p-Brain: An Automated MRI Pipeline for Cerebral Perfusion, Microvasculature, and Blood–Brain Barrier Permeability Estimation

Edis D. Tireli[1,2], Stig P. Cramer[1], Ulrich Lindberg[1], Mark B. Vestergaard[1], Derya Tireli[1,2], and Henrik B. W. Larsson[1,2]

[1]Functional Imaging Unit, Department of Clinical Physiology and Nuclear Medicine, Copenhagen University Hospital, Rigshospitalet, Glostrup, Denmark
[2]Department of Clinical Medicine, Faculty of Health and Medical Sciences, University of Copenhagen, Copenhagen, Denmark

Corresponding author: Edis D. Tireli (edis.tireli@sund.ku.dk)

Abstract

We present p-Brain, an end-to-end neuroimaging analysis framework for reproducible, automated quantitative DCE-MRI analysis at scale. From standard acquisitions, p-Brain estimates baseline relaxation parameters, converts signal to gadolinium concentration, derives arterial and venous input functions using convolutional neural network (CNN) slice selection and ROI segmentation, and produces voxelwise maps with regional and whole-brain summaries. The pipeline implements Patlak graphical analysis to estimate the blood–brain barrier influx constant (K_i) and plasma volume fraction (v_p), and performs model-free residue deconvolution with Tikhonov regularisation to estimate cerebral blood flow (CBF), mean transit time (MTT), and capillary transit-time heterogeneity (CTH) from the same DCE dataset. p-Brain exports analysis-ready outputs, intermediate readouts, structured runtime metadata, and stage-level quality control artifacts to support auditability in batch processing. We evaluate the framework on a technically uniform set of 97 DCE-MRI scans from 58 healthy human participants, and show close agreement between automated Patlak K_i summaries and an established reference workflow. A companion macOS desktop application supports batch execution, job monitoring, and rapid review of curves and maps. p-Brain is open-source and configurable, enabling extension to additional kinetic models.

Keywords: dynamic contrast-enhanced MRI; blood–brain barrier; permeability; perfusion; automation

List of Acronyms

AIF arterial input function

BBB blood-brain barrier

CBF cerebral blood flow

CNN convolutional neural network

CTH capillary transit-time heterogeneity

DCE dynamic contrast-enhanced

DCE-MRI dynamic contrast-enhanced magnetic resonance imaging

EES extravascular-extracellular space

FSL FMRIB Software Library

GM gray matter

GRE gradient-recalled echo

IR inversion recovery

mTBI mild traumatic brain injury

MRI magnetic resonance imaging

MTT mean transit time

NIfTI Neuroimaging Informatics Technology Initiative (file format)

PPCS persistent post-concussion symptoms

rICA right internal carotid artery

ROI region of interest

SPGR spoiled gradient recalled echo

SSS superior sagittal sinus

TE echo time

TR repetition time

U-Net U-shaped convolutional network

VFA variable flip angle

VOF venous output function

WM white matter

1 Introduction

The blood–brain barrier (BBB) regulates cerebral homeostasis, and BBB dysfunction is implicated in many neurological disorders. Accurate quantification of BBB permeability can be used to study neurological diseases such as multiple sclerosis, brain tumors, and neurodegenerative disorders (Cramer et al., 2014; Okuchi et al., 2019; Heye et al., 2014; Chagnot et al., 2021; Raja et al., 2018). Dynamic contrast-enhanced MRI (DCE-MRI) enables in vivo quantification of subtle BBB leakage (Varatharaj et al., 2019; Ingrisch et al., 2012) via tracer-kinetic models such as Patlak graphical analysis (Patlak et al., 1983) and other compartmental models. These approaches estimate permeability-related parameters (e.g., K_i) and complementary metrics. Prior work supports the biological validity and clinical relevance of these measurements: DCE-MRI-derived K_i shows expected gray/white matter differences and lesion sensitivity (Varatharaj et al., 2019), can be reproducible even at low leakage levels (Cramer et al., 2023a), and has demonstrated prognostic value (e.g., optic neuritis to multiple sclerosis conversion) (Cramer et al., 2015). Beyond demyelinating disease, BBB permeability mapping has been applied in acute ischemic stroke and reperfusion injury (Villringer et al., 2017; Arba et al., 2021), traumatic brain injury (Ware et al., 2022), seizure disorders such as new-onset refractory status epilepticus (Li et al., 2023), autoimmune and systemic inflammatory disease including autoimmune encephalitis and systemic lupus erythematosus (Ji et al., 2024; Chi et al., 2019), infectious neuroinflammation such as neuroborreliosis (Lindland et al., 2023), and HIV-associated neurocognitive impairment (Chaganti et al., 2019). Together, these findings motivate automated, user-independent tools for estimating BBB permeability and related perfusion parameters at scale.

Despite the potential of DCE-MRI, analysis often remains labor-intensive. Manual ROI definition (e.g., arterial input functions and tissue curves) and supervision of model fitting introduce operator variability and limit scalability (Varatharaj et al., 2024a; Beresford et al., 2006; Calamante, 2013; Keil et al., 2017). This motivates fully automated pipelines that integrate preprocessing, ROI extraction, and pharmacokinetic modelling while preserving quantitative fidelity (Varatharaj et al., 2024a; Barnes et al., 2015; Nalepa et al., 2020).

To overcome these challenges, we introduce *p*-Brain, a fully automated pipeline that integrates convolutional neural network (CNN)-driven ROI segmentation, adaptive automated arterial input extraction, and tracer-kinetic modeling (Patlak graphical analysis) alongside model-free deconvolution for perfusion and microvascular transit metrics. The design targets complete removal of user interaction while preserving agreement with established analyses. In addition to BBB permeability and plasma volume, *p*-Brain estimates cerebral blood flow (CBF) and residue-derived descriptors of mean transit time (MTT) and capillary transit-time heterogeneity (CTH), enabling interrogation of both bulk perfusion and microvascular dispersion within the same workflow. The framework is designed to be extensible, enabling additional kinetic models to be incorporated as needed. Our pipeline delivers robust, reproducible metrics at multiple analytical scales (voxel-wise, regional, whole-brain) without requiring user interaction.

The workflow is designed to be configurable, with key analysis choices (e.g., input-function strategy, curve extraction, and optional normalization) documented online.

The emergence of fast, deep learning-based neuroimaging pipelines (e.g., FastSurfer (Henschel et al., 2020)) enables whole-brain structural segmentation within minutes, offering a potential solution for consistent ROI definition across subjects. *p*-Brain leverages these advances to automatically generate tissue masks and parcellations for regional summarisation, allowing however the user to use their own choice of segmentation software as well. In this paper, we describe the implementation of *p*-Brain and validate its performance against reference methods and literature values.

Data from DCE-MRI are obtained according to clinical protocols but at minimum requires a 3D T_1 weighted image, a series of inversion times, T_I, recovery sequences for the estimation of T_1 and M_0 parameters, as well as the DCE-MR image. Temporal resolution and duration should be optimized to ensure precise kinetic modeling. MRI data undergo axial reslicing of structural volumes to the native DCE frame.

2 Results

We applied p-Brain to a technically uniform subset of **97** DCE-MRI scans from **58** participants with mild traumatic brain injury (mTBI) but no macroscopic lesions on structural MRI (PPCS cohort; Methods). Each dataset was processed end-to-end with identical settings in fully automatic mode. Deliverables include voxelwise, regional, and whole-brain outputs of BBB influx (K_i), plasma volume (v_p), cerebral blood flow (CBF), capillary transit-time heterogeneity (CTH), and mean transit time (MTT), and structured run summaries saved alongside the maps.

Figures 1–9 show representative outputs from a fully automatic run. Voxelwise maps capture within-scan spatial heterogeneity, while parcel- and cohort-level summaries support group comparisons and downstream statistics.

Voxelwise maps

Voxelwise maps quantify model-derived parameters at the native spatial scale, preserving continuous structure and enabling localized assessment of heterogeneity that can be obscured by regional aggregation.

To illustrate the voxelwise approach, we turn first to K_i, a parameter by default within the pipeline estimated from Patlak graphical analysis (see Patlak graphical analysis) describing the unidirectional transfer of contrast agent from plasma to tissue, reflecting the blood–brain barrier permeability. The map in Figure 1 displays the voxelwise estimation of this parameter across the entire brain volume, providing a spatially continuous quantification of barrier function.

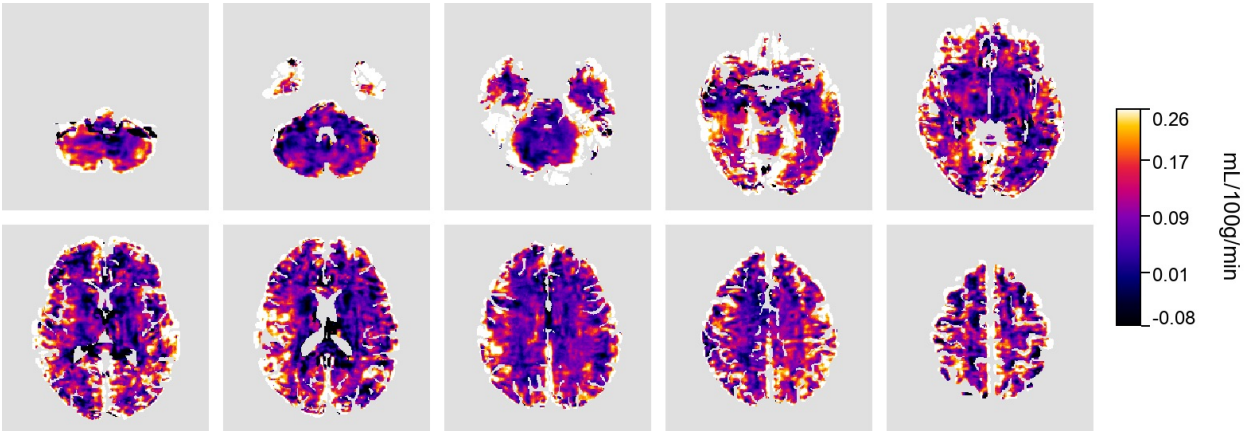

Figure 1: Automated voxelwise K_i map in units of mL/100g/min of the representative example. Local heterogeneities can be seen within and across tissue.

On maps such as these, focal or asymmetric K_i elevations could signify localised blood–brain barrier leakage or microvascular dysfunction, which may be relevant in both diffuse and focal pathologies. Turning to perfusion, Figure 2 shows voxelwise cerebral blood flow (CBF) estimated by model-free residue deconvolution (elaborated in Modelling).

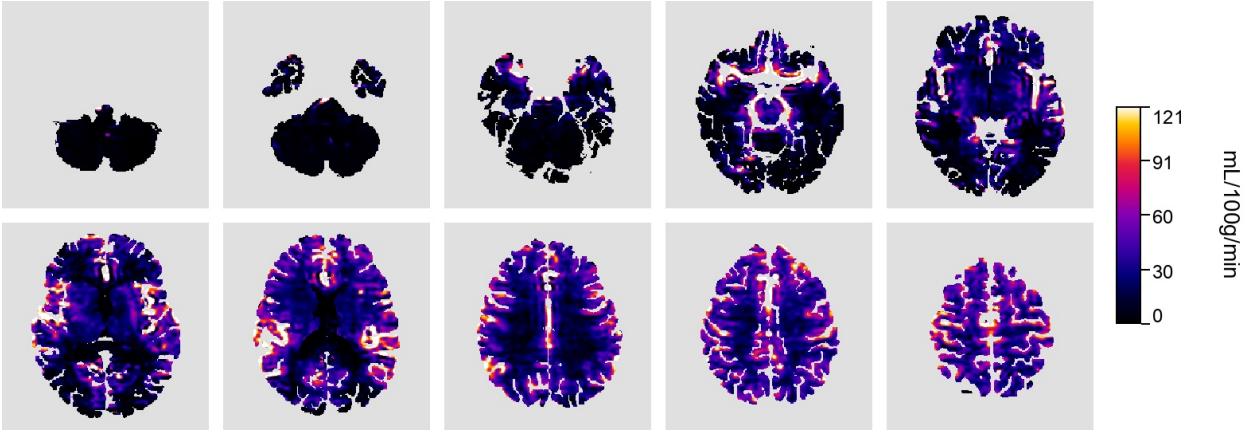

Figure 2: Voxelwise CBF from model-free deconvolution of the representative participant in units of mL/100g/min. Local heterogeneities can be seen as well as elevated CBF around the neurovasculature; notably in this dataset, the arteries forming the circle of Willis are visible in the lower slices.

The voxelwise plasma volume map (Figure 3) reports the fractional intravascular volume v_p estimated from the Patlak intercept. In practice, v_p highlights regions where the tissue signal is dominated by blood rather than extravascular uptake, and it is therefore sensitive to vascular density, large-vessel partial-volume effects, and residual vascular signal in the tissue curves. Interpreting v_p alongside K_i helps separate true exchange from apparent leakage driven by vascular volume (e.g., high v_p with low K_i vs. low v_p with elevated K_i), and provides a complementary context for perfusion measures such as CBF.

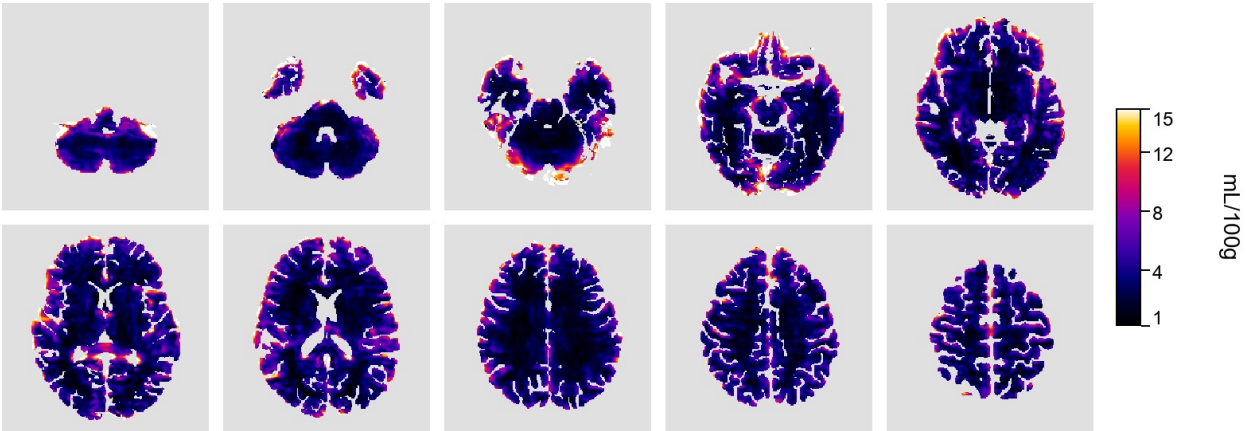

Figure 3: Voxelwise plasma volume (v_p; Patlak intercept) for the representative dataset in units of mL/100g.

