

Neural fields for tissue attenuation curve reconstruction in sparsely sampled time-resolved CT

Lucas de Vries^{1,2,3}

Rudolf L.M. van Herten^{2,3}

P. Matthijs van der Sluijs⁴

Ivana Išgum^{1,2,3}

Bart J. Emmer¹

Charles B.L.M. Majoie¹

Henk A. Marquering^{*1,2}

Efstratios Gavves^{*3}

LUCAS.DEVRIES@AMSTERDAMUMC.NL

R.L.M.VANHERTEN@AMSTERDAMUMC.NL

P.VANDERSLUIJS@ERASMUSMC.NL

I.ISGUM@AMSTERDAMUMC.NL

B.J.EMMER@AMSTERDAMUMC.NL

C.B.MAJOIE@AMSTERDAMUMC.NL

H.A.MARQUERING@AMSTERDAMUMC.NL

E.GAVVES@UVA.NL

¹ *Amsterdam UMC, location University of Amsterdam, Radiology and Nuclear Medicine, Meibergdreef 9, Amsterdam, 1105 AZ, The Netherlands*

² *Amsterdam UMC location University of Amsterdam, Biomedical Engineering and Physics, Meibergdreef 9, Amsterdam, 1105 AZ, The Netherlands*

³ *Informatics Institute, University of Amsterdam, Amsterdam, The Netherlands*

⁴ *Department of Radiology and Nuclear Medicine, Erasmus MC, University Medical Center Rotterdam, Rotterdam, The Netherlands*

Editors: Under Review for MIDL 2025

Abstract

Time-resolved CT imaging can aid acute ischemic stroke diagnosis by visualizing contrast agent transport through the brain (micro)vasculature. CT perfusion imaging, while widely used for stroke diagnosis, requires approximately 30 sequential scans, leading to extensive radiation exposure and motion sensitivity. As an alternative to CTP perfusion imaging, some hospitals opt for multiphase CT angiography for time-resolved analysis with reduced radiation dose. However, multiphase CT angiography lacks standardized perfusion analysis capabilities, making it more challenging to interpret than CT perfusion imaging. We present Sparse Temporal Attenuation Reconstruction (STAR), a novel approach using conditional neural fields that reconstructs tissue attenuation curves from sparse observations, allowing for reduced radiation exposure and motion sensitivity with CT perfusion, while enabling perfusion analysis from multiphase CT angiography. Our method generates full tissue attenuation curves using only 4 out of 30 observations. The results show that perfusion maps from reconstructed data match the reference perfusion maps, potentially reducing radiation and allowing recovery of motion-corrupted images. Moreover, STAR enables perfusion analysis in centers using multiphase CT angiography. Consequently, STAR has the potential to improve the stroke imaging work-up while making perfusion analysis more widely accessible.

Keywords: conditional neural fields, CT perfusion, multiphase CT angiography, acute ischemic stroke

* Contributed equally

1. Introduction

Assessment of hemodynamic consequences of stroke is commonly performed by time-resolved CT imaging such as CT perfusion or multiphase CT angiography, where the transport of the contrast agent through the intracranial (micro) vasculature is assessed. CT perfusion is widely regarded as the standard method for the hemodynamic analysis in most hospitals. This technique involves acquiring approximately 30 CT scans at 1–3 second intervals after a contrast agent is administered. The resulting images are subsequently analyzed by software to generate *perfusion maps*, which summarize tissue hemodynamics such as blood flow and transit times. However, CT perfusion faces several limitations in clinical practice. Firstly, each acquisition exposes the patient to radiation that builds up to a significant radiation dose with 30 acquisitions. Secondly, patient movement can cause severe motion artifacts in acquisitions (Fahmi et al., 2013). If a significant number of images is affected by motion artifacts, the CT perfusion analysis fails. Previously, multiphase CT angiography has been introduced to deal with extensive radiation exposure (Menon et al., 2015). Multiphase CT angiography extends CT angiography with two additional acquisitions, offering temporal information with minimal protocol changes and little additional costs and scan time (Menon et al., 2015; Dundamadappa et al., 2021). However, multiphase CT angiography does not allow for a straightforward generation of perfusion maps, which have superior sensitivity for detecting perfusion defects, particularly for distal occlusions (Benali et al., 2023).

We propose an approach that tackles CT perfusion’s main drawbacks: radiation dose and motion sensitivity. Furthermore, we show that our approach allows for performing perfusion analysis utilizing multiphase CT angiography data, all while maintaining compatibility with commercial off-the-shelf perfusion analysis software. The core of our method lies in reconstructing the complete temporal evolution of contrast enhancement, also known as tissue attenuation curves, using only a limited number of measurements. Previous research focused on the interpolation of uniformly undersampled CT perfusion data to reduce radiation (Bae et al., 2024). Others used multiphase CT angiography acquisitions to estimate perfusion parameter maps through methods such as calculating the per-voxel slope through three attenuation measurements (McDougall et al., 2020) or diffusion models (Cai et al., 2024). However, these methods are unable to extrapolate beyond acquisitions, restricted to interpolation between temporally close samples (Bae et al., 2024), and focused on generating software-specific perfusion maps (Son et al., 2024; Cai et al., 2024) rather than reconstructing tissue attenuation curves, which can be used with any analysis software.

We propose STAR: Sparse Temporal Attenuation Reconstruction, which learns continuous voxel-wise temporal representations of attenuation curves using conditional neural fields to enable attenuation curve reconstruction from sparse measurements. Consequently, STAR can estimate the tissue attenuation curve from only 4 out of 30 observations (13%). These four acquisitions correspond to time points that align with non-contrast CT and multiphase CT angiography acquisition protocols. The generated attenuation data can subsequently be analyzed with off-the-shelf CT perfusion analysis software. We show that the obtained perfusion maps are on par with the perfusion maps obtained from the fully sampled CT perfusion source data. STAR allows for perfusion analysis of attenuation data with 87% fewer CT acquisitions, which could lead to a significant radiation reduction. Additionally, this method could be used to reconstruct parts of the sequence when a significant number

of acquisitions are corrupted by motion artifacts. Furthermore, we demonstrate that STAR can reconstruct sequential attenuation data using non-contrast CT and multiphase CT angiography measurements, enabling perfusion analysis based on multiphase CT angiography.

