

PyPeT: A Python Perfusion Tool for Automated Quantitative Brain CT and MR Perfusion Analysis

Marijn Borghouts¹  and Ruisheng Su¹ 

¹Department of Biomedical Engineering, Eindhoven University of Technology,
Eindhoven, The Netherlands

Correspondence: m.m.borghouts@tue.nl

Abstract. Computed tomography perfusion (CTP) and magnetic resonance perfusion (MRP) are widely used in acute ischemic stroke assessment and other cerebrovascular conditions to generate quantitative maps of cerebral hemodynamics. While commercial perfusion analysis software exists, it is often costly, closed source, and lacks customizability. This work introduces PyPeT, an openly available **Python Perfusion Tool** for head CTP and MRP processing. PyPeT is capable of producing cerebral blood flow (CBF), cerebral blood volume (CBV), mean transit time (MTT), time-to-peak (TTP), and time-to-maximum (Tmax) maps from raw four-dimensional perfusion data. PyPeT aims to make perfusion research as accessible and customizable as possible. This is achieved through a unified framework in which both CTP and MRP data can be processed, with a strong focus on modularity, low computational burden, and significant inline documentation. PyPeT’s outputs can be validated through an extensive debug mode in which every step of the process is visualized. Additional validation was performed via visual and quantitative comparison with reference perfusion maps generated by three FDA-approved commercial perfusion tools and a research tool. These comparisons show a mean SSIM around 0.8 for all comparisons, indicating a good and stable correlation with FDA-approved tools. The code for PyPeT is openly available at our GitHub <https://github.com/Marijn311/CT-and-MR-Perfusion-Tool>

Keywords: CT Perfusion · MR Perfusion · Acute Ischemic Stroke · Perfusion Analysis

1 Introduction

1.1 Background

Perfusion imaging is most commonly employed in the acute management of ischemic stroke, but it also has applications in other cerebrovascular conditions, including vasospasms, brain tumors, and head trauma. CTP and MRP are perfusion imaging techniques in which a contrast agent is injected into the cerebral

circulation, while the brain is continuously imaged using CT or MR as the contrast agent passes through. By isolating the contrast time curves (CTC) and deconvolving them with the arterial input function (AIF), one can obtain the residue function (R). From this residue function, one can extract the CBF, CBV, MTT, and Tmax to visualize the cerebral hemodynamics, see Figure 1 and Equations 1-6. Obtaining stable perfusion maps from raw perfusion images, however, require a significant amount of processing steps.




Fig. 1. Contrast time curves (CTC), residue function, and the derived perfusion parameters.

$$R(t) = \frac{1}{\text{CBF}} \cdot (\text{CTC}(t) \oslash \text{AIF}(t)) \quad (1)$$

$$\text{TTP} = \arg \max_t (\text{AIF}(t)) \quad (2)$$

$$\text{CBF} = \max_t R(t) \quad (3)$$

$$\text{CBV} = \int R(t) dt \quad (4)$$

$$\text{Tmax} = \arg \max_t R(t) \quad (5)$$

$$\text{MTT} = \frac{\text{CBV}}{\text{CBF}} \quad (6)$$

Variable definitions.

CTC(t) Concentration time curve, representing the measured contrast concentration in the tissue as a function of time.

AIF(t) Arterial input function, describing the CTC in an upstream artery.

R(t) Residue function, denoting the fraction of contrast agent remaining in the tissue at time t following an idealized impulse injection.

CBF Cerebral blood flow, defined as the maximum value of the residue function and expressing the volume of blood delivered per unit tissue per unit time.

CBV Cerebral blood volume, computed as the integral of the residue function and representing the volume of blood contained in a unit mass of tissue.

TTP Time to peak of the arterial input function, indicating the arrival time of contrast in the supplying artery.

Tmax Time to the maximum of the residue function, indicating the arrival time of contrast in the tissue.