Capillary transit-time heterogeneity (CTH) summarises the dispersion of microvascular transit times estimated from the residue/outflow function (Angleys et al., 2015), shown in Figure 4. In practice, elevated CTH reflects a broader distribution of capillary passage times and can indicate increased microvascular flow heterogeneity, which may reduce effective oxygen extraction even when bulk flow is preserved (Larsson et al., 2017a; Vestergaard et al., 2023a). CTH maps can therefore be relevant when probing microvascular dysfunction in settings such as small-vessel disease, inflammatory or metabolic stress, and peri-lesional hemodynamic remodeling, and they are most interpretable when considered jointly with CBF and MTT (e.g., heterogeneous vs. uniformly delayed transit).

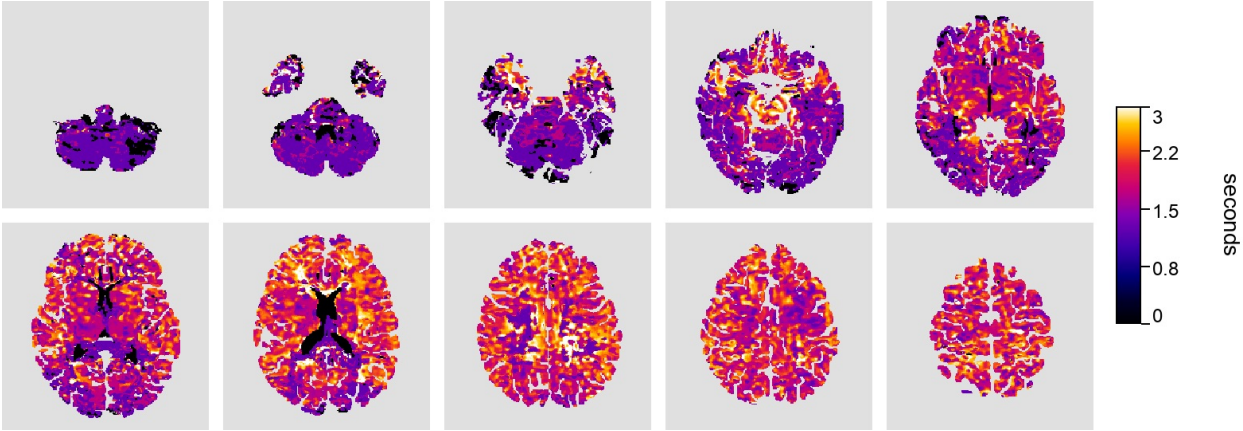

Figure 4: Voxelwise CTH derived from the normalized outflow $h(t) = -r'(t)/\int(-r')$ with units in s.

Fine-scale spatial texture in voxelwise CTH can reflect physiologic heterogeneity but is also sensitive to noise and deconvolution regularization; *p*-Brain therefore exports voxelwise maps alongside parcel and cohort summaries to support reproducible interpretation across subjects.

Mean transit time (MTT) summarizes the average duration of contrast passage through the microvasculature as derived from the first moment of the residue function (Figure 5). Elevated MTT indicates delayed transit and may be observed with reduced perfusion pressure, increased vascular resistance, or reliance on collateral pathways, whereas lower MTT is consistent with faster passage. MTT can be useful for characterizing perfusion delay in cerebrovascular disease and for contextualizing permeability or tissue injury measures; interpreted alongside CBF and CTH it helps distinguish slow-but-uniform flow from heterogeneous microvascular transit.

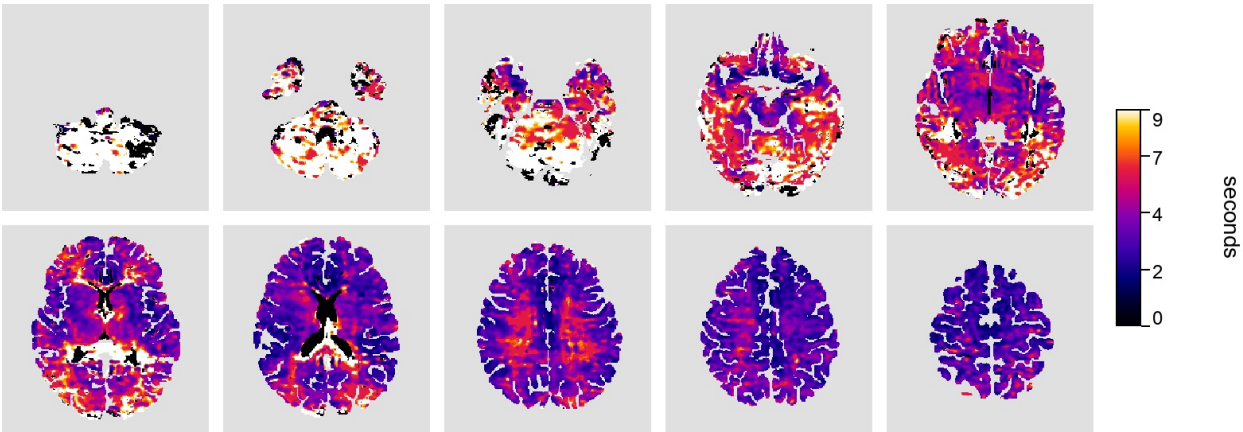

Figure 5: Voxelwise MTT derived from the first moment of the residue function with units in s.

Regional and parcellated organization

Voxelwise maps preserve fine spatial detail, but they are cumbersome for cohort statistics and can be disproportionately influenced by local noise, partial volume, and small residual misregistrations. If a high resolution 3D T1-weighted scan is available, *p*-Brain converts every quantitative output into a common anatomical representation by propagating the parcellation from segmentation[1] into DCE space and summarising each map by parcel; the mean or median value (user configured) is found within the ROI of the parcel. This produces, for every scan, a standardised regional feature set that can be analysed directly (e.g., longitudinal

[1]In this case, segmentation was used from FastSurfer, but in principle any segmentation software may work

tracking, group comparisons, and multi-site harmonisation) without additional scripting.

Figure 6 shows parcel-level CBF, illustrating how the workflow turns a voxelwise perfusion field into a regional map that is immediately comparable across participants. Figure 7 shows the corresponding parcel-wise Patlak K_i, providing an interpretable regional summary of BBB leakage that supports localisation of effects and downstream statistics. Additional parcel-level maps (CTH, MTT, and v_p) are provided in the Appendix (Figures 25–27).

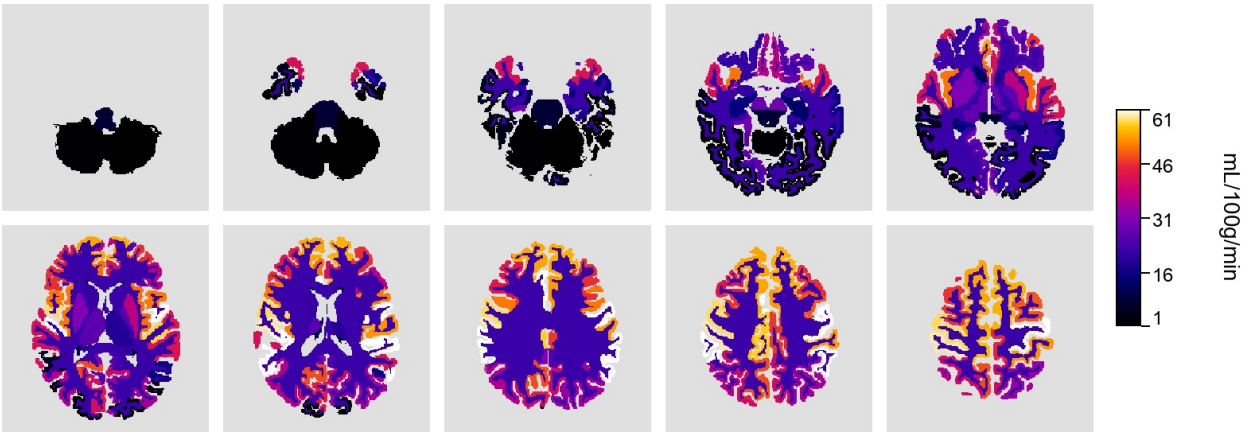

Figure 6: Parcel-level CBF in the representative participant. Units: mL/100g/min.

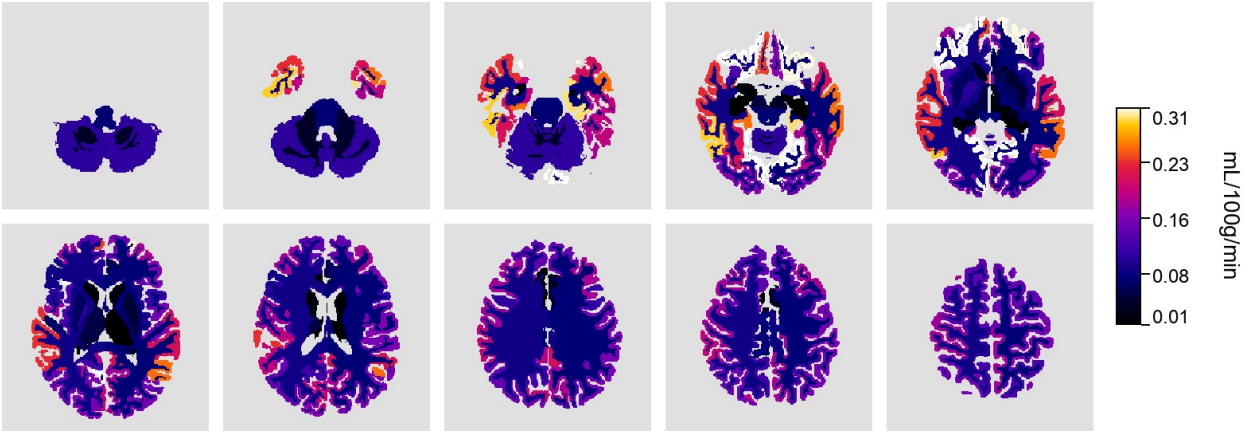

Figure 7: Parcel-level K_i within the representative participant.

These regional exports are therefore more than visual aids: they constitute ready-made datasets that can be pulled directly into statistical software, shared across centres, or revisited as new hypotheses emerge. By automating the parcellation, *p*-Brain lets researchers focus on biological interpretation—mapping where perfusion, permeability, and microvascular timing diverge—rather than on the logistics of producing the maps.

Cohort-level atlas projection

Group-level summaries are essential for quantitative DCE-MRI research because they enable comparisons against control populations, between subgroups within a disease, and across disease entities or phenotypes. *p*-Brain automatically aggregates parcel-level summaries across subjects and projects cohort means onto a common reference segmentation to provide cohort-level context for permeability and perfusion metrics.

Figure 8 shows the cohort-mean Patlak K_i projected onto the reference atlas. Additional cohort-mean atlas projections (CBF, CTH, and MTT) are provided in the Appendix (Figures 22–24).

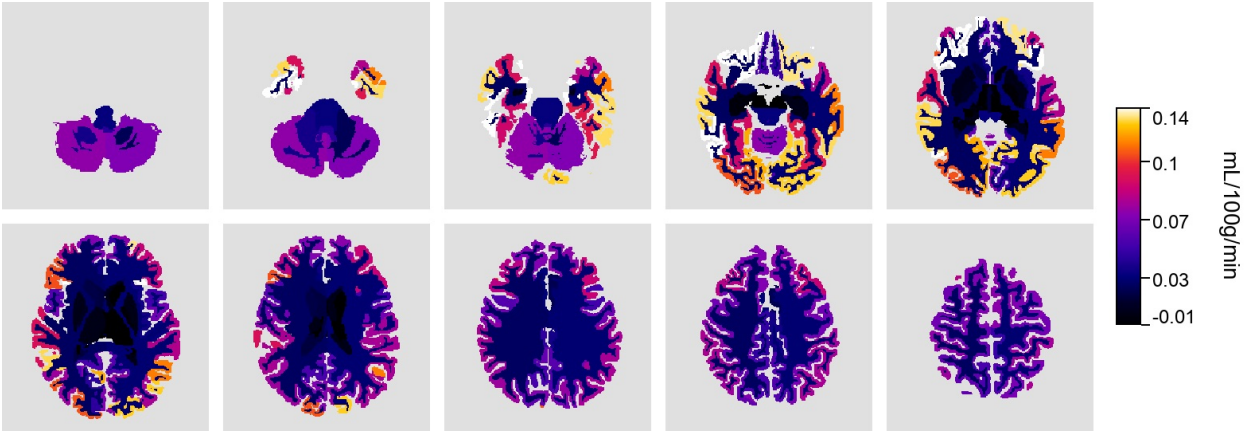

Figure 8: Cohort-mean K_i projected onto a reference atlas. Units: mL/100g/min.

Summarising these maps by anatomical parcels reduces the influence of voxelwise noise and small registration errors, emphasises regionally coherent effects, and enables consistent group-level comparisons.

Whole-brain medians

For scenarios requiring concise descriptors, *p*-Brain reports tissue-specific whole-brain medians derived from robust aggregate curves. Figure 9 shows that these medians preserve the gray/white ordering while condensing each scan to a handful of values. Such summaries can populate electronic health records, feed downstream statistical models, or serve as quality-control references alongside the richer voxelwise and parcelwise outputs.

