNLR; clinical maps are spatially concordant with DWI-defined cores, while baselines show noise or over-diffusion. These qualitative observations are consistent with the improved detection sensitivity reported in Table 1.

Table 2. Ablation on digital phantom data (CBF/MTT NMAE across ROIs). ROIs: GMR (gray matter reduced), GMSR (gray matter severely reduced), WMR (white matter reduced), WMSR (white matter severely reduced). †: unstable training excluded from ranking.

Method / Strategy	CBF					MTT				
	GMR	GMSR	WMR	WMSR	Avg	GMR	GMSR	WMR	WMSR	Avg
Full (proposed)	<u>0.66</u>	<u>1.73</u>	<u>1.27</u>	2.76	<u>1.61</u>	0.59	0.41	0.64	0.51	0.54
w/o Adaptive-Hash	0.94	1.83	1.74	<u>2.67</u>	1.80	0.73	1.35	0.64	1.06	0.94
w/o AIF-Pre.†	5.49	9.33	7.89	12.90	8.90	0.27	0.16	0.35	0.10	0.22
w/o Annealing†	7.51	5.49	4.86	3.88	5.44	4.91	6.99	3.93	5.52	5.34
w/o CBV-Param.	0.82	2.20	1.79	3.14	1.99	<u>0.71</u>	<u>0.79</u>	<u>0.82</u>	<u>0.87</u>	<u>0.80</u>
w/o Phys-init.	0.38	1.65	1.00	3.30	1.59	0.86	2.06	0.98	1.75	1.41




Fig. 4. Training dynamics and residual uncertainty. Left: iterative CBF evolution across methods. Right: final CBF comparison between EPPINN and PINN (EPPINN without evidential loss), with aleatoric, epistemic, and total residual uncertainty maps. $\hat{\sigma}_{\text{ale}}$ dominates under low SNR, while $\hat{\sigma}_{\text{epi}}$ is elevated near vessel/tissue boundaries.

3.5 Ablation Studies

Figure 4 illustrates training dynamics and residual uncertainty. The left panel shows that EPPINN exhibits smoother and more stable CBF convergence with reduced high-frequency fluctuations compared with the PINN baselines. The right panel compares EPPINN with its variant without the evidential loss, which produces noisier and less coherent maps; the aleatoric, epistemic, and total uncertainty localize elevated uncertainty near vascular structures and tissue boundaries.

We ablate engineering and evidential components on digital phantom data at PSNR 18 and $\delta t = 1$. Performance is evaluated using NMAE normalized by the ROI mean (Table 2). Across ROIs, the full model achieves the best MTT performance among stable configurations, while removing CBV-based parameterization consistently yields the second-best MTT results. For CBF, removing PI initialization improves performance in several regions, whereas the full model remains second-best on average with fewer convergence failures across cohorts, highlighting a stability/accuracy trade-off.

4 Conclusion

We presented EPPINN, an uncertainty-aware CT perfusion framework integrating evidential deep learning with physics-informed neural networks. Through evidential residual modeling with physiologically constrained parameterization, EPPINN provides voxel-wise aleatoric and epistemic confidence estimates with stable per-case optimization. Across digital phantom and clinical datasets, EPPINN achieves competitive or improved performance relative to classical deconvolution and recent PINN baselines, indicating enhanced accuracy and reliability for CT perfusion analysis in acute stroke assessment. Bayesian/ensemble UQ baselines and multicenter validation remain future work.

Acknowledgments. This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2026-25479661) (K.S.C.); the SNUH Research Fund (No. 04-2024-0600; No. 04-2025-2060) (K.S.C.); and the Korea Health Technology R&D Project through the Korea Health Industry Development Institute (KHIDI) grant funded by the Ministry of Health&Welfare (No. RS-2024-00439549) (K.S.C.).

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

References

1. Aichert, A., Manhart, M., Navalpakkam, B.K., Grimm, R., Hutter, J., Maier, A., Hornegger, J., Doerfler, A.: A realistic digital phantom for perfusion c-arm CT based on MRI data. In: Nuclear Science Symposium and Medical Imaging Conference. pp. 1–2. IEEE (2013). <https://doi.org/10.1109/NSSMIC.2013.6829168>
2. Albers, G.W., Marks, M.P., Kemp, S., Christensen, S., Tsai, J.P., Ortega-Gutierrez, S., McTaggart, R.A., Torbey, M.T., Kim-Tenser, M., Leslie-Mazwi, T., et al.: Thrombectomy for stroke at 6 to 16 hours with selection by perfusion imaging. *N Engl J Med* **378**(8), 708–718 (2018). <https://doi.org/10.1056/NEJMoa1713973>
3. Alves, J., Carneiro, A., Xavier, J.: Reliability of CT perfusion in the evaluation of the ischaemic penumbra. *The Neuroradiology Journal* **27**(1), 91–95 (2014). <https://doi.org/10.15274/NRJ-2014-10010>
4. Amini, A., Schwarting, W., Soleimany, A.P., Rus, D.: Deep evidential regression. In: Advances in Neural Information Processing Systems. vol. 33 (2020), <https://arxiv.org/abs/1910.02600>
5. Bennink, E., Oosterbroek, J., Kudo, K., Viergever, M., Velthuis, B., de Jong, H.W.A.M.: Fast nonlinear regression method for CT brain perfusion analysis. *Journal of Medical Imaging* **3**(2), 026003–026003 (2016). <https://doi.org/10.1117/1.JMI.3.2.026003>
6. de Vries, L., van Herten, R.L.M., Hoving, J.W., Išgum, I., Emmer, B., Majioie, C., Marquering, H., Gavves, E.: Spatio-temporal physics-informed learning: A novel approach to CT perfusion analysis in acute ischemic stroke. *Medical Image Analysis* **90**, 102971 (2023). <https://doi.org/10.1016/j.media.2023.102971>, <https://doi.org/10.1016/j.media.2023.102971>