2. Method

In the following, we introduce STAR. Figure 1 presents a visual overview of the method.

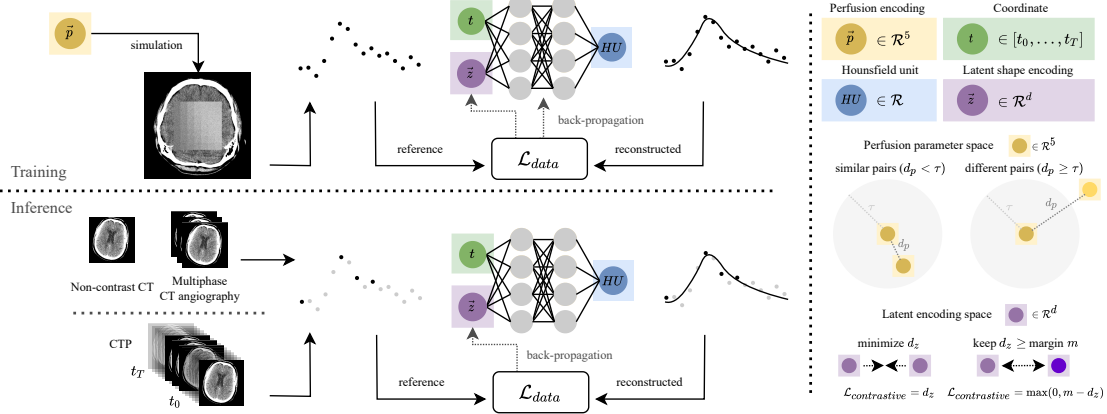


Figure 1: Overview of STAR. During training, a conditional neural field is trained on fully sampled simulated CT perfusion data. At inference, only the latent variable is updated based on the observed data on the subsampled domain.

2.1. Conditional neural fields

In physics, a field f is defined as a scalar or vector quantity over a particular domain. For example, we can view the dynamic contrast attenuation as a field defined on the temporal domain. A *neural field* is a field represented by a neural network f_θ with parameters θ (Xie et al., 2022). In our method, we use the simplest one-dimensional and scalar form of a neural field $f_\theta(t) : \mathbb{R} \rightarrow \mathbb{R}$, because we want to learn the neural field of the scalar attenuation (given in Hounsfield Units) only on the temporal domain $t \in [t_0, \dots, t_T]$. The goal is to infer the complete tissue attenuation curve using only limited measurements. Since attenuation curves often share similar patterns, training a separate neural field for each curve fails to leverage these shared characteristics. Moreover, training individual neural fields is computationally inefficient, as each curve requires its own complete training process. Therefore, we employ *conditional* neural fields.

In a conditional neural field, we enhance the neural field by adding a d -dimensional latent variable $z \in \mathbb{R}^d$ as input, allowing us to represent many different curves with a single field $f_\theta(t, z) : \mathbb{R}^{1+d} \rightarrow \mathbb{R}$. This latent variable z is essentially an embedding of the curve’s shape or a *shape conditional*. In practice, a latent z exists for each instance in the data set.

During training, we optimize both the neural field parameters θ and the latent variables z to minimize the difference between our predictions and the training data. At inference,

we keep the neural field parameters θ fixed and iteratively adjust the latent variable \mathbf{z} to minimize the difference between the neural field predictions and the available data points for a particular number of iterations. The attenuation data are not required to be densely sampled on the temporal domain: we may employ a sub-sampled set. We thus optimize the latents \mathbf{z} with data points from the sampled domain. Once we have the optimized latent variable, we can reconstruct the complete tissue attenuation curve by sampling $f_\theta(t, \mathbf{z})$ at all $t \in [t_0, \dots, t_T]$, allowing us to fill the gaps in the temporal data.

Training All latent variables $\mathbf{z} \in \mathbb{R}^d$ are randomly initialized from $\mathcal{N}(0, \frac{1}{\sqrt{d}})$, where $d = 32$ (Park et al., 2019). The network $f_\theta(t, \mathbf{z})$ consists of a single hidden layer of dimension 128 with tanh activations and outputs the predicted value through a final sigmoid activation function. We use a compound loss function consisting of three losses. The data reconstruction loss $\mathcal{L}_{\text{data}}$ is a ℓ_2 -loss based on the difference between our predictions and the training data. We use a regularization loss $\mathcal{L}_{\text{reg}} = \sum_i \|z_i\|_2$ to prevent the latents from growing arbitrarily large. Finally, we use a contrastive loss $\mathcal{L}_{\text{contrastive}}$ to assure that latents corresponding to similar attenuation curves are close together in the latent space. For any two curve samples i and j , we measure the similarity of the perfusion parameters that were used to simulate these curves (which we denote as $d_{\mathbf{p}}$) and look at how far apart they are in latent space (denoted as $d_{\mathbf{z}}$). For similar pairs (when $d_{\mathbf{p}} < \tau$), we want their latents to be close, so we directly penalize their distance $d_{\mathbf{z}}$. For different pairs (when $d_{\mathbf{p}} \geq \tau$), we want their latents to be at least margin m apart. If they are too close, we penalize with $\max(0, m - d_{\mathbf{z}})$.

$$\mathcal{L}_{\text{contrastive}} = \frac{1}{|P|} \sum_{i,j} \begin{cases} d_{\mathbf{z}}(i,j) & \text{if } d_{\mathbf{p}}(i,j) < \tau \text{ (similar pairs)} \\ \max(0, m - d_{\mathbf{z}}(i,j)) & \text{if } d_{\mathbf{p}}(i,j) \geq \tau \text{ (different pairs)} \end{cases} \quad (1)$$

To make all comparisons fair, \mathbf{p} and \mathbf{z} are normalized. We average over all possible pairs (the total number of pairs is denoted as $|P|$). The threshold $\tau = 0.1$ decides what we consider similar, while $m = 1.0$ defines how far apart different latents should be. Figure 1 visualizes the contrastive loss. The total training loss, with weights empirically set, is:

$$\mathcal{L} = 100 \cdot \mathcal{L}_{\text{data}} + 0.1 \cdot \mathcal{L}_{\text{reg}} + 10 \cdot \mathcal{L}_{\text{contrastive}} \quad (2)$$

We optimize both the model parameters and the latents using Adam and employ a cosine learning rate schedule that decays the initial learning rate 10^{-2} to 10^{-6} . We train for 16,000 iterations with batches of 1,000 densely sampled attenuation curves.