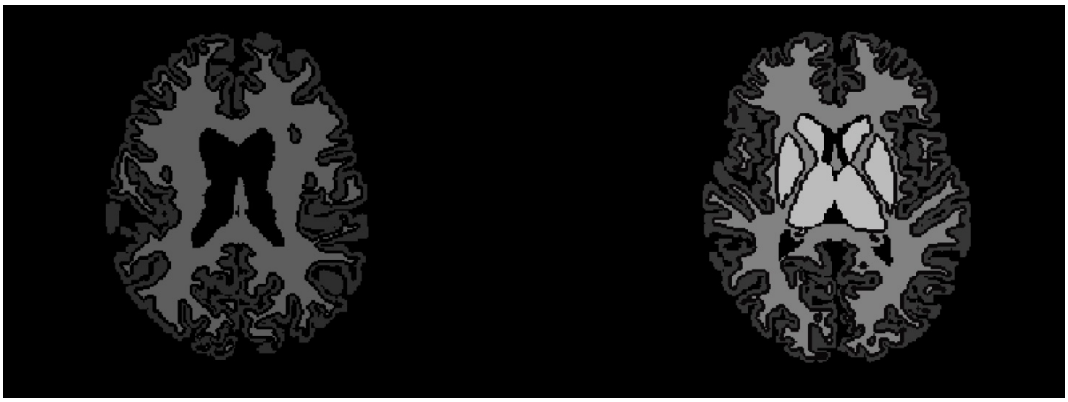

Figure 9: Whole-brain K_i medians per tissue provide scan-level summaries while preserving the GM>WM ordering.

End-to-end transparency

Automated analysis is made auditable by exporting a compact visual summary of intermediate outputs (Figure 10) together with structured stage-level QC reports (described in Methods). These artifacts enable rapid identification of failed or atypical runs without manual per-frame inspection.

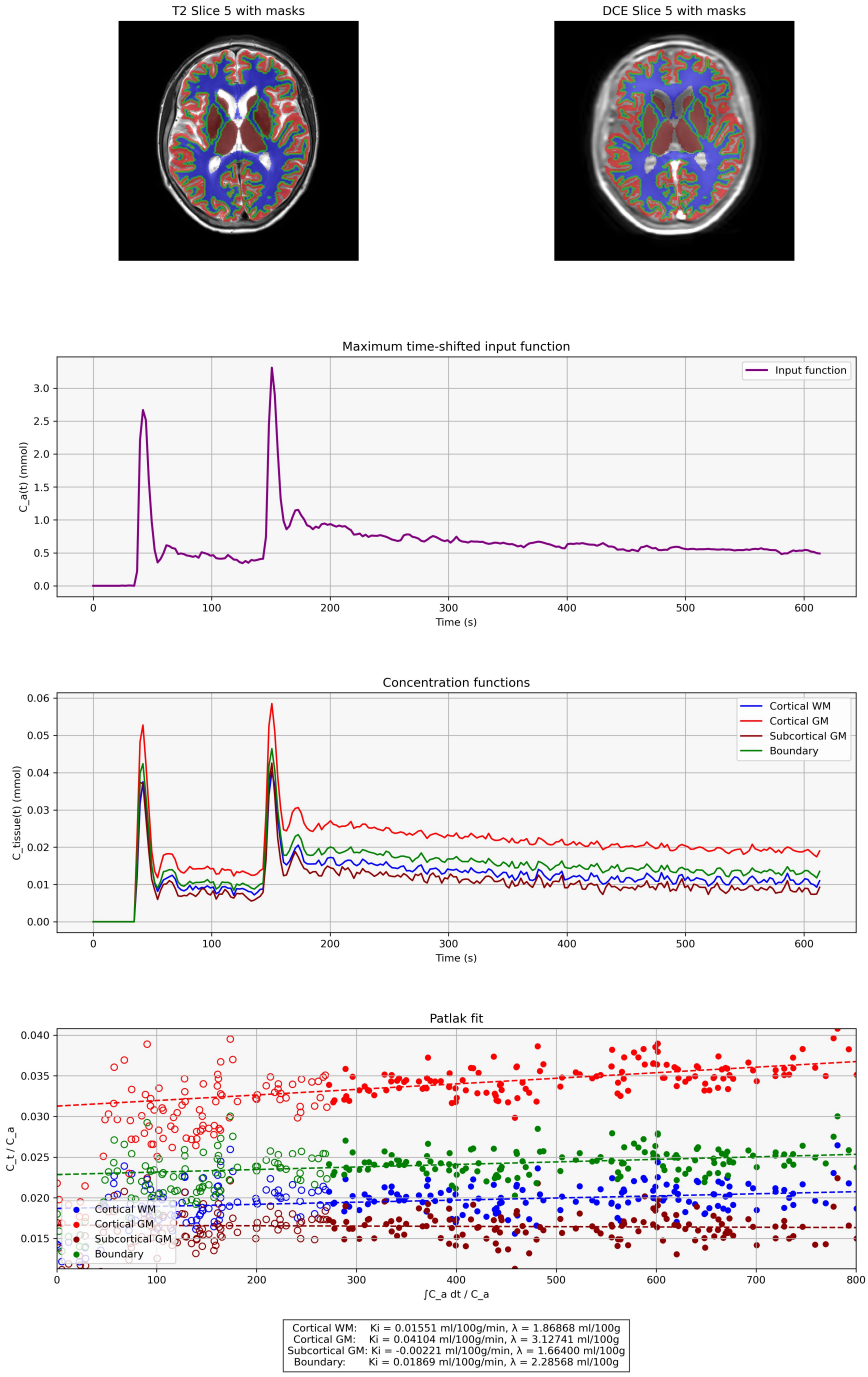

Figure 10: End-to-end transparency panel for a representative scan. We show segmentation, input functions, tissue curves, Patlak fits, and resulting K_i and λ (plasma volume, v_p).

Comparison to a reference implementation

We evaluated agreement between p-Brain and an established reference software workflow for Patlak-derived BBB permeability (K_i). For this comparison, the reference workflow used manually defined arterial/venous input-function ROIs and tissue ROIs, and p-Brain was run in both manual and fully automatic modes.

Figure 11 compares gray- and white-matter Patlak-derived permeability (K_i) summaries between p-Brain and the reference workflow across eight subjects with complete automatic outputs. In the reference workflow, tissue ROIs were manually placed in the right prefrontal cortex and right thalamus on a mid-brain slice (approximately slices 4–5 in the presented cohort), whereas p-Brain automatic measurements were derived from full-volume tissue masks (subcortical gray matter and cortical white matter). The between-method distributions showed no statistically significant differences in either tissue class.

Agreement was quantified using Pearson correlation and Bland–Altman analysis. For Patlak K_i, agreement with the reference workflow was high (GM: $r = 0.978$, bias $= -5.11 \times 10^{-4}$ mL/100g/min, 95% limits of agreement $[-0.0141, 0.0131]$; WM: $r = 0.987$, bias $= -2.22 \times 10^{-3}$ mL/100g/min, 95% limits $[-0.00896, 0.00453]$), indicating small systematic offsets and tight subject-level agreement for scan-level K_i summaries in this dataset.
In addition, voxelwise comparisons against the reference MATLAB program showed very good agreement for the perfusion map outputs and for the T_1/M_0 fitting outputs on a representative slice (Appendix Figures 20 and 21). In these comparisons, p-Brain was run in fully automatic mode, whereas the reference program relied on manual steps.

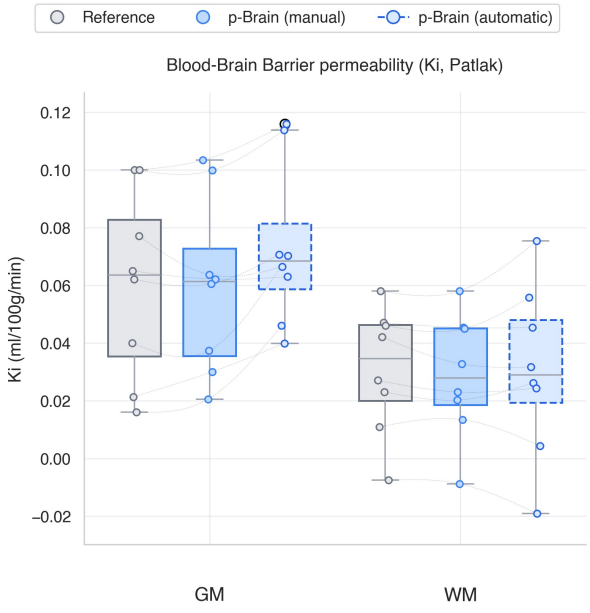

(a) Patlak K_i in GM and WM: reference workflow, p-Brain (manual), and p-Brain (automatic) across eight subjects.

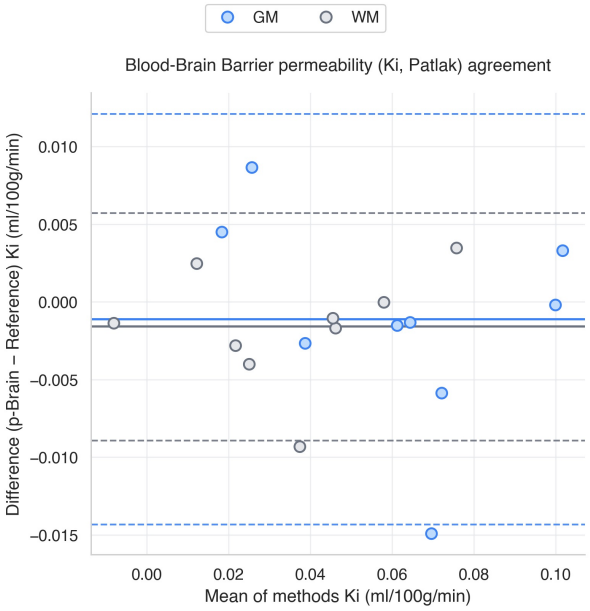

(b) Bland–Altman agreement for Patlak K_i in GM and WM: p-Brain (manual) vs. reference across eight subjects.

Figure 11: Automated comparison to an established manual reference workflow for whole-brain Patlak K_i. Left: between-method distributions; right: Bland–Altman agreement.

3 Discussion

The present work demonstrates that p-Brain can replace a traditionally manual, multi-software workflow for dynamic contrast-enhanced MRI with a reproducible automated pipeline. Direct agreement with reference implementations indicates that the automation does not sacrifice quantitative fidelity while eliminating operator-dependent variability that has historically hindered longitudinal and multi-site studies of blood–brain barrier physiology (Varatharaj et al., 2024a; Cramer et al., 2023a). By orchestrating segmentation, pharmacokinetic modeling, and multi-scale reporting in a single run, p-Brain offers a practical path toward routine deployment of quantitative DCE-MRI in research hospitals and, eventually, clinical radiology departments. Beyond providing reviewable figures, the pipeline exports analysis-ready outputs across scales—voxelwise NIfTI maps, parcel- and tissue-level summaries, cohort-level atlas projections, and scan-level whole-brain descriptors—together with structured logs and runtime metadata. This design supports downstream statistics and sharing across centres without requiring bespoke post-processing scripts.

The comparison to the reference workflow highlights expected sources of between-method variation, including differences in numerical implementation (e.g., optimization details) and the operator dependence of manually drawn ROIs. ROI definition and size can also influence stability: as ROI size increases, pharmacokinetic estimates generally become more stable due to averaging over noise and partial-volume effects, whereas K_i can appear inflated when estimated from small hand-drawn tissue ROIs compared with larger, atlas-derived tissue masks. Despite these factors, the observed bias and limits of agreement indicate close concordance of scan-level K_i summaries in this technically uniform dataset.

Our approach contrasts with previous semi-automated toolkits that still require labor-intensive arterial input selection, ROI drawing, or custom scripting to obtain voxel-wise and parcel-wise summaries. For example, ROCKETSHIP and related frameworks provide modular modeling utilities but expect considerable user supervision and do not guarantee harmonized outputs across centers (Barnes et al., 2015). Similarly, recent deep learning pipelines for tumor imaging often focus on specific pathologies and omit generalized whole-brain reporting or residue analysis (Nalepa et al., 2020). By comparison, p-Brain integrates CNN-based vascular and tissue segmentation, Patlak modeling, and model-free deconvolution to deliver BBB leakage, perfusion, and microvascular dispersion metrics in a unified exportable format.

In addition to the open-source Python implementation, we distribute a macOS desktop application that provides project organisation, batch job management, and responsive end-to-end execution on local machines, lowering the barrier to routine use by non-expert users while keeping sensitive data on-device.

The value of this integration is evident in the derived residue-based parameters. Capillary transit-time heterogeneity and mean transit time summarize complementary aspects of microvascular transport that influence oxygen delivery, and their inclusion enables interrogation of subtle hemodynamic alterations linked to vascular reserve and metabolic efficiency (Larsson et al., 2017b; Vestergaard et al., 2023b; Angleys et al., 2015). Because these summaries are produced alongside whole-brain and parcel medians, the same outputs can be reused for longitudinal monitoring, clinical decision support, or cross-cohort harmonization without revisiting the raw DCE data.

Automated anatomical segmentation underpins many of these advances. FastSurfer yields consistent gray- and white-matter parcellations within minutes (Henschel et al., 2020), and the released CNN models for vascular ROI extraction extend this principle to arterial input and venous outflow definition (Tireli, 2025). The resulting standardization reduces the learning curve for non-expert operators and mitigates variability arising from manual ROI placement (Cramer et al., 2023a). These gains are particularly relevant in large-scale studies of aging or neurodegeneration, where mild BBB leakage and perfusion changes must be detected against a backdrop of small effect sizes (Verheggen et al., 2020).

Finally, the pipeline preserves flexibility for future methodological advances. Because p-Brain exports intermediate concentration time series and segmentation masks, investigators can incorporate alternative kinetic models or machine learning predictors without reimplementing the ingestion and QC stages. Incorporating

sequence harmonization and uncertainty quantification are natural next steps to improve robustness across sites and patient populations.

Limitations

Despite these strengths, several practical constraints remain. First, the CNN models for anatomical and vascular segmentation were trained on relatively homogeneous data and can underperform in the presence of atypical anatomy, implants, or extensive pathology as well as sequences with varying contrast relationships, voxel dimensions or number of slices. When automated ROI extraction is not desired or is unsuitable for a given dataset, *p*-Brain can also be run with manual ROI delineation (vascular and/or tissue ROIs) via the GUI. For the segmentation used here, and as default, FastSurfer itself notes reduced reliability around large lesions, making visual inspection or targeted fine-tuning necessary when transferring to outlier populations (Henschel et al., 2020). Second, accurate kinetic modeling assumes physiologic regimes that may not hold universally. Patlak analysis presumes negligible backflux, and deconvolution-based perfusion estimates can be sensitive to input-function quality, bolus timing, and signal-to-noise; these assumptions can break in severe inflammatory states or other atypical physiology, biasing permeability and perfusion estimates (Varatharaj et al., 2019). Incorporating model selection logic, uncertainty quantification, or Bayesian approaches remains future work.