MTT Mean transit time, defined as CBV/CBF and reflecting the average time blood spends within the tissue.

1.2 Problem Statement

Several FDA-approved commercial software solutions are available for perfusion image processing (e.g., Icobrain CVA, Viz CTP, iSchemaView RAPID, Canon MR-Clinical-Suite [3]). While these tools are effective, they are often expensive, closed-source, and limited in flexibility, making them less suitable for research. Consequently, there is a need for open alternatives that allow for easy customization and transparency.

1.3 Related Works

Some open approaches to perfusion image processing have been developed. For CTP, examples include the codebases described in [18,12], which both implement a singular value decomposition (SVD)-based deconvolution method. Another contribution is the work of [1], who proposed a box nonlinear regression technique to improve upon SVD-based deconvolution for generating perfusion maps. More recently, machine learning approaches have been explored for CTP processing. For instance, [19] employed physics-informed neural networks, while [6] applied a U-Net architecture to estimate perfusion maps directly from raw CTP data.

For MRP, open approaches include the works of [5,16,7] who developed SVD-based deconvolution implementations. Additionally, there is the work by [14] who adapted the work of [16]. There is also the project of [9] which aims to collect and validate open-source code fragments for MRP processing. More recently, there have been machine learning based approaches for MRP processing as well, such as the work by [13], which uses GAN to generate perfusion maps.

Although these open approaches are all valuable contributions to the field, and this list is surely not all-encompassing, each comes with its own limitations. Some represent only basic implementations and therefore lack features expected from a comprehensive perfusion analysis toolbox. Others provide limited or insufficient documentation, reducing their usability and customizability. Certain tools demand substantial computational resources, requiring high-performance CPUs or GPUs to function effectively. Some tools lack proper validation strategies. Moreover, all of these solutions are restricted to a single imaging modality. As a result, there remains a need for further progress in the development of open perfusion processing tools.

Therefore, PyPeT aims to contribute the first open and unified framework capable of processing both CTP and MRP data within a single tool. The software includes extensive inline documentation and is accompanied by this manuscript, which explains its design and functionality in detail. In addition, PyPeT offers a comprehensive debug mode that enables early detection of potential processing errors and includes validation against multiple FDA-approved perfusion analysis toolboxes.

1.4 Design Philosophy

The primary objective of this work is to make perfusion image processing as accessible as possible for researchers. To this end, we introduce PyPeT, an open **Python Perfusion Tool**. Python was chosen as the implementation language because of its widespread usage, clean syntax, expressiveness, and ecosystem of libraries [4]. Furthermore, Python is freely available, in contrast to commercial alternatives such as MATLAB. PyPeT employs an SVD-based deconvolution method to generate perfusion maps. Unlike deep learning approaches, SVD-based methods do not require training datasets or expensive hardware. The entire PyPeT pipeline can be executed on a mobility laptop in seconds. PyPeT’s codebase was deliberately structured in a modular fashion with extensive comments to support flexibility, customization, and future extensions by the research community. Finally, PyPeT introduces a unified framework for handling both CTP and MRP data. This dual compatibility broadens its applicability across different modalities and research contexts, making it a versatile tool for perfusion imaging studies. This report provides both a technical overview of the perfusion processing pipeline and practical guidance for researchers aiming to customize or validate its outputs.

2 Methodology

2.1 Perfusion Processing Steps

PyPeT’s implementation was loosely based upon the publicly available codebases of [18,16], which provide pipelines for head CT and MR perfusion processing, respectively. The main steps of the PyPeT pipeline are described stepwise in the following sections.

First, a raw four-dimensional perfusion image is loaded from a `.nii.gz` file and preprocessed. The volumes are reoriented to a standard RAS+ orientation. Next, the volumes get downsampled to the desired shape if they exceed a user-defined threshold, to speed up computation. After which all 3D volumes are registered to the first volume using ANTsPy’s QuickRigid method, to correct for patient motion. The volumes are then smoothed using a Gaussian filter. This smoothing step helps stabilize the results by mitigating noise in temporal measurements. Such noise can arise from small residual patient movements that may remain after motion correction.