Inference All latent variables \mathbf{z} are randomly initialized from $\mathcal{N}(0, \frac{1}{\sqrt{d}})$. We optimize the latents \mathbf{z} with $\mathcal{L}_{\text{data}}$ and \mathcal{L}_{reg} using only the observed data with an ℓ_1 -loss, while keeping the neural field parameters θ fixed. We optimize for 1000 iterations using Adam (<1 minute on consumer hardware with a NVIDIA RTX 2080S) with learning rate 10^{-2} with a cosine learning rate schedule decaying the learning rate to 10^{-6} . To obtain the attenuation values at all time points we infer the network $f_\theta(t, \mathbf{z})$ at all $t \in [t_0, \dots, t_T]$.

2.2. Datasets

We use simulated phantom data to train, and validate with both phantom and patient data. We refer to Appendix A for more details regarding the data and preprocessing.

Phantom CT perfusion data In perfusion analysis, cerebral blood flow (CBF), cerebral blood volume (CBV), mean transit time (MTT), time-to-maximum (TMAX), and the DELAY are the main parameters of interest. The phantom developed by Kudo et al. (2012) included CBV values of 1–5 ml/100g. Here, we expand the phantom to cover a wider range of 1–20 ml/100g such that our method learns to operate on attenuation curves corresponding to non-tissue voxels like vessels. In total, we trained with 735 perfusion parameter combinations.

Patient CT perfusion data We evaluate with CT perfusion data sets from 17 patients from the Ischemic Stroke Lesion Segmentation Challenge (ISLES) 2024 (de la Rosa et al., 2024). We used one patient for validation and the other 16 patients for testing.

Patient multiphase CT angiography data We curated a dataset from Erasmus MC of five patients with matched non-contrast CT, multiphase CT angiography, and CT perfusion.

2.3. Baseline: gamma variate model

We fit a voxel-by-voxel gamma variate model to the attenuation curves (Thompson et al., 1964) using constrained curve-fitting. Initialization and bounds are based on the expected contrast dynamics (details provided in Appendix B). The model parameters represent physical properties: peak enhancement, arrival time, rise time, and washout rate. Given the inherent noise in the data, we use this parametric model over more flexible approaches.

3. Experiments

Phantom study We compare the perfusion parameters obtained from the complete CT perfusion scan with different subsampling scenarios. The first two scenarios correspond to uniform temporal subsampling at 15/30 ($t \in [t_0, t_2, \dots]$) and 7/30 ($t \in [t_0, t_4, \dots]$). The third scenario corresponds to time points aligning with multiphase CT angiography acquisition times. These include the first frame $t = t_0$ and frames corresponding to the peak arterial phase (the time point where the arterial input function peaks) and two delayed phases (peak venous and late venous): $t \in [t_{pa}, t_{pv}, t_{lv}]$. The selected time points are at: $t_0 = 0$ seconds (the first frame of the sequence), $t_{pa} = 16$ seconds, $t_{pv} = 24$ seconds, and $t_{lv} = 32$ seconds, using similar intervals between phases as with multiphase CT angiography acquisition protocols (Menon et al., 2015). Lastly, as a fourth scenario, we use $t \in [t_0, t_{pa}]$. We use commercially available CT perfusion analysis software SYNGO.VIA (version VB60S; Siemens Healthcare, Erlangen, Germany) to obtain the perfusion parameters, and qualitatively and quantitatively compare the perfusion parameters obtained from STAR-reconstructions with the fully sampled CT perfusion data. Moreover, we quantitatively analyze the gamma variate model as a baseline. We calculate the mean absolute error as an evaluation metric.

Patient study We extract sparse measurements $t \in [t_0, t_{pa}, t_{pv}, t_{lv}]$ from fully sampled CT perfusion data. Since, at inference time, the arterial peak is not available, we use the same time points that we extracted from the phantom data. With two commercially available CT perfusion analysis software solutions: SYNGO.VIA and STROKEVIEWER (version 3.2.11; Nicolab, Amsterdam, The Netherlands), we qualitatively compare perfusion maps generated from fully sampled and STAR-derived CT perfusion data from these sparse measurements and assess the volumetric agreement between penumbra and ischemic core volumes.

Table 1: Mean absolute error in perfusion parameter estimation from phantom data.

Scenario	Total time pts.	CBF [ml/100g/s]	CBV [ml/100g]	MTT [s]	TMAX [s]	DELAY [s]
$t \in [t_0, t_2, \dots]$	15/30	1.6	0.2	2.1	0.9	0.6
$t \in [t_0, t_4, \dots]$	7/30	6.4	0.3	3.7	1.6	0.8
$t \in [t_0, t_{pa}, t_{pv}, t_{lv}]$ – proposed	4/30	4.5	0.3	3.8	1.6	0.8
$t \in [t_0, t_{pa}, t_{pv}, t_{lv}]$ – curve fit	4/30	10.3	0.4	6.6	4.4	3.8
$t \in [t_0, t_{pa}]$	2/30	12.7	2.6	3.2	1.8	0.8

Multiphase CT angiography study As a proof-of-concept, we qualitatively assess agreement between SYNGO.VIA and STROKEVIEWER perfusion maps from STAR-derived data reconstructed from multiphase CT angiography with those from CT perfusion.