Motion and acquisition variability present additional challenges. Misregistration between the structural volume and the DCE time series directly perturbs the arterial input function and parcel propagation, degrading voxel-wise accuracy (Cramer et al., 2023a). Although the current pipeline performs rigid alignment and logs QC metrics, severe head motion or low signal-to-noise ratio datasets may still require reacquisition or more sophisticated correction. Finally, like most data-driven systems, *p*-Brain inherits domain shift sensitivity: scanner upgrades, protocol adjustments, or application to pediatric and geriatric cohorts can alter intensity distributions and vascular dynamics. Systematic validation across vendors and longitudinal recalibration of the CNN components and kinetic priors are essential before broad clinical deployment (Varatharaj et al., 2024a).

4 Materials and methods

4.1 Participants and MRI acquisition

We analysed DCE-MRI data acquired at Copenhagen University Hospital (Rigshospitalet) in a study of persistent post-concussion symptoms (PPCS) following mild traumatic brain injury (mTBI). Inclusion in the PPCS group required $\geq$3 months since mTBI with ongoing symptoms. All participants provided written informed consent. Ethical approval was granted by the Danish National Research Ethics Committee (H-21047857) and all procedures complied with the Declaration of Helsinki.

All MRI data were acquired on a Philips 3 T dSTREAM Achieva scanner with a 32-channel head coil. A high-resolution 3D T_1-weighted turbo field echo sequence was acquired with TE = 5.13 ms, TR = 11.18 ms, flip angle = 8°, and voxel size $\approx 0.67\,\text{mm} \times 0.67\,\text{mm} \times 0.70\,\text{mm}$.

The dynamic contrast-enhanced acquisition was a 4D time series with $256 \times 256 \times 10 \times 250$ samples (ten slices, 250 time points), voxel size $\approx 0.898\,\text{mm} \times 0.898\,\text{mm} \times 9.50\,\text{mm}$, echo time TE = 1.93 ms, and temporal sampling $\Delta t \approx 2.463\,\text{s}$ (total duration $\approx 616\,\text{s}$). Although evaluated on this specific protocol, *p*-Brain does not assume a fixed DCE geometry: spatial dimensions and temporal sampling are read from NIfTI headers and JSON sidecars when available, with configurable overrides via environment settings to support other acquisition geometries. The pipeline supports both single- and double-bolus injections; in our data, the blood signal can exhibit two bolus peaks, and the default settings are tuned to this regime.

For baseline relaxation mapping, an inversion-recovery series was acquired at inversion times $T_I \in \{120\,\text{ms}, 300\,\text{ms}, 600\,\text{ms}, 1000\,\text{m}$ with matching in-plane geometry to the DCE slices. A gadolinium-based contrast agent was administered intravenously according to the local clinical/research protocol.

4.2 Workflow overview

We summarize the end-to-end workflow in Figure 12. The gray shading marks fully-automatic steps. White boxes run with assisted automation and may, if desired, require human input for ROI drawing and review.

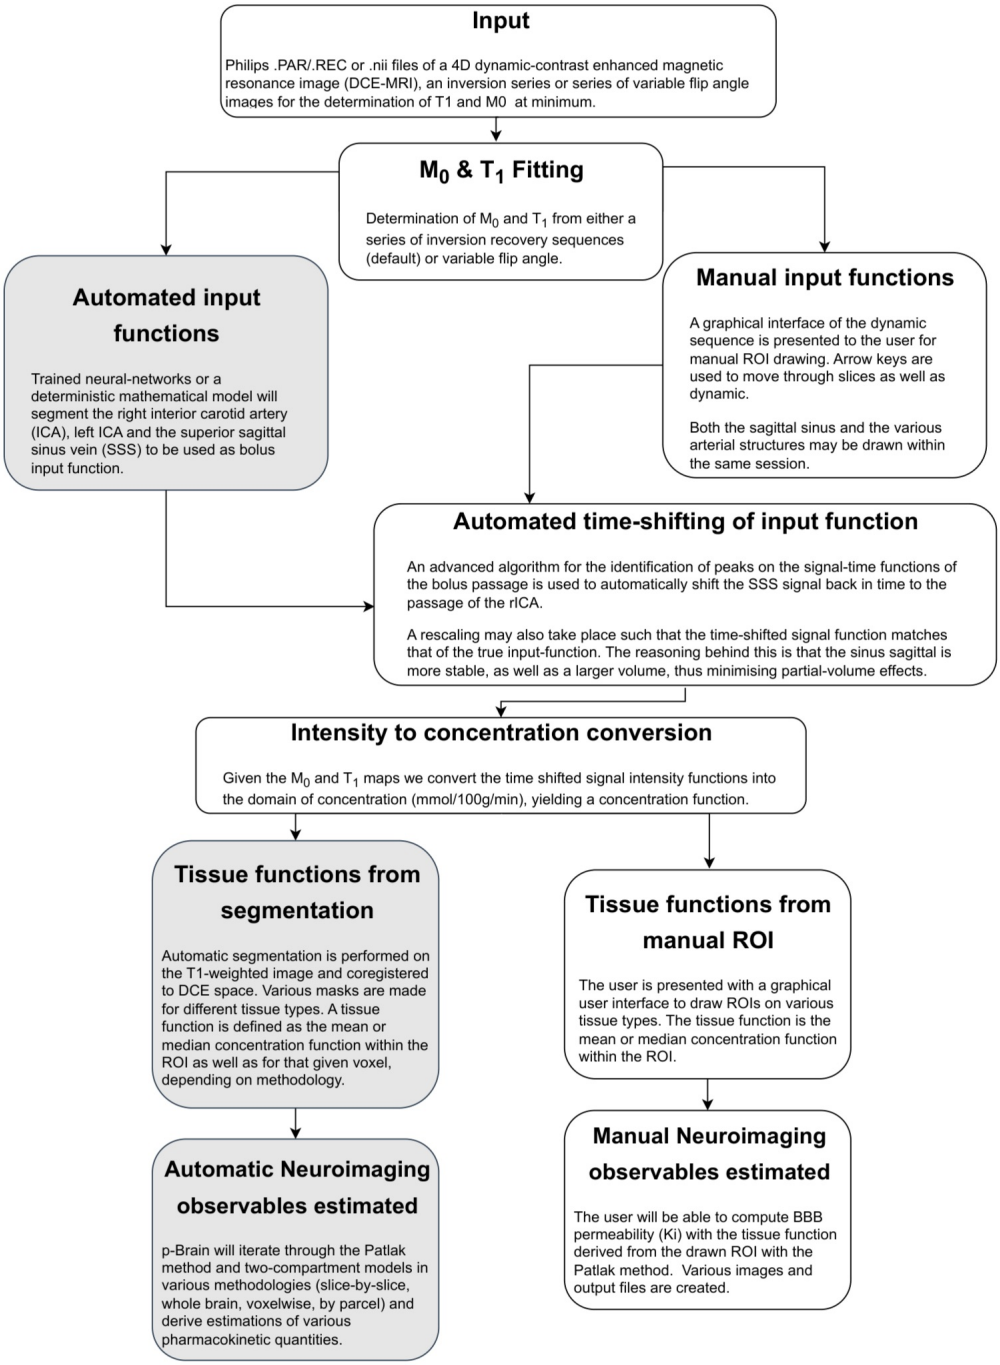

Figure 12: End-to-end p-Brain workflow. From T_1/M_0 fitting and automated AIF/VOF extraction to signal–to–concentration conversion, tissue-curve aggregation (manual or segmentation-based), and pharmacokinetic modeling (Patlak; optional deconvolution), yielding voxelwise, regional, slice-wise, and whole-brain outputs. The gray boxes denote fully automatic steps and white boxes indicate user-assisted stages.

The pipeline is organized into the following steps (see Figure 12):

1. **Inputs.** Required sequences are an inversion-recovery series for T_1/M_0 and a 4D DCE-MRI time series. A 3D T_1-weighted sequence will allow the segmentation module to report outputs per parcel or tissue.

2. **Preprocessing.** Standard clinical imaging inputs (e.g., DICOM; NIfTI) are ingested and aligned to a common geometry. Structural volumes are rigidly aligned to the DCE acquisition. Temporal sampling remains as acquired. Further spatial corrections and intensity normalization are omitted in the present build.

3. **T_1 and M_0 maps.** T_1 and M_0 are estimated voxelwise by nonlinear least squares. The implementation supports both inversion–recovery and saturation–recovery signal models; acquisition-specific timing parameters are described in the *Participants + acquisition* section.

4. **Input functions.** The internal carotid artery (default reference: rICA) and the superior sagittal sinus are identified by the CNN-based slice classifier and ROI segmentation models when the fully automatic workflow is selected. Manual delineation of vascular ROIs is also supported via the GUI when AI-based ROI extraction is not used. By default, modelling uses an SSS-derived effective input function (TSCC) constructed by aligning the venous curve to the selected ICA reference, with an optional rescaling to account for transit delay and dispersion; the ICA-derived arterial input function can be used directly if required.

5. **Tissue ROIs.** Tissue ROIs may be hand drawn within the presented user-interface. In the fully automatic mode, the segmentation module, by default FastSurfer (Henschel et al., 2020), runs an anatomical segmentation on the T_1-weighted image whereafter the segmentation module delineates cortical, subcortical and cerebellar gray matter, white matter, brainstem, and optional boundary masks defined as the boundary between subcortical grey matter and cortical white matter. An affine transformation aligns these masks to DCE space.

6. **Signal to concentration.** The pipeline converts signal–time curves into gadolinium concentration using the fitted T_1 and M_0 maps together with the acquisition timings from frames to seconds. Built-in safeguards reject invalid logarithmic terms.

7. **Modeling.** Patlak analysis yields K_i and v_p, both available from manual analysis. Model-free Tikhonov deconvolution provides CBF, mean transit time (MTT), and capillary transit-time heterogeneity (CTH). The framework is designed to be extensible, enabling additional kinetic models to be incorporated as needed.

8. **Outputs.** The pipeline produces quantitative outputs at multiple scales, including voxelwise maps, regional metrics derived from anatomical parcellations, slice-wise summaries, and whole-brain summaries. Outputs are exported in standard research-friendly formats for direct use in statistical workflows. Parcellated summaries for CBF, MTT, and CTH are produced analogously to K_i, using FastSurfer anatomical labels propagated into DCE space.

9. **Quality control and auditability.** Each stage emits reviewable artifacts (figures, masks, maps) and a structured quality control (QC) report that checks the presence and basic validity of its expected outputs (e.g., non-empty ROI masks, segmentation outputs, modelling summaries). QC is stage-scoped and can optionally be enforced to stop a run when a critical output is missing.

Each component of the workflow is detailed below in compact form, with full mathematical derivations and implementation settings provided in the Supplementary Methods.

4.3 T_1 and M_0 estimation

Voxelwise T_1 and M_0 were estimated from the inversion–recovery (IR) dataset by nonlinear least squares (Branch et al., 1999; Coleman and Li, 1996). For compatible protocols, *p*-Brain can alternatively estimate baseline T_1/M_0 from a variable flip-angle (VFA) spoiled GRE/SPGR series (multiple flip angles), using

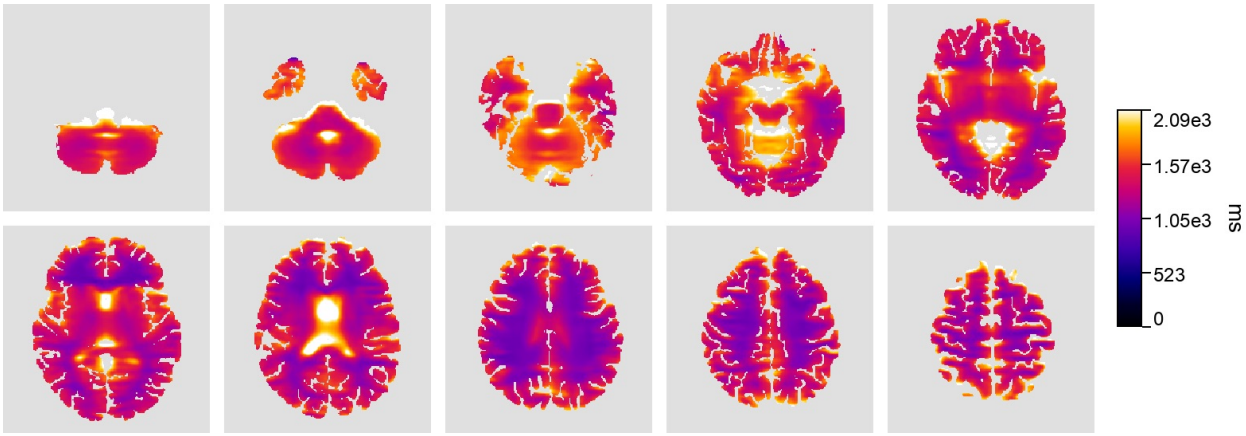

Figure 13: Voxelwise T_1 (ms) values from a representative ten-slice inversion-recovery dataset.

per-series repetition time and flip-angle values from the acquisition protocol; this option was not used for the data analysed for the presented data however.

For inversion times $T_I = \{T_{I,1}, \ldots, T_{I,N}\}$ and observed voxel signals y_i, the expected signal was expressed as

$$\mathscr{F}(T_I; A, B, T_1) = A - B\, e^{-T_I/T_1}, \tag{1}$$

a form corresponding to the standard inversion–recovery signal model

$$S(T_I) = \kappa M_0 \left(1 - 2\eta\, e^{-T_I/T_1}\right), \tag{2}$$

where η is the inversion efficiency (typically ≤ 1) and κ is a receive–gain scale. In *p*-Brain we fit Eq. 1 and do not estimate η or κ separately; their effects are absorbed into A and B (with $A = \kappa M_0$ and $B = 2\eta\kappa M_0$). This model describes recovery from inverted magnetisation ($-M_0$) to equilibrium. Inversion times for the dataset analysed here are given in *Participants and MRI acquisition*.