Next, a brain mask is generated. In case of CTP images, a combination of intensity thresholding and a fast-marching algorithm is applied to grow the mask from the largest connected components, which should be the skull. For MRP images, an intensity thresholding and active contour approach is applied. Alternatively, users may supply a precomputed brain mask with a tool of their choice.

The onset of the contrast bolus is determined by tracking the mean signal intensity over time. To ensure consistent absolute values, the mean signal is normalized to the first time point. Bolus arrival is then identified using a sliding window of three time points: when the mean normalized signal within the window changes by more than a user-defined threshold (default 5%) compared to the preceding window, the bolus onset is marked. The contrast baseline image (S_0) is calculated by averaging all volumes acquired prior to bolus arrival. Contrast time curves (CTCs) are derived differently for CTP and MRP. In CTP, the CTC is obtained as $CTC = S - S_0$, reflecting the approximately linear relationship between iodine concentration and x-ray attenuation, which directly affects signal intensity. In MRP, the CTC is computed as $CTC = \frac{1}{\text{echo_time}} \cdot \log\left(\frac{S}{S_0}\right)$, since the contrast agent influences relaxation times rather than directly altering signal intensity. TTP maps mark the time at which each CTC reaches its peak, with $t=0$ set to the contrast bolus onset.

The arterial input function (AIF) is extracted using an adaptive thresholding approach. First, the distribution of the TTP and the area under the CTC are computed. A series of progressively relaxed distribution-percentile-based thresholds for TTP and AUC are then applied to identify candidate voxels likely to represent an arterial signal. Arterial voxels are expected to exhibit low TTP and high AUC. Connected component analysis is used to isolate distinct clusters of candidate AIF voxels. For each cluster, the mean CTC is extracted. Clusters are excluded if their mean CTC is excessively noisy or if the cluster size exceeds user-defined thresholds. The mean CTC of the remaining clusters is then fitted with a gamma variate function. Finally, a scoring function that combines cluster volume, curve peak, and mean fitting error is applied to select the final AIF segmentation.

After that, the residue function (R) is determined by deconvolving the CTC with the AIF. In practice, this is done via singular value decomposition (SVD). PyPeT offers an implementation for standard SVD, block-circulant SVD with Simpson’s rule integration [20], block-circulant SVD with manual matrix construction, and oscillation index SVD with adaptive regularization. Next, CBF is determined as the maximum of R. Tmax is the time index corresponding to this maximum. CBV is defined as the integral of R. Lastly, MTT is calculated by dividing CBV by CBF. This information is summarized in Figure 1 and Equations 1-6.

2.2 Debug Mode for Intermediate Visualizations

PyPeT provides a debugging mode that can be enabled by setting the `DEBUG` flag in `main.py` to `True`. When activated, the pipeline generates intermediate

visualizations at key processing stages. These figures allow users to verify that each step produces the expected output, making it easier to identify and correct errors before they propagate to the final results.

The first viewer displays the preprocessed perfusion volumes in their original 4D form as well as maximum or minimum intensity projections, depending on the modality. Users can scroll through both spatial and temporal frames (example in Figure 2). The second viewer shows the generated 3D brain mask together with the brain segmentation, with spatial and temporal scrolling enabled (example in Figure 3). Another viewer shows the normalized mean signal intensity over time along with the detected bolus arrival index S_0 (example in Figure 4). In addition, the CTCs are visualized with both temporal and spatial scrolling (example in Figure 5). Finally, the AIF segmentation is displayed as an overlay on the CTC image. This viewer also shows the mean CTC within the AIF region and its fitted gamma variate model, again with temporal and spatial scrolling supported (example in Figure 6).