Additional studies With UMAP (McInnes et al., 2018), a dimensionality reduction technique, we visualize the latent space to assess if physiologically similar curves are encoded proximally in the latent space and if the contrastive loss stimulates this space to be well-structured. Furthermore, as ablation study, we investigate loss function combinations and $t_{pa} \leftrightarrow t_{pv} \leftrightarrow t_{lv}$ intervals on the phantom data perfusion parameter accuracy.

4. Results

Phantom study All perfusion parameters derived from STAR reconstructions show strong agreement with those from complete CT perfusion data. Table 1 lists the mean absolute error. As we subsample more extensively, we see that the errors in the temporal perfusion parameters increase. For the scenario with four time points, the gamma variate model performs poorly, with particularly large errors in temporal parameters. Figure 4 in Appendix C confirms the error increase with more sub-sampling by comparing the CBF and TMAX maps.

Patient study Figure 2 presents SYNGO.VIA and STROKEVIEWER perfusion maps. We observe strong visual correspondence between the perfusion maps from the STAR-derived and fully sampled CT perfusion data. The main deviations are in the temporal perfusion parameter TMAX. Table 2 in Appendix D lists the volumetric agreement between infarct core and penumbra estimates. When comparing STAR-derived versus fully sampled CT perfusion data, the median (IQR) volumetric differences are -5.5 (-18.1 – 0.1) ml (SYNGO.VIA) and 0.0 (-3.5 – 4.5) ml (STROKEVIEWER) for the infarct core. For reference, the inter-software differences (STROKEVIEWER–SYNGO.VIA) for infarct core volume in the case of fully sampled CT perfusion data are -10.5 (-10.2 – (-6.5)) ml.

Multiphase CT angiography study Figure 2 demonstrates the correspondence between perfusion maps from STAR-derived attenuation data from clinical multiphase CT angiography data and CT perfusion data. We refer to Appendix E for more examples.

Additional studies Figure 3 visualizes the latent space under different loss functions. The figure shows that contrastive loss groups similar perfusion patterns closer together. We

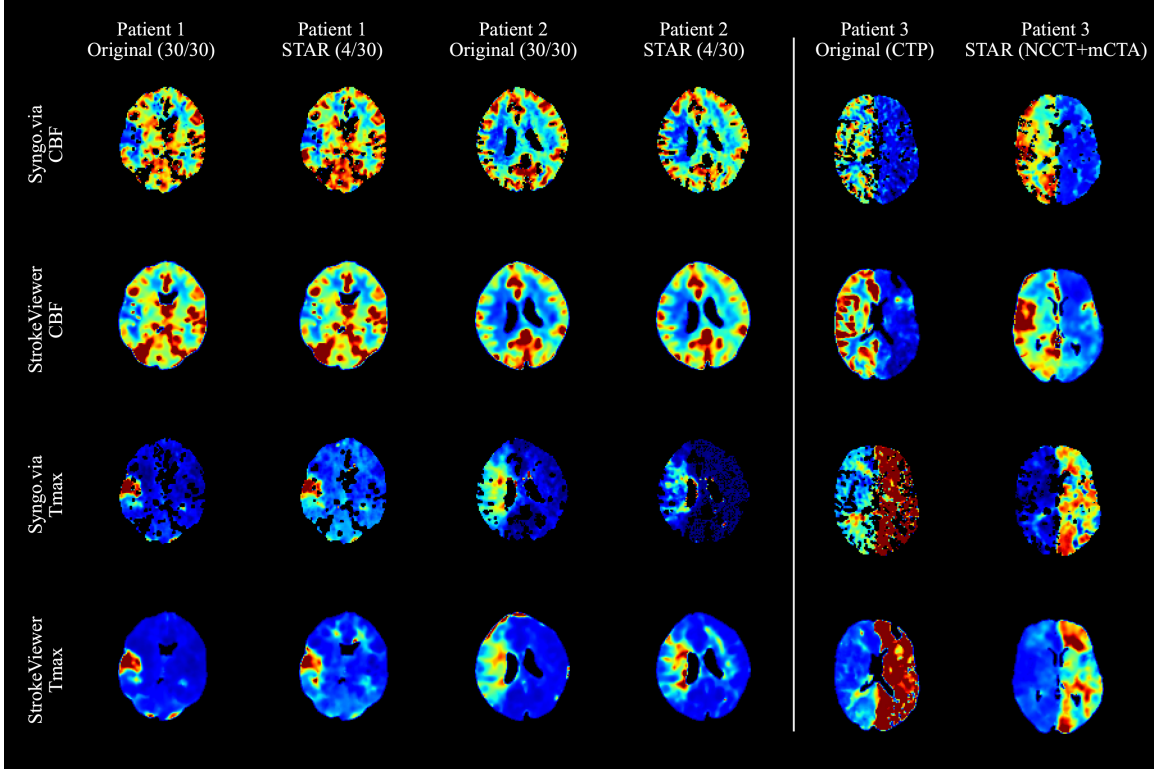


Figure 2: Perfusion maps from both SYNGO.VIA and STROKEVIEWER comparing: (left) fully sampled CT perfusion versus four timepoints and (right) CT perfusion versus non-contrast CT with multiphase CT angiography.

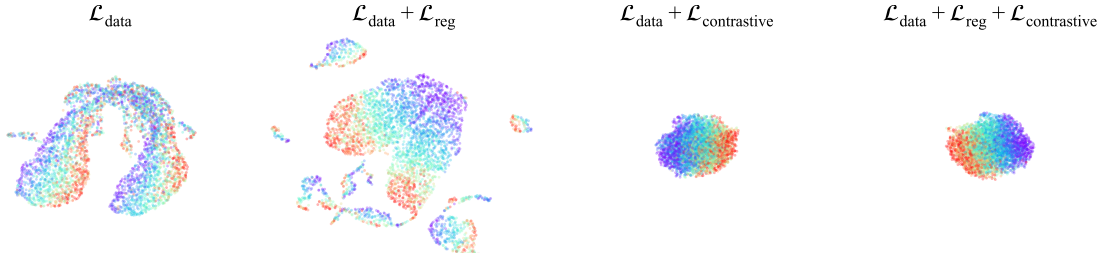


Figure 3: Visualization of the latent space under different loss function combinations. The color of each dot represents its CBV value.

refer to Appendix F for more details. The ablation studies in Appendix G show the impact of different loss function combinations and different $t_{pa} \leftrightarrow t_{pv} \leftrightarrow t_{lv}$ intervals on perfusion parameter accuracy on the phantom data, demonstrating that contrastive learning improves perfusion parameter accuracy and that an interval of 6 or 8 seconds is appropriate.