Parameter estimation proceeded voxelwise by minimising the residual energy

$$\mathscr{J}(A, B, T_1) = \sum_{i=1}^{N} \left| y_i - \mathscr{F}(T_{I,i}; A, B, T_1) \right|^2, \tag{3}$$

subject to $T_1 \in [100, 6000]$ ms. In the released implementation, A and B are constrained to be positive; the signed IR signal can cross zero around the null point ($T_I \approx T_1 \ln 2$) when $B > A$. This ensures unbiased T_1 estimates across tissues where the acquisition samples both sides of the inversion null. Initial guesses were $A = \max_i y_i$, $B = A/2$, and $T_1 = 750$ ms; voxels with negligible signal were excluded.

The fitted maps $\hat{T}_1$ and $\hat{A}$ retain the geometry of the original inversion–recovery images (Figs. 13 and 14). An M_0 proxy can be recovered as $\hat{M}_0 = \hat{A}/\kappa$ when receive calibration is available. These maps provide baseline relaxation parameters required for quantitative DCE-MRI analysis, enabling conversion of signal intensity to gadolinium concentration.

In the current released implementation, voxelwise gadolinium concentration curves are derived using the fitted T_1 and M_0 maps via the saturation-recovery-based conversion (see *Signal-to-concentration conversion*).

4.4 Vascular and tissue ROI extraction

Regions of interest (ROIs) support two core tasks in the pipeline: (i) deriving vascular input functions (arterial and venous) and (ii) aggregating tissue concentration–time curves for downstream modeling.

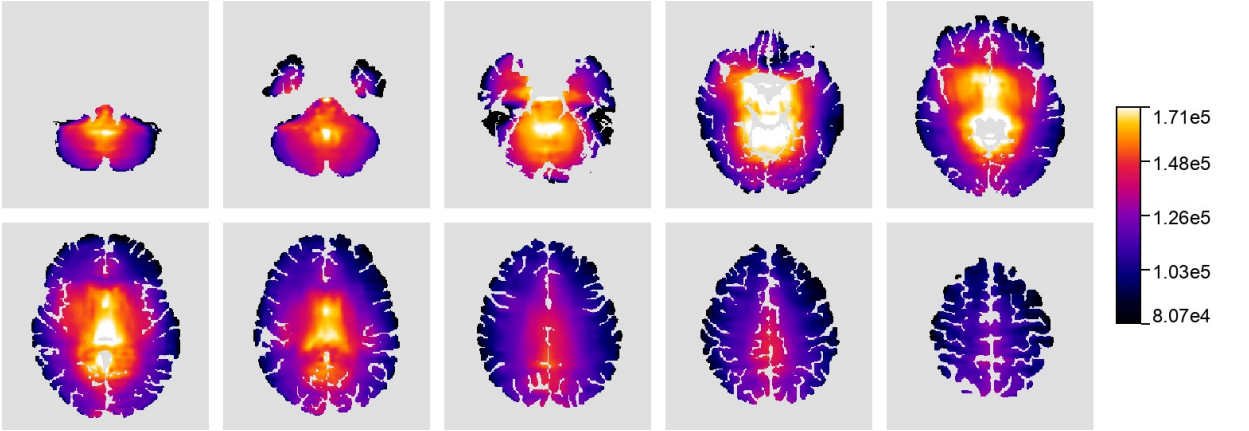

Figure 14: Voxelwise M_0 magnetisation values from a representative ten-slice inversion-recovery dataset.

For the arterial input function (AIF), we extract the signal from an internal carotid artery (default: rICA; lICA can be selected if required, e.g., if the rICA curve is degraded), which provides a robust measure of contrast arrival to the brain. The superior sagittal sinus (SSS) serves as the venous output function (VOF), reflecting contrast washout. As described in Workflow Step 4, the venous curve is aligned to the arterial peak via cross-correlation and amplitude rescaling to compensate for transit delays and dispersion (Cramer et al., 2023b; Manning et al., 2021; Thrippleton et al., 2019). In the fully automated implementation, we employ a two-stage CNN-based approach:

1. **Slice Classification CNN:** Identifies MRI slices containing relevant vascular and tissue structures based on intensity profiles and anatomical context.

2. **ROI Segmentation CNN:** Precisely delineates ROIs around the ICA and SSS for reliable extraction of input functions (for lICA selection, the ICA model can be applied on left–right mirrored slices and the predicted mask mirrored back).

Training used manually labelled examples of the ICA and SSS. The ROI segmentation networks are 2D U-Nets (Ronneberger et al., 2015); training and evaluation details are provided in the Supplementary Methods. Trained model weights are distributed with the software and are also available on Zenodo (Tireli, 2025).

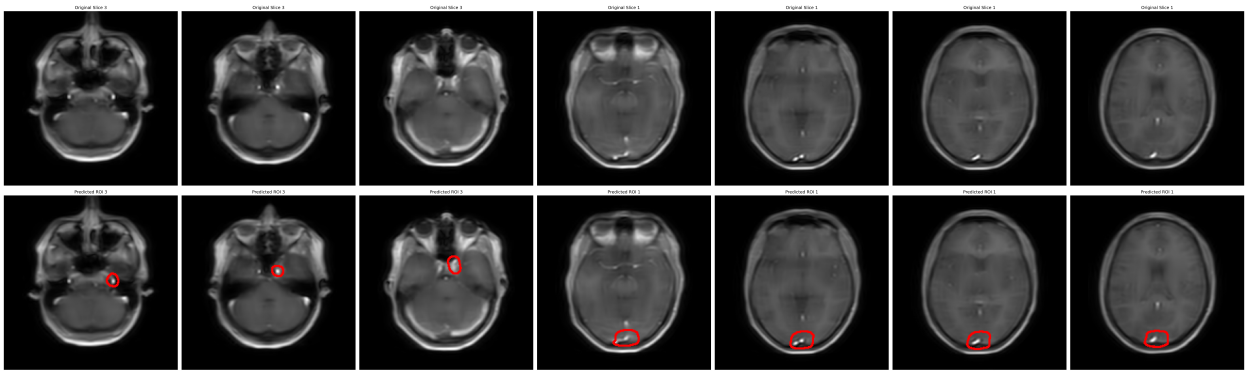

Figure 15: CNN slice selection and ROI delineation for the right internal carotid artery and superior sagittal sinus.

Manual ROI definition is supported via the GUI (hand-drawn ROIs with user-selected slices) when AI-based ROI extraction is not used or when additional review is required.

For automated tissue ROI extraction, we employ anatomical segmentation using FastSurfer (Henschel et al., 2020) and FSL (Jenkinson et al., 2012), which generate masks for anatomical parcels directly from segmentation and further *p*-Brain aggregates cortical and subcortical gray matter, white matter, cerebellum, and brainstem through fully automated processing.

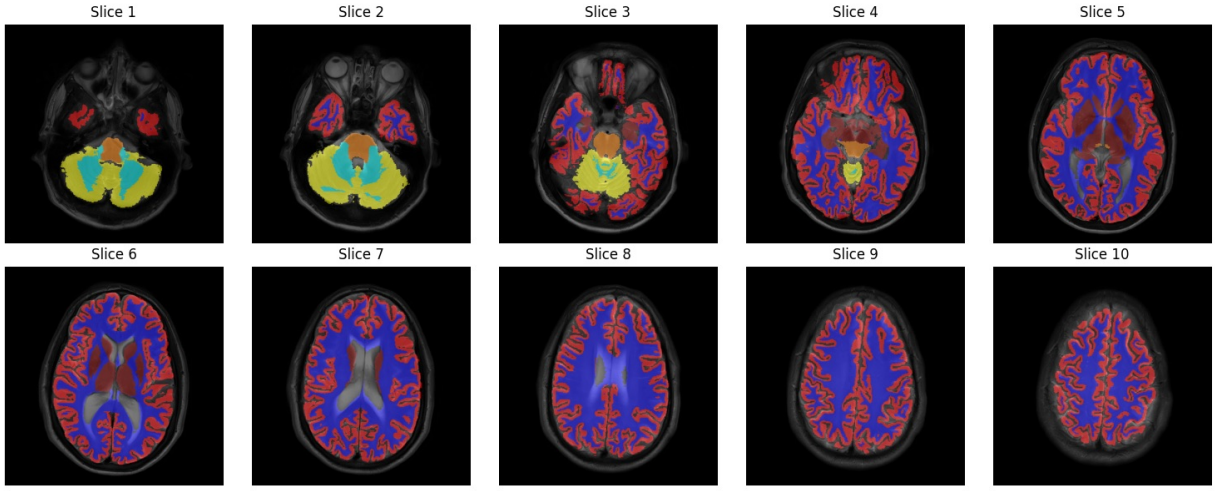

Figure 16: Anatomical segmentation overlaid on an axial T1-weighted image. Cortical gray matter (red), subcortical gray matter (dark red), white matter (blue), brainstem (orange), cerebellar white matter (light blue), and cerebellar gray matter (yellow).

A boundary mask between gray and white matter is computed by overlapping dilated masks and eroding the intersection to a few-voxel thickness[2]. All segmentation labels are then propagated into DCE-MRI space using affine registration, ensuring reliable alignment despite patient motion or anatomical variability. These propagated masks provide the foundation for accurate voxelwise, regional, and parcellated pharmacokinetic analyses.

4.5 Signal–to–concentration conversion

DCE-MRI yields a four-dimensional dataset in which each voxel has a signal–time curve $S(t)$. After defining arterial, venous, and tissue ROIs, we extract representative curves for modeling. For vascular ROIs, the default is to select the single brightest voxel within the mask (maximum over time) to reduce partial-volume contamination; alternatively, users can average the signal across ROI voxels. When using the max-voxel strategy, *p*-Brain also supports an *adaptive-max* variant in which the brightest voxel is re-selected independently at each time frame within the ROI to reduce sensitivity to in-plane motion. For tissue ROIs, we average voxelwise signals within the mask to reduce noise and improve stability.

[2]The boundary is derived as the erosion of the overlap between the dilated gray- and white-matter masks.

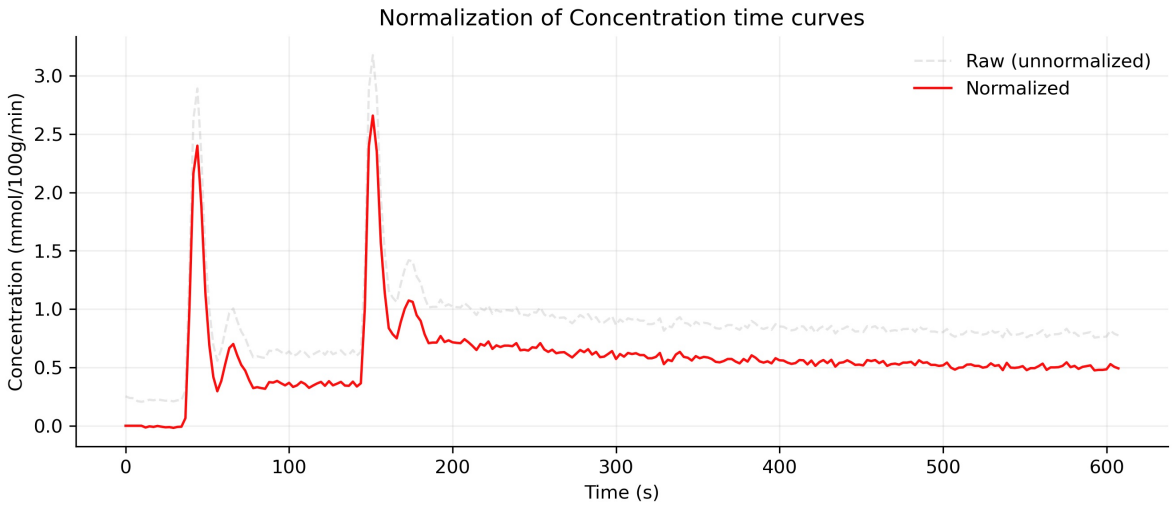

Figure 17: Example tissue concentration–time curve before and after the optional baseline/robust normalization used for diagnostics and visualization. Baseline frames are set to zero after subtracting the baseline mean, and the remaining curve is rescaled by a robust percentile to emphasise shape.

We convert signal–time curves $S(t)$ to gadolinium concentration $\mathscr{C}(t)$ using the voxelwise T_1 and M_0 maps via the saturation-recovery conversion used in the released p-Brain implementation:

$$\mathscr{C}(t) = -\frac{1}{r_1}\left[\frac{1}{T_D}\ln\left(1 - \frac{S(t)}{M_0\sin(\alpha)}\right) + \frac{1}{T_1}\right], \tag{4}$$

with T_1 and T_D expressed in seconds, α denoting the α-pulse flip angle, prepulse-to-readout delay $T_D = 120\,\mathrm{ms}$ (here for saturation recovery), and relaxivity $r_1 = 4\,\mathrm{s}^{-1}\,\mathrm{mM}^{-1}$ (blood at $3\,\mathrm{T}$). Sequence timing and flip-angle parameters were taken from the acquisition protocol. For the DCE protocol analysed here, the nominal readout flip angle was $30°$. To ensure numerical robustness, we clamp the logarithm argument to $(0, 1)$.

For spoiled gradient-echo (SPGR) acquisitions with variable flip angles, p-Brain also supports a VFA-based conversion that first estimates $T_1(t)$ from the SPGR steady-state equation

$$S(t;\alpha) \;=\; M_0\,\frac{\big(1 - E_1(t)\big)\sin\alpha}{1 - E_1(t)\cos\alpha}, \qquad E_1(t) = \exp\left(-\frac{\mathrm{TR}}{T_1(t)}\right), \tag{5}$$

using two or more flip angles. A common linearisation writes

$$\frac{S(t;\alpha)}{\tan\alpha} \;=\; E_1(t)\,\frac{S(t;\alpha)}{\sin\alpha} + M_0\,\big(1 - E_1(t)\big), \tag{6}$$

so that $E_1(t)$ (slope) and $M_0(1 - E_1(t))$ (intercept) can be estimated by regression across flip angles at each time point, yielding $T_1(t) = -\mathrm{TR}/\ln E_1(t)$. Concentration then follows from the relaxivity relation $1/T_1(t) = 1/T_{1,0} + r_1\,\mathscr{C}(t)$, i.e.

$$\mathscr{C}(t) = \frac{1}{r_1}\left(\frac{1}{T_1(t)} - \frac{1}{T_{1,0}}\right). \tag{7}$$

This VFA/SPGR branch is provided for compatible protocols but was not used for the data analysed here.