2.3 Technical Details

PyPeT was developed using Python 3.10 and relies on widely used packages such as NumPy, SciPy, scikit, and matplotlib (see requirements.txt on our GitHub repository). The code was developed on Windows 11 but should run on most systems without requiring a GPU or high-performance CPU. The GitHub repository includes one demo CTP image with dimensions $200 \times 200 \times 23 \times 45$ (in-plane x, in-plane y, number of slices, number of time points). Processing this image takes about 17 seconds on an Intel i5 processor, of which 11 seconds are spent on registering the 45 three-dimensional volumes.




Fig. 2. Left: snapshot of the debug viewer showing a raw 4D input CTP image. Right: corresponding maximum intensity projection of the middle third of slices.




Fig. 3. Left: snapshot of the debug viewer showing the generated mask. Middle: the corresponding perfusion image. Left: The corresponding brain segmentation.




Fig. 4. Snapshot of the debug viewer showing normalized mean signal intensity over time, with the start of the bolus highlighted.




Fig. 5. Left: snapshot of the debug viewer showing the CTC image. Right: corresponding maximum intensity projection of the middle third of slices.




Fig. 6. Left: slice with the AIF segmentation overlay on the CTC image. Middle: the CTC images without the AIF overlay. Right: the mean signal in the AIF segmentation, and its gamma-variate fit.

3 Validation Experiments

The performance of PyPeT was validated by comparing its output perfusion maps against reference maps obtained from established commercial and research pipelines. PyPeT-derived perfusion maps (CBF, CBV, MTT, TTP, and Tmax) were compared against the reference maps both visually and quantitatively using the structural similarity index (SSIM).

3.1 CT Perfusion Validation

Three independent reference sources were used for CTP. These are 1) Icobrain CVA (FDA-approved, commercially available): ten CTP scans were selected from the publicly available ISLES2024 dataset [17]. A dataset originally created for the Ischemic Stroke Lesion Segmentation Challenge at MICCAI 2024, which contains perfusion maps generated with Icobrain CVA; 2) iSchemaView Rapid (FDA-approved, commercially available): ten CTP scans were selected from the ISLES2018 dataset [8,2]. This dataset is also publicly available and was developed for stroke lesion segmentation challenges; 3) UniToBrain (Published in-house research pipeline): ten CTP scans were selected from the UniToBrain dataset [15]. This dataset was released by the University of Turin and is also publicly accessible. In this dataset, perfusion maps were generated using an in-house pipeline based on the boxNLR model, as described in [1].

3.2 MR Perfusion Validation

One independent reference source was used for MRP, which is Olea Sphere 3.0 (FDA-approved, commercially available). Eight MRP scans were selected from an internal dataset at Inselspital, Bern, Switzerland. This dataset contains stroke patients who received DSC-MRI imaging for diagnostic purposes.

4 Results

Figure 7 presents visual comparisons between the perfusion maps generated by PyPeT and *Icobrain CVA*. Note that all maps were normalized by the whole-brain mean to obtain relative perfusion values, since deriving absolute values is challenging and absolute values as well as reported units vary across commercial software packages [10,11]. The perfusion maps generated by PyPeT show a strong correspondence with those from *Icobrain CVA*. Additional visual comparisons with the other reference maps can be seen in the appendices.

A notable observation concerns the MTT maps. The PyPeT-derived MTT appears different from the corresponding *Icobrain CVA* map. This discrepancy becomes clearer when examining the color bar ranges. Although the overall spatial patterns are similar, differences in color scaling create the impression of a larger divergence than is actually present. This is supported by the SSIM values

reported in Table 1, which confirm high structural similarity between the two maps.

Table 1 summarizes the mean SSIM and corresponding standard deviations across datasets. The results are consistent, with values typically ranging between 0.75 and 0.85. The main exception is the MTT comparison for *Icobrain CVA*, which shows a lower agreement. Overall, the standard deviations are rather small, varying between 0.01 and 0.05.