5. Discussion and Conclusion

STAR enables a CT perfusion subsampling approach that aligns with the *As Low As Reasonably Achievable (ALARA)* principle. STAR allows a significant reduction in radiation exposure while maintaining diagnostic quality. The high undersampling rate suggests that the method can also reconstruct CT perfusion data with severe motion corruption, allowing for perfusion analysis even in such cases. Moreover, STAR shows promise in reconstructing full-sequence attenuation data from multiphase CT angiography. Since STAR completes CT attenuation curves rather than directly generating perfusion maps, clinicians can use off-the-shelf CT perfusion analysis software without modifying existing workflows.

We note three observations that need further discussion. First, we observed expected deviations in the temporal perfusion maps, as sparse sampling naturally misses exact bolus arrival time and wash-out (Table 1). The deviations increase with more aggressive subsampling protocols. Second, STAR’s training on tissue contrast attenuation curves creates a clear constraint: the model struggles to reconstruct higher attenuation values typical of arterial and venous structures, despite our expanded phantom’s cerebral blood volume range. Consequently, this limitation affects the automated selection of arterial input and venous outflow locations by the CT perfusion software, impacting perfusion estimates. Using population-based arterial input and venous outflow attenuation curves could offer a solution. Finally, replacing CT perfusion attenuation data with the combination of non-contrast CT and multiphase CT angiography data results in increased perfusion parameter deviations. While we can detect infarcts qualitatively, the perfusion maps lack the precision needed for accurate core-penumbra volume calculations. Differences in acquisition protocols, tube current, and kilovolt-peak settings likely impact attenuation values. These findings highlight the need to further explore alignment between non-contrast CT and multiphase-CT angiography to match CT perfusion source data.

STAR only considers the temporal domain for reconstruction. Future work could incorporate spatial context through neural fields that operate in both spatial and temporal domains (Dupont et al., 2022; Bauer et al., 2023). However, for more complex spatio-temporal signals, conditional neural fields with global conditioning scale poorly (Dupont et al., 2022; Bauer et al., 2023; Papa et al., 2024; Xie et al., 2022). Recent work on equivariant neural fields could provide a solution with a geometry-informed latent space (Wessels et al., 2024).

The contrastive loss provides regularization of the latent space, ensuring reconstructed curves maintain physiologically plausible shapes even from very sparse measurements – a clear advantage over the gamma-variate model, which lacks flexibility to capture the full range of attenuation patterns. Future work could explore other approaches like physics-informed losses to constrain solutions toward physiologically valid curves, particularly when working with extremely limited temporal samples (De Vries et al., 2023).

In conclusion, we demonstrated that reconstructing full perfusion attenuation curves from as few as four measurements is possible, potentially allowing both substantial radiation dose reduction and correction of motion-corrupted acquisitions. Our method enables perfusion measurements from standard multiphase CT angiography acquisition protocols while maintaining compatibility with existing clinical software and workflows.

Acknowledgments

This work was part of the Artificial Intelligence for Early Imaging-Based Patient Selection in Acute Ischemic Stroke (AIRBORNE) project. This project was supported by Top Sector Life Sciences & Health and Nicolab B.V. We gratefully acknowledge Erasmus MC for providing access to the imaging data used in this study.

References

- Juyoung Bae, Elizabeth Tong, and Hao Chen. Conditional Diffusion Model for Versatile Temporal Inpainting in 4D Cerebral CT Perfusion Imaging . In *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, volume LNCS 15002. Springer Nature Switzerland, 2024.
- Matthias Bauer, Emilien Dupont, Andy Brock, Dan Rosenbaum, Jonathan Richard Schwarz, and Hyunjik Kim. Spatial Functa: Scaling Functa to ImageNet Classification and Generation. *arXiv e-prints*, page arXiv:2302.03130, 2023. doi: 10.48550/arXiv.2302.03130.
- Faysal Benali, Jianhai Zhang, Najratun Nayem Pinky, Fouzi Bala, Ibrahim Alhabli, Rotem Golan, Luis A Souto Maior Neto, Ibukun Elebute, Chris Duszynski, Wu Qiu, and Bijoy K Menon. Validation of a Novel Multiphase CTA Perfusion Tool Compared to CTP in Patients With Suspected Acute Ischemic Stroke. *Stroke: Vascular and Interventional Neurology*, 3(3):e000811, 5 2023. doi: 10.1161/SVIN.122.000811. URL <https://doi.org/10.1161/SVIN.122.000811>.
- Yuxin Cai, Jianhai Zhang, Lei He, Aravind Ganesh, and Wu Qiu. Masked Residual Diffusion Probabilistic Model with Regional Asymmetry Prior for Generating Perfusion Maps from Multi-phase CTA . In *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, volume LNCS 15002. Springer Nature Switzerland, 2024.
- Ezequiel de la Rosa, Ruisheng Su, Mauricio Reyes, Roland Wiest, Evamaria O Riedel, Florian Kofler, Kaiyuan Yang, Hakim Baazaoui, David Robben, Susanne Wegener, Jan S Kirschke, Benedikt Wiestler, and Bjoern Menze. ISLES’24: Improving final infarct prediction in ischemic stroke using multimodal imaging and clinical data, 2024. URL <https://arxiv.org/abs/2408.10966>.
- Lucas De Vries, Rudolf L M Van Herten, Jan W Hoving, Ivana Išgum, Bart J Emmer, Charles B L M Majoie, Henk A Marquering, and Efstratios Gavves. Spatio-temporal physics-informed learning: A novel approach to CT perfusion analysis in acute ischemic stroke. *Medical Image Analysis*, 90:102971, 2023. ISSN 1361-8415. doi: <https://doi.org/10.1016/j.media.2023.102971>. URL <https://www.sciencedirect.com/science/article/pii/S1361841523002311>.
- S Dundamadappa, K Iyer, A Agrawal, and D J Choi. Multiphase CT Angiography: A Useful Technique in Acute Stroke Imaging-Collaterals and Beyond. *AJNR. American*