4.5.1 Curve normalization and input-function selection

Baseline enforcement and curve scaling. After signal–to–concentration conversion, the pre-contrast concentration should be zero: gadolinium is absent before injection and any baseline fluctuations primarily reflect

noise and modelling/fit error. p-Brain therefore applies a lightweight baseline correction to concentration–time curves (arterial, venous, and tissue) by subtracting a pre-bolus baseline estimate and enforcing a zero baseline window.

Let c_i denote the sampled curve at time index i and let B be the number of baseline frames (default $B = 5$). We estimate the pre-bolus baseline as the mean over the initial frames,

$$\bar{c}_0 \;=\; \frac{1}{B}\sum_{i=0}^{B-1} c_i, \tag{8}$$

then baseline-correct the curve and enforce a zero baseline window,

$$c'_i \;=\; c_i - \bar{c}_0, \qquad c'_i \leftarrow 0 \text{ for } i < B. \tag{9}$$

To reduce sensitivity to outliers and inter-voxel amplitude differences, we rescale by a robust deviation scale (95th percentile),

$$s \;=\; P_{95}\big(|c'|\big), \qquad \tilde{c}_i \;=\; \frac{c'_i}{\max(s,\epsilon)}. \tag{10}$$

To optionally emphasise curve *shape* (timing and relative bolus features) over absolute amplitude in visualisations, we rescale by a robust percentile as above, yielding the behaviour illustrated in Figure 17. This rescaling can be disabled via configuration.

Input-function selection (modelling). For kinetic modelling, p-Brain uses an SSS-derived effective input function constructed by time-shifting the venous curve to a selected ICA reference (default: rICA; lICA can be selected if required, e.g., if the rICA curve is degraded) via cross-correlation peak alignment. An optional amplitude rescaling then produces a time-shifted (and optionally rescaled) concentration curve (TSCC). Rescaling can be disabled (time-shift only), or computed either by matching peak height (peak-based) or by matching the total curve "volume" (area-under-curve based). The ICA-derived arterial input function can also be used directly.

Input functions follow the automated extraction described in Workflow Step 4: the arterial curve is drawn from the right internal carotid artery and the venous curve from the superior sagittal sinus, with cross-correlation alignment and amplitude rescaling to compensate for transit delays and dispersion (Varatharaj et al., 2024b). We further apply adaptive signal tracking to mitigate in-plane motion when deriving the AIF.

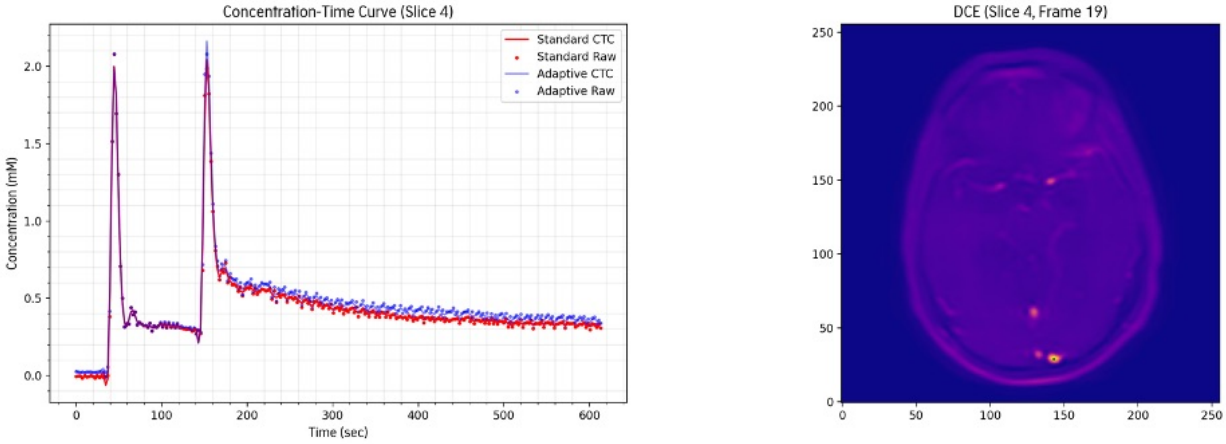

Figure 18: Time-shifted venous concentration–time curve for a double-bolus acquisition. The AIF is derived with adaptive signal tracking to correct minor in-plane motion. The concentration curves in this plot are passed through a Butterworth filter (Butterworth, 1930) for better interpretability.

4.6 Tracer–kinetic models and parameter estimation

4.6.1 Patlak graphical analysis

We quantify low-level BBB leakage with the Patlak model, which assumes negligible backflux from the extravascular–extracellular space (EES) to plasma ($k_2 \approx 0$) and yields a linear relationship between tissue and input concentrations (Patlak et al., 1983). Let $C_a(t)$ denote the selected input function (default: TSCC; optional: ICA-derived AIF) and $C_t(t)$ the tissue concentration. The model reads

$$C_t(t) = v_p\, C_a(t) + K_i \int_0^t C_a(\tau)\, d\tau, \tag{11}$$

with K_i the influx constant and v_p the fractional plasma volume. Dividing by $C_a(t)$ gives a linear form,

$$\frac{C_t(t)}{C_a(t)} = v_p + K_i \frac{\int_0^t C_a(\tau)\, d\tau}{C_a(t)}. \tag{12}$$

For discrete sampling at t_i, define

$$x_i = \frac{\sum_{j=0}^{i} C_{a,j}\, \Delta t_j}{C_{a,i}}, \qquad y_i = \frac{C_{t,i}}{C_{a,i}}. \tag{13}$$

We fit $y_i = v_p + K_i x_i$ by ordinary least squares over the late (approximately linear) epoch of the Patlak plot; in the released implementation this is selected as the upper fraction of x_i (default: $x_i \in [\frac{1}{3}x_{\max}, x_{\max}]$) after excluding invalid samples (e.g., $C_a(t_i) = 0$ or NaNs), yielding

$$K_i = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sum_i (x_i - \bar{x})^2}, \qquad v_p = \bar{y} - K_i \bar{x}. \tag{14}$$

We report uncertainty as

$$\mathrm{SD}_{K_i} = \sqrt{\frac{\sum_i \left(y_i - (v_p + K_i x_i)\right)^2}{(n-2)\sum_i (x_i - \bar{x})^2}}. \tag{15}$$

In the released p-Brain implementation, the cumulative integral in x_i is evaluated with a left Riemann sum on the acquired grid (with $x_0 = 0$), i.e.

$$x_i = \frac{\sum_{j=0}^{i-1} C_{a,j}\, \Delta t_j}{C_{a,i}}. \tag{16}$$

The fitted slope and intercept follow legacy reporting conventions: K_i is converted from s^{-1} to mL/100g/min by multiplying by 6000, and v_p is converted from a unitless fraction to mL/100g by multiplying by 100.

4.6.2 Model-free deconvolution and microvascular transit metrics

We model the tissue concentration–time curve $C_t(t)$ as a convolution of the arterial input function $C_a(t)$ with a causal impulse response $g(t) = F\, R(t)$, where F is cerebral blood flow and $R(t)$ is the residue function with $R(0) = 1$:

$$C_t(t) \;=\; \int_0^t C_a(\tau)\, g(t-\tau)\, d\tau \;=\; F \int_0^t C_a(\tau)\, R(t-\tau)\, d\tau.$$

Sampling on a uniform grid t_i with step Δt, recovering g from discrete measurements is ill-posed. In the released p-Brain implementation, we represent the impulse response in a cubic B-spline basis and solve for basis coefficients with Tikhonov regularisation. Let B be the B-spline basis matrix evaluated on the acquired grid, so that $\mathbf{g} \approx B^\top \mathbf{x}/\Delta t$. With A the lower–triangular Toeplitz convolution matrix built from the (optionally time-shifted) input function, define $D = AB^\top$ and solve

$$\hat{\mathbf{x}} \;=\; \arg\min_{\mathbf{x}} \left\|D\mathbf{x} - \mathbf{C_t}\right\|_2^2 \;+\; \lambda^2 \left\|L\mathbf{x}\right\|_2^2,$$

where L is a first-difference operator on the coefficient vector and $\lambda > 0$ the regularisation weight. Candidate λ values are evaluated on a fixed grid (default: 121 values spanning $[0.05, \sigma_{\max}(A_0)]$, where A_0 is the Toeplitz matrix built from the unshifted input and $\sigma_{\max}$ its largest singular value). For each voxel, λ is selected by maximising the curvature of the L-curve in $\log \|D\hat{\mathbf{x}} - \mathbf{C_t}\|$ versus $\log \|L\hat{\mathbf{x}}\|$ space. Optional per-voxel temporal offsets are handled by shifting $C_a(t)$ within the forward operator via a shape-preserving cubic interpolation; offsets can be quantised into small groups for efficiency.

Flow, blood volume, and transit metrics. After estimating $\hat{\mathbf{x}}$, we reconstruct the impulse response $\hat{\mathbf{g}}$ and estimate flow as the early peak of the impulse response,

$$\hat{F} = \max_{0 \le i < w} \hat{g}(t_i),$$

with a fixed peak window w (default: 50 samples). Cerebral blood flow is then reported in mL/100g/min using a density scaling (default tissue density $\rho = 1.04\,\mathrm{g\,mL^{-1}}$); when the input function is plasma-derived, an additional $(1 - \mathrm{Hct})$ factor is applied.
The released implementation reports an apparent distribution volume

$$v_d = \frac{\int C_t(t)\,dt}{\int C_a(t)\,dt},$$

where the denominator uses the time-shifted input when offsets are enabled (numerical integration by the trapezoidal rule). Mean transit time is then computed as

$$\mathrm{MTT} = \frac{v_d}{\hat{F}},$$

which is consistent with $\int_0^\infty R(t)\,dt$ under the convolution model. The central volume relationship $\mathrm{CBV} = \hat{F}\,\mathrm{MTT}$ therefore holds by construction.

CTH from the constrained residue. For microvascular dispersion, we form a normalised residue $R(t) = \hat{g}(t)/\hat{F}$ and apply stabilising shape constraints prior to differentiation: the residue is peak-aligned (so the impulse peak is at the start), renormalised to satisfy $R(0) = 1$, clipped to $[0, 1]$, enforced to be non-increasing, and forced to taper smoothly to zero at the tail. We then compute

$$h(t) = \frac{-R'(t)}{\int_0^\infty \big(-R'(u)\big)\,du}, \qquad -R'(t) \leftarrow \max\{-R'(t), 0\},$$

using finite differences on the acquired grid and trapezoidal-rule normalisation on midpoints. Finally,

$$\mathrm{CTH} = \sqrt{\int_0^\infty (t-\mu)^2 h(t)\,dt}, \qquad \mu = \int_0^\infty t\,h(t)\,dt.$$

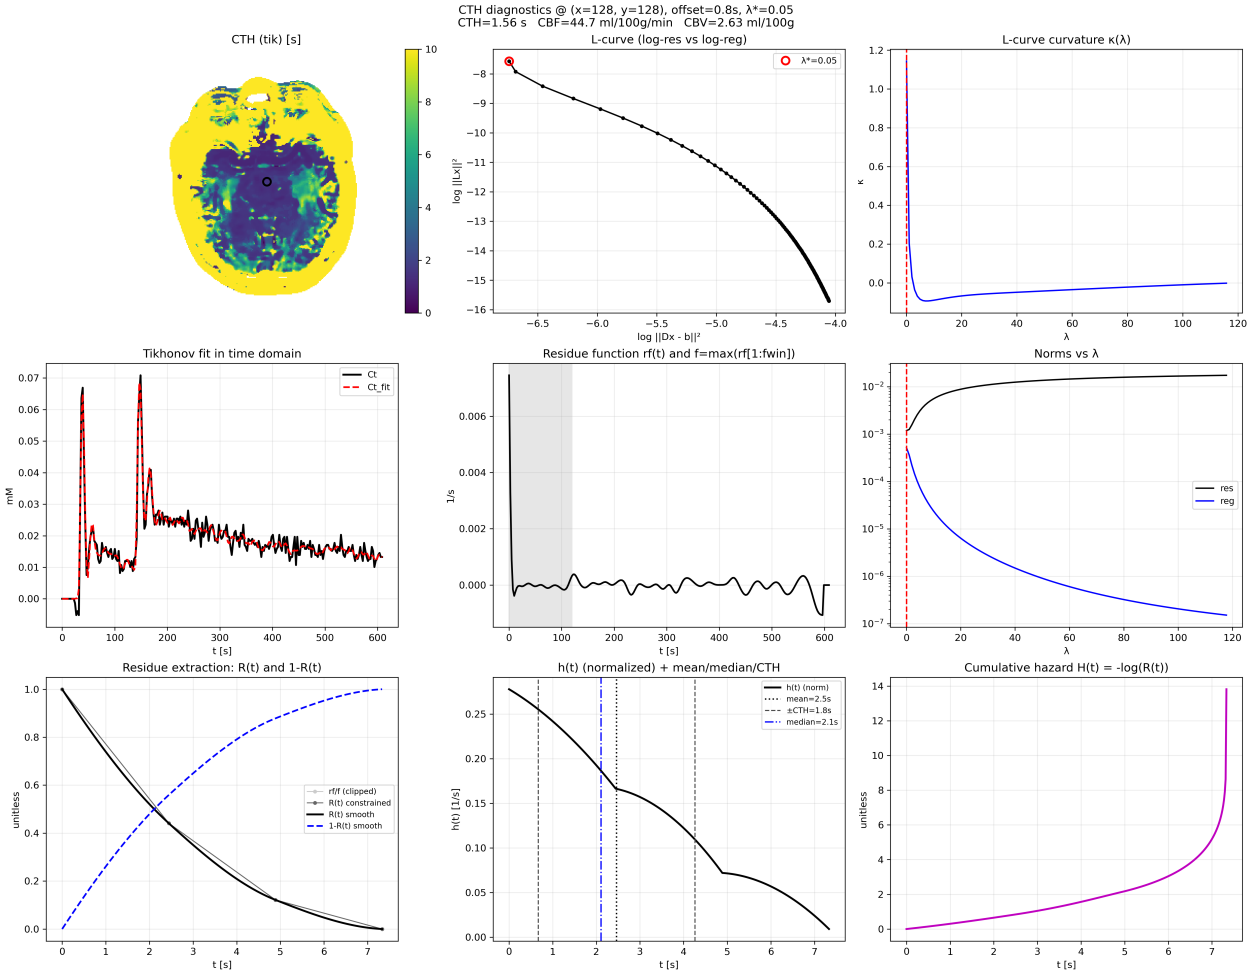

Figure 19: Example deconvolution diagnostics from the released p-Brain Tikhonov implementation for a representative voxel. The panel shows the L-curve (log residual norm vs log regularisation norm), curvature-based selection of the optimal λ, the resulting time-domain fit, the reconstructed impulse response and constrained residue, the derived outflow distribution $h(t)$, and the corresponding MTT/CTH summaries.