7. de Vries, L., van Herten, R.L.M., Hoving, J.W., Išgum, I., Emmer, B., Majoie, C.B., Marquering, H., Gavves, S.: Accelerating physics-informed neural fields for fast CT perfusion analysis in acute ischemic stroke. In: Medical Imaging with Deep Learning. Proceedings of Machine Learning Research, vol. 250, pp. 1606–1626. PMLR (2024)
8. Fieselmann, A., Kowarschik, M., Ganguly, A., Hornegger, J., Fahrigr, R.: Deconvolution-based CT and MR brain perfusion measurement: Theoretical model revisited and practical implementation details. *International Journal of Biomedical Imaging* **2011**(1), 1–20 (2011). <https://doi.org/10.1155/2011/467563>
9. Goyal, M., Ospel, J., Menon, B., Almekhlafi, M., Jayaraman, M., Fiehler, J., Psychogios, M., Chapot, R., van der Lugt, A., Liu, J., et al.: Challenging the ischemic core concept in acute ischemic stroke imaging. *Stroke* **51**(10), 3147–3155 (2020). <https://doi.org/10.1161/STROKEAHA.120.030620>
10. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: Advances in Neural Information Processing Systems. vol. 30 (2017), <https://arxiv.org/abs/1612.01474>
11. Livne, M., Rieger, J., Aydin, O., Taha, A., Akay, E.M., Kossen, T., Sobesky, J., Kelleher, J.D., Hildebrand, K., Frey, D., et al.: A U-Net deep learning framework for high performance vessel segmentation in patients with cerebrovascular disease. *Frontiers in Neuroscience* **13**, 97 (2019). <https://doi.org/10.3389/fnins.2019.00097>
12. Maier, O., Menze, B.H., von der Gablentz, J., Häni, L., Heinrich, M., Liebrand, M., Winzeck, S., Basit, A., Bentley, P., Chen, L., et al.: ISLES 2015 - a public evaluation benchmark for ischemic stroke lesion segmentation from multispectral MRI. *Medical Image Anal.* **35**, 250–269 (2017). <https://doi.org/10.1016/j.media.2016.07.009>
13. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics* **41**(4), 102:1–102:15 (2022). <https://doi.org/10.1145/3528223.3530127>
14. Müller, T.: tiny-cuda-nn (4 2021), <https://github.com/NVlabs/tiny-cuda-nn>
15. Murphy, B.D., Chen, X., Lee, T.: Serial changes in CT cerebral blood volume and flow after 4 hours of middle cerebral occlusion in an animal model of embolic cerebral ischemia. *AJNR Am J Neuroradiol* **28**(4), 743–749 (2007), <https://www.ajnr.org/content/28/4/743>
16. Østergaard, L., Weisskoff, R., Chesler, D., Gyldensted, C., Rosen, B.: High resolution measurement of cerebral blood flow using intravascular tracer bolus passages. part i: Mathematical approach and statistical analysis. *Magnetic Resonance in Medicine* **36**(5), 715–725 (1996). <https://doi.org/10.1002/MRM.1910360510>
17. Raissi, M., Perdikaris, P., Karniadakis, G.: 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). <https://doi.org/10.1016/J.JCP.2018.10.045>
18. Riordan, A., Prokop, M., Viergever, M., Dankbaar, J., Smit, E., de Jong, H.W.A.M.: Validation of CT brain perfusion methods using a realistic dynamic head phantom. *Medical Physics* **38**(6Part1), 3212–3221 (2011). <https://doi.org/10.1118/1.3592639>
19. Saver, J.L.: Time is brain—quantified. *Stroke* **37**(1), 263–266 (2006). <https://doi.org/10.1161/01.str.0000196957.55928.ab>
20. Sensoy, M., Kandemir, M., Kaplan, L.M.: Evidential deep learning to quantify classification uncertainty. In: Advances in Neural Information Processing Systems. vol. 31 (2018), <https://arxiv.org/abs/1806.01768>

21. Sitzmann, V., Martel, J.N.P., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 7462–7473 (2020), <https://arxiv.org/abs/2006.09661>
22. Taha, A., Hanbury, A.: Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. *BMC Medical Imaging* **15**(1), 29 (2015). <https://doi.org/10.1186/s12880-015-0068-x>
23. Wu, O., Østergaard, L., Weisskoff, R., Benner, T., Rosen, B., Sorensen, A.: Tracer arrival timing-insensitive technique for estimating flow in MR perfusion-weighted imaging using singular value decomposition with a block-circulant deconvolution matrix. *Magnetic Resonance in Medicine* **50**(1), 164–174 (2003). <https://doi.org/10.1002/mrm.10522>
24. Yang, L., Meng, X., Karniadakis, G.: B-PINNs: Bayesian physics-informed neural networks for forward and inverse PDE problems with noisy data. *Journal of Computational Physics* **425**, 109913 (2021). <https://doi.org/10.1016/j.jcp.2020.109913>