- journal of neuroradiology*, 42(2):221–227, 1 2021. ISSN 1936-959X (Electronic). doi: 10.3174/ajnr.A6889.
- Emilien Dupont, Hyunjik Kim, S. M. Ali Eslami, Danilo Rezende, and Dan Rosenbaum. From data to functa: Your data point is a function and you can treat it like one. *arXiv e-prints*, page arXiv:2201.12204, 1 2022. doi: 10.48550/arXiv.2201.12204.
- F Fahmi, L F M Beenen, G J Streekstra, N Y Janssen, H W de Jong, A Riordan, Y B Roos, C B Majoie, E Vanbavel, and H A Marquering. Head movement during CT brain perfusion acquisition of patients with suspected acute ischemic stroke. *European journal of radiology*, 82(12):2334–2341, 12 2013. ISSN 1872-7727 (Electronic). doi: 10.1016/j.ejrad.2013.08.039.
- Kohsuke Kudo, Soren Christensen, Makoto Sasaki, Leif Østergaard, Hiroki Shirato, Kuniaki Ogasawara, Max Wintermark, and Steven Warach. Accuracy and reliability assessment of cT and Mr Perfusion analysis software Using a Digital Phantom 1. *Radiology*, 267: 201, 2012. doi: 10.1148/radiol.12112618/-/DC1. URL <http://radiology.rsna.org/lookup/suppl/doi:10.1148/radiol.12112618/-/DC1>.
- C C McDougall, L Chan, S Sachan, J Guo, R G Sah, B K Menon, A M Demchuk, M D Hill, N D Forkert, C D d’Esterre, and P A Barber. Dynamic CTA-Derived Perfusion Maps Predict Final Infarct Volume: The Simple Perfusion Reconstruction Algorithm. *AJNR. American journal of neuroradiology*, 41(11):2034–2040, 11 2020. ISSN 1936-959X (Electronic). doi: 10.3174/ajnr.A6783.
- Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger. UMAP: Uniform Manifold Approximation and Projection. *Journal of Open Source Software*, 3(29):861, 2018. doi: 10.21105/joss.00861. URL <https://doi.org/10.21105/joss.00861>.
- Bijoy K Menon, Christopher D d’Esterre, Emmad M Qazi, Mohammed Almekhlafi, Leszek Hahn, Andrew M Demchuk, and Mayank Goyal. Multiphase CT Angiography: A New Tool for the Imaging Triage of Patients with Acute Ischemic Stroke. *Radiology*, 275(2): 510–520, 2015. doi: 10.1148/radiol.15142256. URL <https://doi.org/10.1148/radiol.15142256>.
- Samuele Papa, Riccardo Valperga, David Knigge, Miltiadis Kofinas, Phillip Lippe, Jan-Jakob Sonke, and Efstratios Gavves. How to Train Neural Field Representations: A Comprehensive Study and Benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22616–22625, 2024.
- Jeong Joon Park, Peter R Florence, Julian Straub, Richard A Newcombe, and Steven Lovegrove. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In *{IEEE} Conference on Computer Vision and Pattern Recognition, {CVPR} 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 165–174. Computer Vision Foundation / {IEEE}, 2019. doi: 10.1109/CVPR.2019.00025. URL http://openaccess.thecvf.com/content%5C_CVPR%5C_2019/html/Park%5C_DeepSDF%5C_Learning%5C_Continuous%5C_Signed%5C_Distance%5C_Functions%5C_for%5C_Shape%5C_Representation%5C_CVPR%5C_2019%5C_paper.html.

- Moo Hyun Son, Juyoung Bae, Elizabeth Tong, and Hao Chen. Progressive Knowledge Distillation for Automatic Perfusion Parameter Maps Generation from Low Temporal Resolution CT Perfusion Images . In *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, volume LNCS 15010. Springer Nature Switzerland, 2024.
- Howard K Thompson, C Frank Starmer, Robert E Whalen, and Henry D McIntosh. Indicator Transit Time Considered as a Gamma Variate. *Circulation Research*, 14(6):502–515, 6 1964. doi: 10.1161/01.RES.14.6.502. URL <https://doi.org/10.1161/01.RES.14.6.502>.
- Kilian M Treurniet, Natalie E LeCouffe, Manon Kappelhof, Bart J Emmer, Adriaan C G M van Es, Jelis Boiten, Geert J Lycklama, Koos Keizer, Lonneke S F Yo, Hester F Lingsma, Wim H van Zwam, Inger de Ridder, Robert J van Oostenbrugge, Aad van der Lugt, Diederik W J Dippel, Jonathan M Coutinho, Yvo B W E M Roos, and Charles B L M Majoie. MR CLEAN-NO IV: intravenous treatment followed by endovascular treatment versus direct endovascular treatment for acute ischemic stroke caused by a proximal intracranial occlusion-study protocol for a randomized clinical trial. *Trials*, 22(1):141, 2 2021. ISSN 1745-6215 (Electronic). doi: 10.1186/s13063-021-05063-5.
- David R Wessels, David M Knigge, Samuele Papa, Riccardo Valperga, Sharvaree Vadgama, Efstratios Gavves, and Erik J Bekkers. Grounding Continuous Representations in Geometry: Equivariant Neural Fields. *arXiv e-prints*, page arXiv:2406.05753, 2024. doi: 10.48550/arXiv.2406.05753.
- Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural Fields in Visual Computing and Beyond. *Computer Graphics Forum*, 2022. ISSN 1467-8659. doi: 10.1111/cgf.14505.