4.7 Quality control and end-to-end auditability

p-Brain is designed to produce standardized outputs while remaining auditable in batch operation. In addition to visual summaries (e.g., per-slice montage panels of masks, curves, and fits), the pipeline runs a lightweight quality-control (QC) layer after key stages. QC emits one structured report per stage with an overall status (pass, warning, or fail), together with brief messages that point to the relevant outputs for review. Optionally, QC can be enforced to stop a run when a stage reports failure.

The current QC checks are intentionally conservative: they verify that each stage produced its expected outputs and that key masks/curves are non-empty, rather than attempting to judge physiological correctness. Specifically, the following stage-level checks are performed:

- T_1/M_0 **fitting.** Checks that baseline relaxation mapping completed successfully and produced the expected parameter maps. Missing outputs yield fail.
- **Input functions.** Checks that arterial/venous ROIs were produced and are non-empty (i.e., contain voxels) so that AIF/VOF curves can be extracted. Unreadable outputs yield warning; missing or universally empty ROIs yield fail.
- **Time shifting.** Checks that the time-shift/alignment step produced readable selection metadata

for the chosen ICA reference and (when configured) the TSCC alignment used for modelling. When modelling uses the arterial curve directly, missing TSCC metadata is downgraded to warning.

- **Segmentation.** Checks that the anatomical segmentation completed and produced the expected tissue/parcellation outputs for downstream aggregation. Missing segmentation yields fail.
- **Tissue curve export.** Checks that tissue concentration–time curves were exported for the configured tissue ROIs. Missing outputs yield fail.
- **Modelling.** Checks that model summaries were written and that the requested modelling outputs (summaries and, when enabled, maps) were produced. Missing summaries yield fail; missing optional maps yield warning.

In practice, QC reports complement the visual end-to-end montage (Figure 10) by highlighting runs that warrant closer review.

5 Conclusion

Our fully automated platform provides a standardized, reproducible framework for generating quantitative metrics of cerebral perfusion and permeability. Here we validate technical agreement of Patlak-derived K_i estimates in a uniform dataset; generalization across scanners, sites, pathologies, and clinical workflows remains to be established.

6 Code and Data Availability

The pipeline is freely available via the *p*-Brain repository on GitHub, complete with detailed installation instructions and thorough documentation. A macOS desktop application for local project organisation and job execution (no cloud upload, data remains on-device) is available at `github.com/edtireli/p-brain-platform`. Trained CNN weights for slice selection and ROI segmentation are hosted at Zenodo (Tireli, 2025).

7 Acknowledgements

The authors thank Antonis Asiminas for valuable discussions and encouragement during the preparation of this work.

Appendix

Voxelwise validation against a reference MATLAB implementation

Additional voxelwise comparisons against the reference MATLAB program for perfusion map outputs and for the T_1/M_0 fitting outputs (representative slice 5).

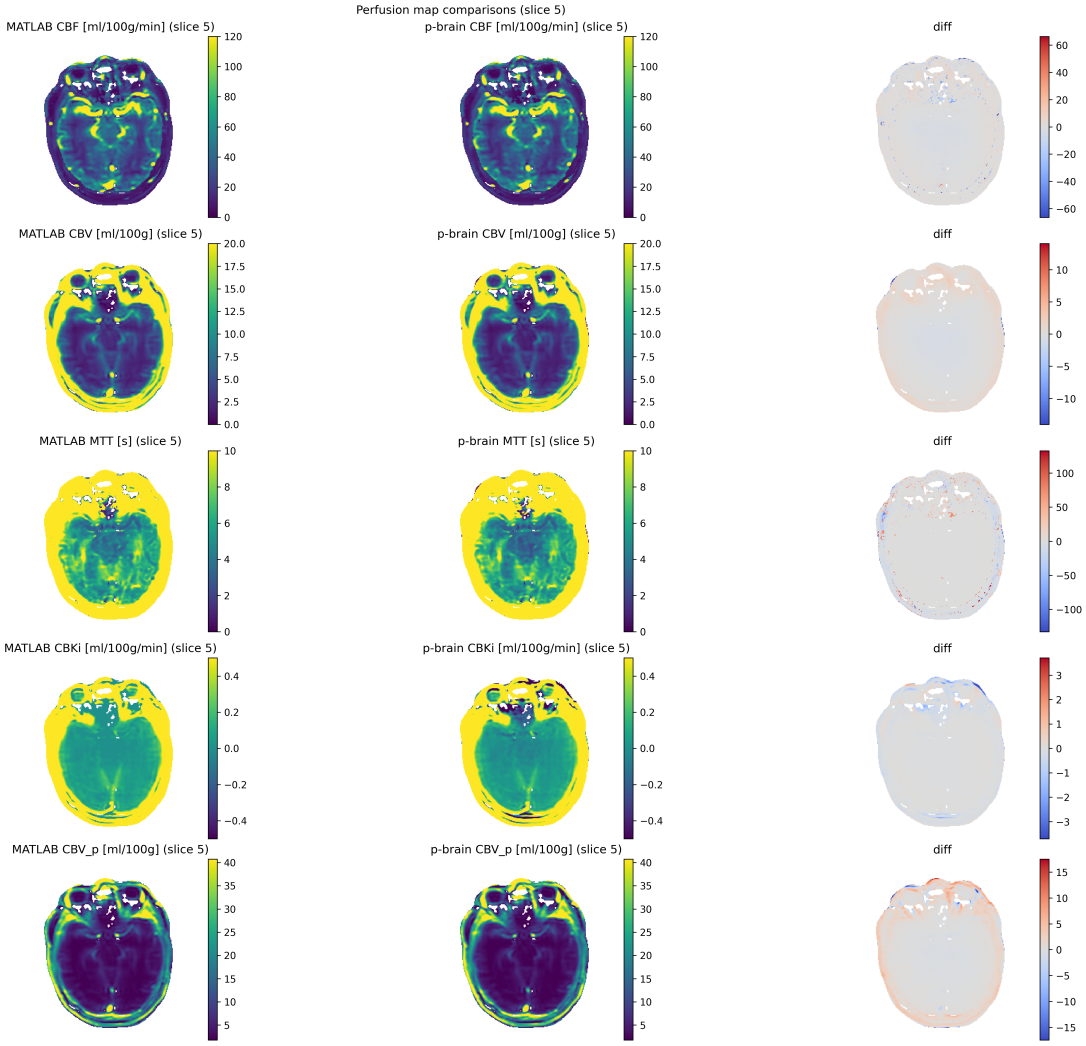

Figure 20: Voxelwise comparison of perfusion map outputs for a representative slice (slice 5): *p*-Brain (fully automatic) versus the reference MATLAB program (manual).

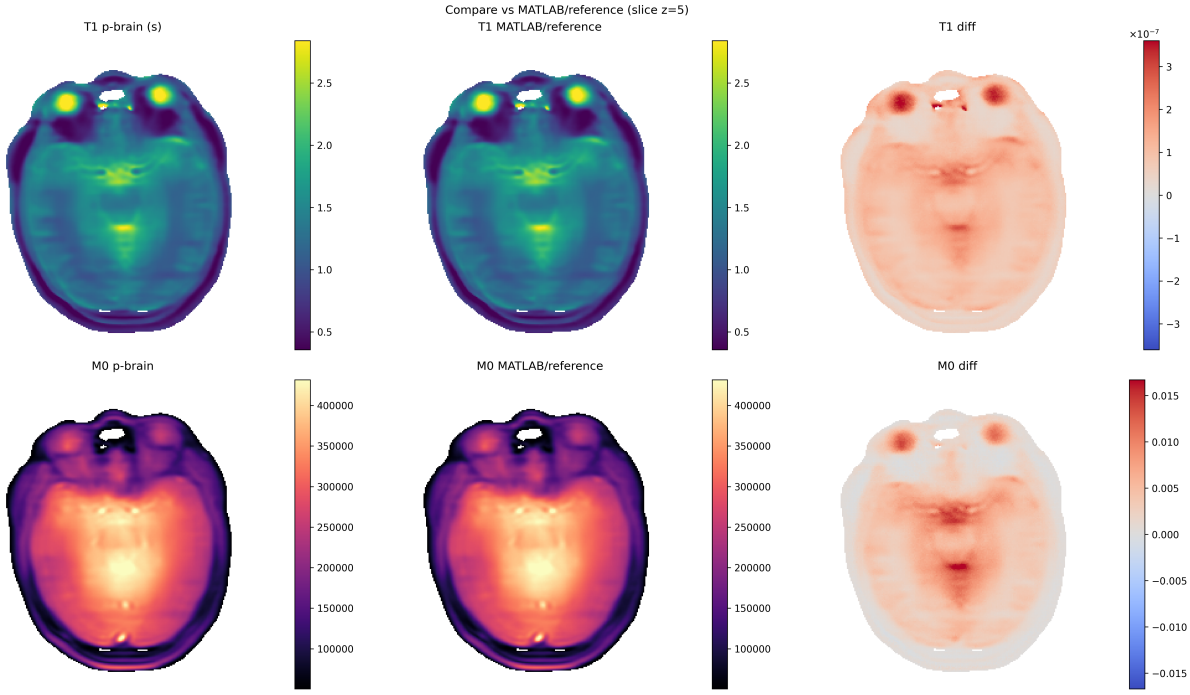

Figure 21: Voxelwise comparison of T_1 and M_0 fitting outputs for a representative slice (slice 5): p-Brain (fully automatic) versus the reference MATLAB program (manual).

Additional cohort-level atlas projections

Additional cohort-mean atlas projections corresponding to the Results section on cohort-level atlas projection.

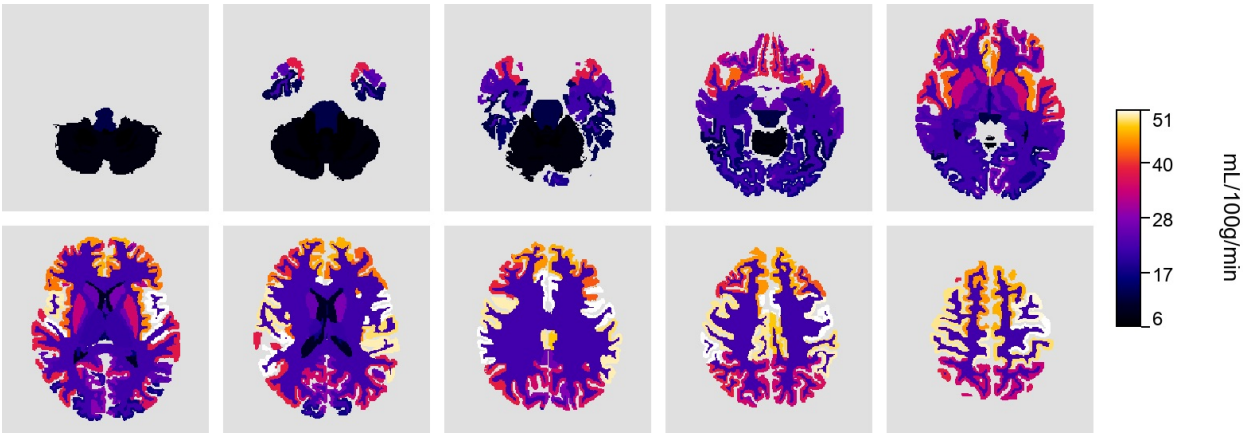

Figure 22: Cohort-mean CBF from model-free deconvolution. Units: mL/100g/min.

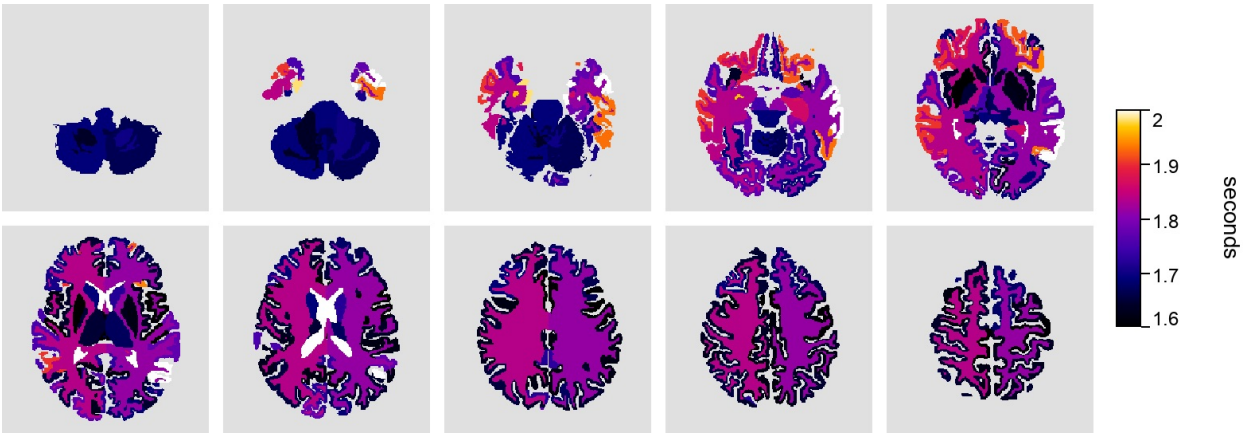

Figure 23: Cohort-mean capillary transit-time heterogeneity (CTH). Units: s.