Table 1. SSIM of the comparisons between PyPeT and the reference maps.

Reference	Modality	Parameter	Mean	SD
Icobrain CVA	CTP	CBF	0.83	0.01
		CBV	0.78	0.03
		MTT	0.66	0.03
		Tmax	0.76	0.04
iSchemaView Rapid	CTP	CBF	0.83	0.05
		CBV	0.80	0.04
		MTT	0.75	0.04
		Tmax	0.80	0.05
UniToBrain	CTP	CBF	0.81	0.03
		CBV	0.85	0.02
		MTT	0.78	0.03
		TTP	0.80	0.02
Olea Sphere 3.0	MRP	CBF	0.80	0.02
		CBV	0.79	0.02
		TTP	0.78	0.02
		Tmax	0.80	0.02




Fig. 7. Visual comparisons of perfusion maps generated by PyPeT with those by Ico-brain CVA. This scan shown is provided as a demo image in our GitHub repository.

5 Discussion and Conclusion

This work introduced PyPeT, an open Python-based tool for generating quantitative perfusion maps from head CTP and MRP images. The tool is designed with a modular architecture, clear documentation, and low computational requirements, aiming to make perfusion imaging research as accessible and customizable as possible. PyPeT provides an extensive debug mode to facilitate transparency in the processing steps. Validation against reference perfusion maps from four different sources demonstrated consistent structural similarity, with SSIM values typically ranging between 0.75 and 0.85. While these results indicate strong correspondence with FDA approved alternatives, PyPeT is intended solely for research use. It has not undergone the level of validation required for clinical deployment.

A major challenge in validation arises from the unknown post-processing applied by commercial software, which limits the degree of visual similarity achievable between PyPeT-generated maps and reference outputs. It is important to note that discrepancies with commercial software do not necessarily reflect inaccuracies in PyPeT-generated maps. In contrast, PyPeT offers fully transparent and reproducible results, allowing researchers to interpret and adjust the outputs more effectively than is possible with commercial tools.

It is worth noting that PyPeT remains under active development. Future work should focus on expanding functionality, improving validation strategies, and incorporating user feedback to strengthen its utility as a research resource. Some concrete points for improvement include 1) **improving absolute perfusion values:** the comparisons presented in this work were based on relative perfusion maps. To ensure that perfusion maps have accurate absolute values, it is important to tune physiological constants such as the tissue density and the hematocrit correction factor. Secondly, one may revisit the quality of arterial input function selection, and the regularization of the deconvolution process; 2) **investigation of MTT discrepancies:** although CBV and CBF maps closely match those produced by the commercial tool, the MTT maps show greater variation, especially for the comparison with Icobrain. Since MTT is theoretically the ratio of CBV to CBF, further investigation is needed to understand the post-processing strategies employed by Icobrain and to reconcile these differences; 3) **additional validation:** by applying PyPeT to additional MRP datasets, we could strengthen confidence in its generalizability and performance, especially for MRP data.

In conclusion, PyPeT lowers the barrier to entry for brain perfusion research and contributes to the growing movement toward reproducible and openly accessible medical imaging analysis.

Acknowledgment. This work is supported by the MIMIC project, funded by the NWO NGF AiNed XS Europe grant (NGF.1609.242.047). We also thank the Inselspital, and Richard McKinley and Roland Wiest in particular, for their discussions and for providing the MRP dataset used for validation.