Appendix A. Dataset description

Phantom CT perfusion data The phantom includes attenuation curves corresponding to cerebral blood volume: $CBV \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 12, 14, 16, 18, 20\}$ ml/100g. The mean transit time and delay perfusion values used for simulation are: $MTT \in \{3.4, 4.0, 4.8, 6.0, 8.0, 12.0, 24.0\}$ s and $DELAY \in \{0.0, 0.5, 1.0, 1.5, 2.0, 2.5, 3.0\}$ s. The corresponding cerebral blood flow (CBF) and time-to-maximum (TMAX) are: $CBF = CBV/MTT$ ml/100g/s and $TMAX = DELAY + \frac{1}{2}MTT$ s. The data also includes a simulated arterial input function and a simulated venous output function. All attenuation data is generated with a temporal resolution of two seconds and with a total acquisition time of 58 seconds (30 time points). The attenuation curves are generated through the convolution of an arterial input function with a box-shaped impulse response function. The phantom is organized in a series of axial slices, where each slice represents a distinct CBV level. Within each slice, the data is arranged in a 7×7 grid of tiles. In this grid, the MTT values increase from right to left, which consequently means that cerebral blood flow ($CBF = CBV/MTT$) increases from left to right. The delay values increase from top to bottom. Each tile contains 32×32 tissue

attenuation curves with different noise realizations. We use a 50/25/25 split for training, validation, and testing. After splitting the data we apply a Gaussian filter with a standard deviation of one voxel to each tile in each axial slice. We scale all attenuation values between 0 and 1. We subtract the attenuation value of the first time point as a baseline and restore the baseline after inferring the attenuation curve.

Patient CT perfusion data We resampled the CT perfusion data sets from 17 patients from the Ischemic Stroke Lesion Segmentation Challenge (ISLES) 2024 (de la Rosa et al., 2024) to a temporal resolution of 2 seconds and motion-corrected all CT perfusion images by registering all frames to the first frame of the sequence. Thereafter, we applied the bilateral smoothing filter ($\sigma_{domain} = 3$ mm, $\sigma_{range} = 10$ HU) to approximate the signal-to-noise of the phantom data.

Patient non-contrast and multiphase CT angiography data The proof-of-concept dataset from Erasmus MC consisted of five patients with matched non-contrast CT, multiphase CT angiography, and CT perfusion source data. The patients participated in the MR CLEAN-NO IV trial (Treurniet et al., 2021). The MR CLEAN-NO IV trial (MEC-2017-368) obtained ethics approval from Erasmus MC University Medical Center in Rotterdam and required written informed consent from all participants. The CT perfusion data were temporally resampled to 2-second intervals to maintain consistency with our previous experiments. We applied identical preprocessing steps as described above, including motion correction through registration to the first frame and bilateral filtering. All non-contrast CT and multiphase CT angiography volumes were spatially aligned to the first frame of the CT perfusion sequence.

Appendix B. Baseline: Gaussian variate curve fit

Following principles from the indicator dilution theory, we implemented a four-parameter gamma variate function to model the contrast agent dynamics:

$$f(t) = a \cdot (t - b)^c \cdot e^{-(t-b)/d} \cdot H(t - b) \quad (3)$$

where $H(t - b)$ is the Heaviside step function. The gamma variate’s parameters capture the key physiological aspects of contrast dynamics: amplitude a represents the peak enhancement reflecting maximum contrast concentration, time offset b indicates contrast arrival time in the tissue, shape parameter c characterizes the rise time corresponding to tissue perfusion rate, and scale parameter d describes the washout rate. Parameter bounds and initialization were defined based on the measured intensity values $C(t)$ at the expected domain $t \in [t_0, \dots, t_T]$:

$$\begin{aligned} 0 &\leq a \leq 5 \max(C(t)) \\ t_0 &\leq b \leq t_{\max(C(t))} \\ 0.1 &\leq c \leq 5.0 \\ t_T/10 &\leq d \leq t_T \end{aligned} \quad (4)$$

Initial parameter estimates were set as:

$$\begin{aligned} a_0 &= \max(C(t)) \\ b_0 &= t_0 \\ c_0 &= 1.0 \\ d_0 &= t_T/3 \end{aligned} \tag{5}$$

These constraints ensure physiologically plausible fits while providing sufficient flexibility to capture varying perfusion patterns. The initialization strategy proved robust across our dataset. We enforce that the attenuation returns to baseline levels after the contrast passes through.

Appendix C. Perfusion maps for phantom data

We focus on the CBF and TMAX perfusion parameters because these typically assess the acute infarct core and the salvageable tissue. We note that the error tends to increase for the temporal perfusion parameters as we progress to more severe subsampling.




Figure 4: Comparison of commercially available CT perfusion analysis software SYNGO.VIA perfusion estimates (CBF and TMAX) derived from complete CT perfusion data and four subsampling protocols.

Appendix D. Quantitative results for patient CT perfusion data

Table 2 lists the infarct core and penumbra volume estimates from STROKEVIEWER and SYNGO.VIA for the fully sampled CT perfusion data and the STAR-reconstructed data, and a comparison for these volumes between the two commercial software packages.

Appendix E. Perfusion maps for multiphase CT angiography-derived tissue attenuation data

Figure 5 present more qualitative results for perfusion maps derived from multiphase CT angiography data. The perfusion maps from STAR reconstructed curves appear smoother.

Table 2: Infarct core and penumbra volume (ml) estimates based on our method’s reconstructed and fully sampled data, analyzed with SYNGO.VIA and STROKEVIEWER. The table lists median (IQR) (absolute) volumetric difference (VD, AVD) over the test patients. Symbols indicate if smaller (\downarrow) or closer to zero (0) values denote better performance.

Data	Volume	SYNGO.VIA		STROKEVIEWER	
		VD (0)	AVD (\downarrow)	VD (0)	AVD (\downarrow)
CTP	Core	-5.5 (-18.1–0.1)	10.8 (3.1–22.7)	0.0 (-3.5–4.5)	5.0 (0.0–15.5)
	Penumbra	1.3 (-22.4–28.6)	28.4 (12.5–47.6)	-1.5 (-11.3–4.8)	10.5 (4–18.0)
STROKEVIEWER – SYNGO.VIA					
Data	Volume	VD (0)	AVD (\downarrow)		
CTP	Core	-10.5 (-10.2–(-6.5))	10.5 (6.5–20.2)		
	Penumbra	-88.5 (-107.9–(-62.0))	88.5 (62.0–107.9)		

While we can observe the lower CBF and elevated TMAX regions, the quality of the estimated perfusion maps is not the same as the reference perfusion maps. The elevated TMAX in the same region as the reference maps suggests the quality may be sufficient for detection. However, the accuracy is inadequate for core and penumbra estimates.