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

References

1. Bennink, E., Oosterbroek, J., Kudo, K., Viergever, M.A., Velthuis, B.K., de Jong, H.W.A.M.: Fast nonlinear regression method for ct brain perfusion analysis. *Journal of Medical Imaging* **3**, 026003 (6 2016). <https://doi.org/10.1117/1.jmi.3.2.026003>
2. Cereda, C.W., Christensen, S., Campbell, B.C., Mishra, N.K., Mlynash, M., Levi, C., Straka, M., Wintermark, M., Bammer, R., Albers, G.W., Parsons, M.W., Lansberg, M.G.: A benchmarking tool to evaluate computer tomography perfusion infarct core predictions against a dwi standard. *Journal of Cerebral Blood Flow and Metabolism* **36**, 1780–1789 (10 2016). <https://doi.org/10.1177/0271678X15610586>
3. Chandrabhatla, A.S., Kuo, E.A., Sokolowski, J.D., Kellogg, R.T., Park, M., Mastorakos, P.: Artificial intelligence and machine learning in the diagnosis and management of stroke: A narrative review of united states food and drug administration-approved technologies. *Journal of Clinical Medicine* **12** (6 2023). <https://doi.org/10.3390/jcm12113755>
4. Ekmekci, B., McAnany, C.E., Mura, C.: An introduction to programming for bioscientists: A python-based primer. *PLoS Computational Biology* **12** (6 2016). <https://doi.org/10.1371/journal.pcbi.1004867>
5. Fernández-Rodicio, S., Ferro-Costas, G., Sampedro-Viana, A., Bazarra-Barreiros, M., Ferreirós, A., López-Arias, E., Pérez-Mato, M., Ouro, A., Pumar, J.M., Mosqueira, A.J., Alonso-Alonso, M.L., Castillo, J., Hervella, P., Iglesias-Rey, R.: Perfusion-weighted software written in python for dsc-mri analysis. *Frontiers in Neuroinformatics* **17** (2023). <https://doi.org/10.3389/fninf.2023.1202156>
6. Gava, U.A., D’Agata, F., Tartaglione, E., Renzulli, R., Grangetto, M., Bertolino, F., Santonocito, A., Bennink, E., Vaudano, G., Boghi, A., Bergui, M.: Neural network-derived perfusion maps: A model-free approach to computed tomography perfusion in patients with acute ischemic stroke. *Frontiers in Neuroinformatics* **17** (2023). <https://doi.org/10.3389/fninf.2023.852105>
7. Gordaliza, P.M., Mateos-Pérez, J.M., Montesinos, P., de Villoria, J.A.G., Desco, M., Vaquero, J.J.: Development and validation of an open source quantification tool for dsc-mri studies. *Computers in Biology and Medicine* **58**, 56–62 (3 2015). <https://doi.org/10.1016/j.combiomed.2015.01.002>
8. Hakim, A., Christensen, S., Winzeck, S., Lansberg, M.G., Parsons, M.W., Lucas, C., Robben, D., Wiest, R., Reyes, M., Zaharchuk, G.: Predicting infarct core from computed tomography perfusion in acute ischemia with machine learning: Lessons from the isles challenge. *Stroke* **52**, 2328–2337 (7 2021). <https://doi.org/10.1161/STROKEAHA.120.030696>
9. van Houdt, P.J., Ragunathan, S., Berks, M., Ahmed, Z., Kershaw, L.E., Gurney-Champion, O.J., Tadimalla, S., Arvidsson, J., Sun, Y., Kallehauge, J., Dickie, B., Lévy, S., Bell, L., Sourbron, S., Thrippleton, M.J.: Contrast-agent-based perfusion mri code repository and testing framework: Ismrm open science initiative for perfusion imaging (osipi). *Magnetic Resonance in Medicine* **91**, 1774–1786 (5 2024). <https://doi.org/10.1002/mrm.29826>