Appendix F. Latent visualization

Figure 3 visualizes the inferred latent space with each combination of training loss functions, for attenuation curves simulated with $MTT = 4$ seconds and $DELAY = 2$ seconds. With the contrastive loss enabled, we see that the latents corresponding to attenuation curves simulated with identical perfusion values are much closer in the latent space. The more compact latent space reduces the likelihood of converging to undesirable local optima.

Appendix G. Ablation studies

Table 3 lists the results for the ablation studies. We observe that the contrastive loss reduces the error on the perfusion parameters. Moreover, we see that a sampling interval of 8 seconds is a good middle ground in terms of performance, balancing the error in CBF and $DELAY$.

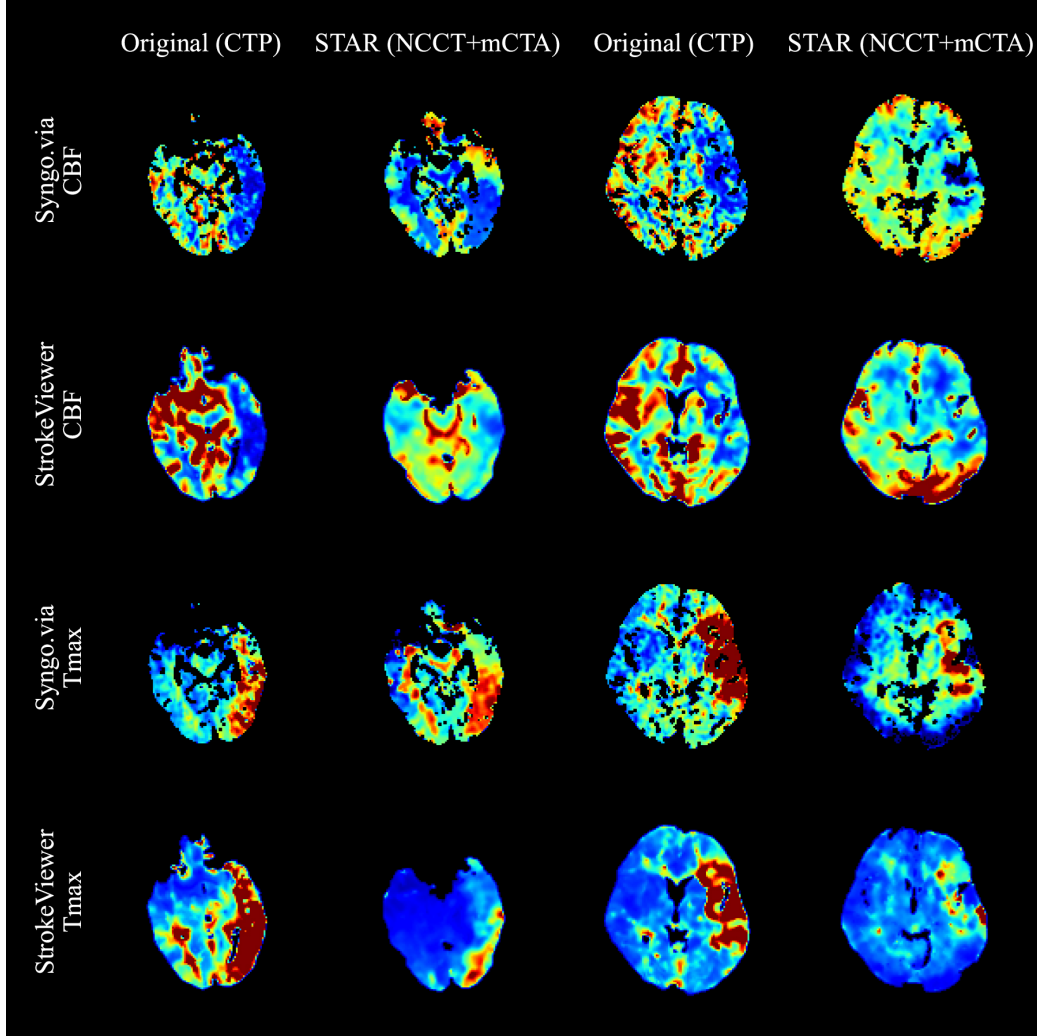


Figure 5: Qualitative results for perfusion maps derived from multiphase CT angiography data.

Table 3: Ablation studies results for different combinations of loss functions and various sampling intervals. Values closer to zero represent better performance.

\mathcal{L} 's used	CBF [ml/100g/s]	CBV [ml/100g]	MTT [s]	TMAX [s]	DELAY [s]
$\mathcal{L}_{\text{data}}$	7.4	1.4	5.4	4.2	3.5
$\mathcal{L}_{\text{data}} + \mathcal{L}_{\text{reg}}$	5.9	0.6	4.8	2.8	2.2
$\mathcal{L}_{\text{data}} + \mathcal{L}_{\text{contrastive}}$	3.3	0.9	4.0	1.3	0.7
$\mathcal{L}_{\text{data}} + \mathcal{L}_{\text{reg}} + \mathcal{L}_{\text{contrastive}}$	4.5	0.3	3.8	1.6	0.8
$t_{\text{pa}} \leftrightarrow t_{\text{pv}} \leftrightarrow t_{\text{lv}}$ interval	CBF [ml/100g/s]	CBV [ml/100g]	MTT [s]	TMAX [s]	DELAY [s]
6 sec.	2.9	0.3	3.7	1.7	1.2
8 sec.	4.5	0.3	3.8	1.6	0.8
10 sec.	6.0	0.4	4.2	1.7	0.7