10. Korfiatis, P., Kline, T.L., Kelm, Z.S., Carter, R.E., Hu, L.S., Erickson, B.J.: Dynamic susceptibility contrast-mri quantification software tool: Development and evaluation. *Tomography* **2**, 448–456 (2016). <https://doi.org/10.18383/j.tom.2016.00172>
11. Kudo, K., Uwano, I., Hirai, T., Murakami, R., Nakamura, H., Fujima, N., Yamashita, F., Goodwin, J., Higuchi, S., Sasaki, M.: Comparison of different post-processing algorithms for dynamic susceptibility contrast perfusion imaging of cerebral gliomas. *Magnetic Resonance in Medical Sciences* **16**, 129–136 (2017). <https://doi.org/10.2463/mrms.mp.2016-0036>
12. Lirette, S.T., Smith, A.D., Aban, I.B.: A tool to visualize and analyze perfusion data: Development and application of the r package “ctp”. *Computer Methods and Programs in Biomedicine* **153**, 11–17 (1 2018). <https://doi.org/10.1016/j.cmpb.2017.09.016>
13. Lohrke, F., Madai, V.I., Kossen, T., Aydin, O.U., Behland, J., Hilbert, A., Mutke, M.A., Bendszus, M., Sobesky, J., Frey, D.: Perfusion parameter map generation from tof-mra in stroke using generative adversarial networks. *NeuroImage* **298** (9 2024). <https://doi.org/10.1016/j.neuroimage.2024.120770>
14. Lévy, S., Roche, P.H., Guye, M., Callot, V.: Feasibility of human spinal cord perfusion mapping using dynamic susceptibility contrast imaging at 7t: Preliminary results and identified guidelines. *Magnetic Resonance in Medicine* **85**, 1183–1194 (3 2021). <https://doi.org/10.1002/mrm.28559>
15. Perlo, D., Tartaglione, E., Gava, U., D’Agata, F., Benninck, E., Bergui, M.: Uniotrain dataset: A brain perfusion dataset. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. vol. 13373 LNCS, pp. 498–509. Springer Science and Business Media Deutschland GmbH (2022). https://doi.org/10.1007/978-3-031-13321-3_44
16. Peruzzo, D., Bertoldo, A., Zanderigo, F., Cobelli, C.: Automatic selection of arterial input function on dynamic contrast-enhanced mr images. *Computer Methods and Programs in Biomedicine* **104** (12 2011). <https://doi.org/10.1016/j.cmpb.2011.02.012>
17. Riedel, A.E.O., de la Rosa, E., Baran, T.A., Petzsche, M.H., Baazaoui, H., Yang, K., Musio, F.A., Huang, H., Robben, D., Seia, J.O., Wiest, R., Reyes, M., Su, R., Zimmer, C., Boeckh-Behrens, T., Berndt, M., Menze, B., Rueckert, D., Wiestler, B., Wegener, S., Kirschke, J.S., Riedel, E.O.: Isles’24-a real-world longitudinal multimodal stroke dataset. *arXiv* (2024). <https://doi.org/10.48550/arXiv.2408.11142>
18. turtleizy: https://github.com/turtleizy/ctp_csvd (2021)
19. de Vries, L., van Herten, R.L., Hoving, J.W., Išgum, I., Emmer, B.J., Majoie, C.B., Marquering, H.A., Gavves, E.: Spatio-temporal physics-informed learning: A novel approach to ct perfusion analysis in acute ischemic stroke. *Medical Image Analysis* **90** (12 2023). <https://doi.org/10.1016/j.media.2023.102971>
20. Zanderigo, F., Bertoldo, A., Pillionetto, G., Cobelli, C.: Nonlinear stochastic regularization to characterize tissue residue function in bolus-tracking mri: Assessment and comparison with svd, block-circulant svd, and tikhonov. *IEEE Transactions on Biomedical Engineering* **56**, 1287–1297 (5 2009). <https://doi.org/10.1109/TBME.2009.2013820>

6 Appendices




Fig. 8. Visual comparisons of perfusion maps generated by PyPeT compared to maps generated by iSchemaView Rapid.




Fig. 9. Visual comparisons of perfusion maps generated by PyPeT compared to maps generated by UniToBrain.




Fig. 10. Visual comparisons of perfusion maps generated by PyPeT compared to maps generated by Olea Sphere 3.0.