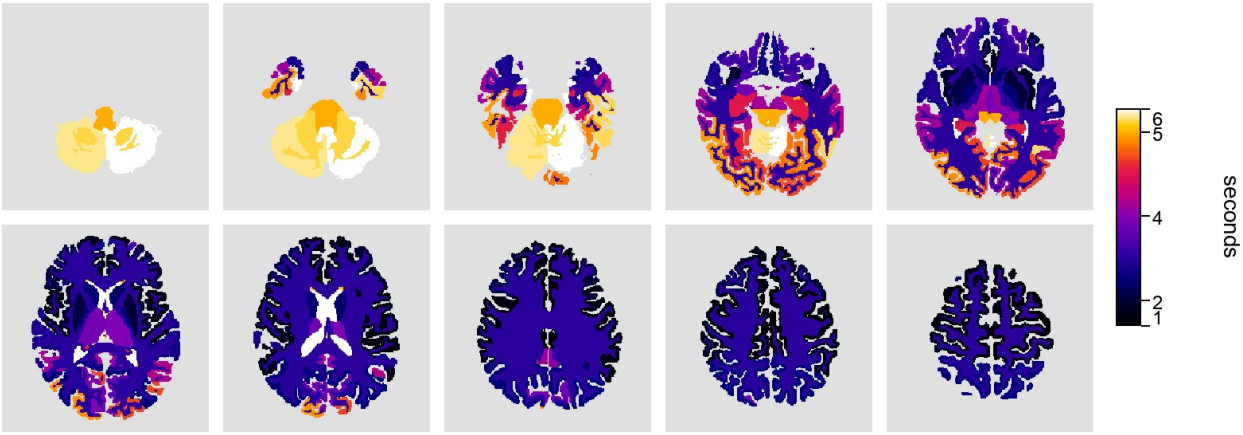

Figure 24: Cohort-mean mean transit time (MTT). Units: s.

Additional parcel-level maps

Additional parcel-level outputs corresponding to the Results section on regional and parcellated organization.

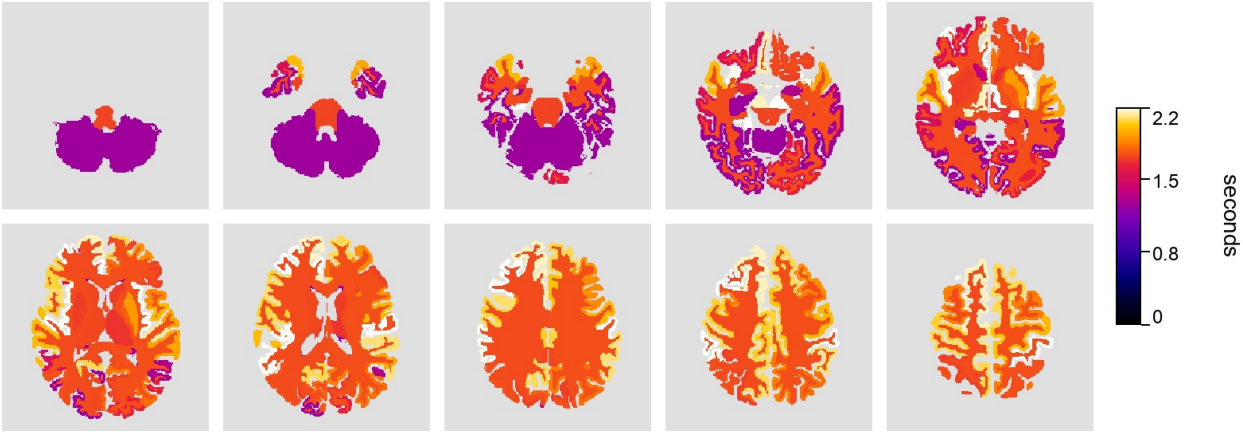

Figure 25: Parcel-level CTH in the representative participant. Units: s.

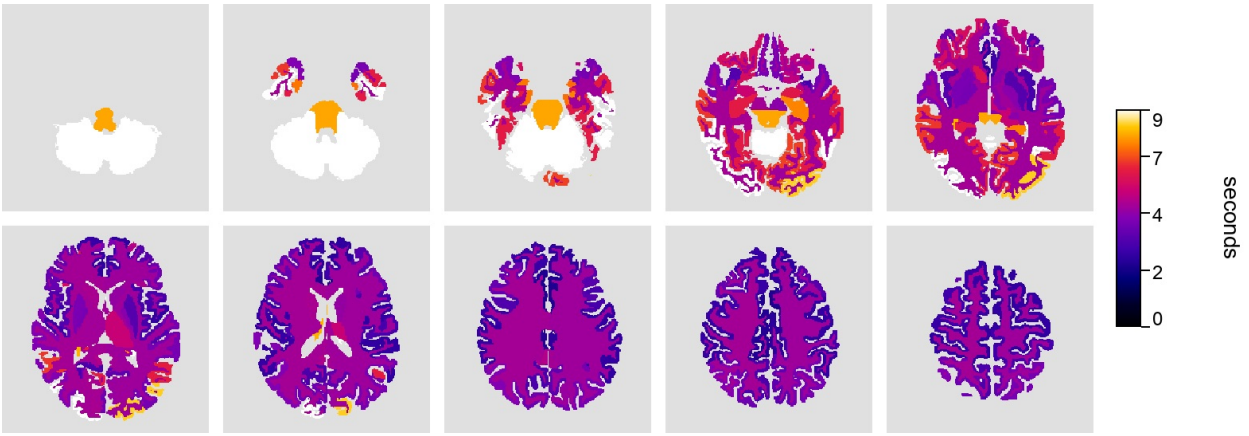

Figure 26: Parcel-level MTT in the representative participant. Units: s.

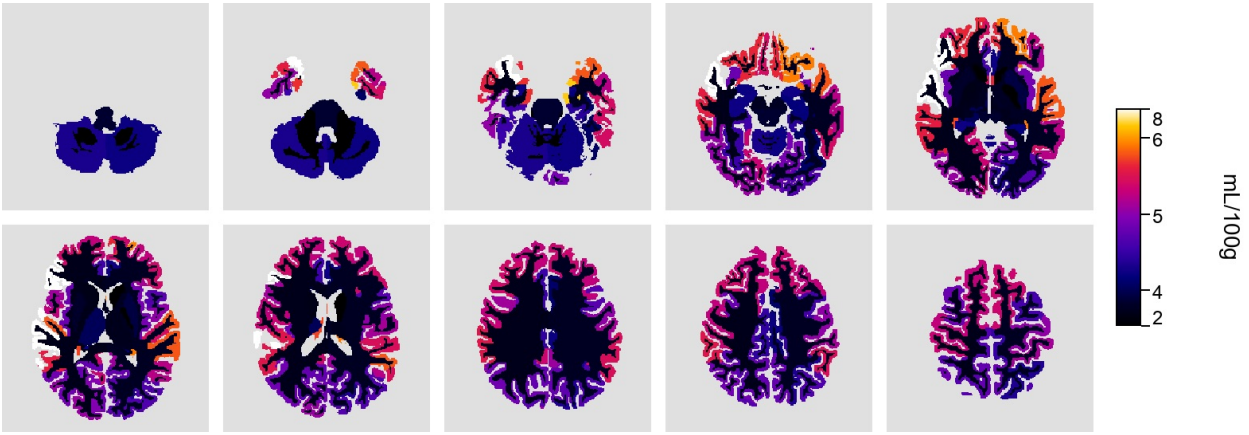

Figure 27: Parcel-level v_p in the representative participant.

References

S.P. Cramer, H. Simonsen, J.L. Frederiksen, E. Rostrup, and H.B.W. Larsson. Abnormal bloodbrain barrier permeability in normal appearing white matter in multiple sclerosis investigated by mri. *NeuroImage: Clinical*, 4:182189, 2014. ISSN 2213-1582. doi: 10.1016/j.nicl.2013.12.001. URL http://dx.doi.org/10.1016/j.nicl.2013.12.001.

Sachi Okuchi, Antonio RojasGarcia, Agne Ulyte, Ingeborg Lopez, Jurgita Usinskiene, Martin Lewis, Sara M Hassanein, Eser Sanverdi, Xavier Golay, Stefanie Thust, Jasmina PanovskaGriffiths, and Sotirios Bisdas. Diagnostic accuracy of dynamic contrastenhanced perfusion mri in stratifying gliomas: A systematic review and metaanalysis. *Cancer Medicine*, 8(12):55645573, August 2019. ISSN 2045-7634. doi: 10.1002/cam4.2369. URL http://dx.doi.org/10.1002/cam4.2369.

Anna K. Heye, Ross D. Culling, Maria del C. Valdes Hernandez, Michael J. Thrippleton, and Joanna M. Wardlaw. Assessment of bloodbrain barrier disruption using dynamic contrast-enhanced mri. a systematic review. *NeuroImage: Clinical*, 6:262274, 2014. ISSN 2213-1582. doi: 10.1016/j.nicl.2014.09.002. URL http://dx.doi.org/10.1016/j.nicl.2014.09.002.

Audrey Chagnot, Samuel R. Barnes, and Axel Montagne. Magnetic resonance imaging of bloodbrain barrier permeability in dementia. *Neuroscience*, 474:1429, October 2021. ISSN 0306-4522. doi: 10.1016/j.neuroscience.2021.08.003. URL http://dx.doi.org/10.1016/j.neuroscience.2021.08.003.

Rajikha Raja, Gary A. Rosenberg, and Arvind Caprihan. Mri measurements of blood-brain barrier function in dementia: A review of recent studies. *Neuropharmacology*, 134:259271, May 2018. ISSN 0028-3908. doi: 10.1016/j.neuropharm.2017.10.034. URL http://dx.doi.org/10.1016/j.neuropharm.2017.10.034.

Aravinthan Varatharaj, Maria Liljeroth, Angela Darekar, Henrik B. W. Larsson, Ian Galea, and Stig P. Cramer. Blood–brain barrier permeability measured using dynamic contrast-enhanced mri: a validation study. *Journal of Physiology*, 597(3):699–709, 2019. doi: 10.1113/JP276887.

Michael Ingrisch, Steven Sourbron, Dominik Morhard, Birgit Ertl-Wagner, Tania Kümpfel, Reinhard Hohlfeld, Maximilian Reiser, and Christian Glaser. Quantification of perfusion and permeability in multiple sclerosis: dynamic contrast-enhanced mri in 3d at 3t. *Investigative Radiology*, 47(4):252–258, 2012. doi: 10.1097/RLI.0b013e31823bfc97.

C. S. Patlak, R. G. Blasberg, and J. D. Fenstermacher. Graphical evaluation of blood-to-brain transfer constants from multiple-time uptake data. *Journal of Cerebral Blood Flow and Metabolism*, 3(1):1–7, 1983. doi: 10.1038/jcbfm.1983.1.

Stig P. Cramer, Henrik B. W. Larsson, Maria H. Knudsen, Helle J. Simonsen, Mark B. Vestergaard, and Ulrich Lindberg. Reproducibility and optimal arterial input function selection in dynamic contrast-enhanced perfusion mri in the healthy brain. *Journal of Magnetic Resonance Imaging*, 57(4):1229–1240, 2023a. doi: 10.1002/jmri.28380.

Stig P. Cramer, Signe Modvig, Helle J. Simonsen, Jette L. Frederiksen, and Henrik B. W. Larsson. Permeability of the blood–brain barrier predicts conversion from optic neuritis to multiple sclerosis. *Brain*, 138(Pt 9): 2571–2583, 2015. doi: 10.1093/brain/awv203.

Kersten Villringer, Borja E. Sanz Cuesta, Ann-Christin Ostwaldt, Ulrike Grittner, Peter Brunecker, Ahmed A. Khalil, Kristina Schindler, Ole Eisenblatter, Heinrich Audebert, and Jochen B. Fiebach. Dce-mri bloodbrain barrier assessment in acute ischemic stroke. *Neurology*, 88(5):433440, January 2017. ISSN 1526-632X. doi: 10.1212/wnl.0000000000003566. URL http://dx.doi.org/10.1212/WNL.0000000000003566.

Francesco Arba, Chiara Rinaldi, Danilo Caimano, Federica Vit, Giorgio Busto, and Enrico Fainardi. Bloodbrain barrier disruption and hemorrhagic transformation in acute ischemic stroke: Systematic review and meta-analysis. *Frontiers in Neurology*, 11, January 2021. ISSN 1664-2295. doi: 10.3389/fneur.2020.594613. URL http://dx.doi.org/10.3389/fneur.2020.594613.

Jeffrey B. Ware, Saurabh Sinha, Justin Morrison, Alexa E. Walter, James J. Gugger, Andrea L.C. Schneider, Cian Dabrowski, Hannah Zamore, Leroy Wesley, Brigid Magdamo, Dmitriy Petrov, Junghoon J. Kim, Ramon Diaz-Arrastia, and Danielle K. Sandsmark. Dynamic contrast enhanced mri for characterization of blood-brain-barrier dysfunction after traumatic brain injury. *NeuroImage: Clinical*, 36:103236, 2022. ISSN 2213-1582. doi: 10.1016/j.nicl.2022.103236. URL http://dx.doi.org/10.1016/j.nicl.2022.103236.

Huiping Li, Xian Liu, Ruihong Wang, Aili Lu, Zhaohui Ma, Shibiao Wu, Hongji Lu, Yaming Du, Kan Deng, Lixin Wang, and Fang Yuan. Bloodbrain barrier damage and new onset refractory status epilepticus: An exploratory study using dynamic contrastenhanced magnetic resonance imaging. *Epilepsia*, 64(6):15941604, March 2023. ISSN 1528-1167. doi: 10.1111/epi.17576. URL http://dx.doi.org/10.1111/epi.17576.

So-Hyun Ji, Roh-Eul Yoo, Seung Hong Choi, Woo Jin Lee, Soon Tae Lee, Young Hun Jeon, Kyu Sung Choi, Ji Ye Lee, Inpyeong Hwang, Koung Mi Kang, and Tae Jin Yun. Dynamic contrast-enhanced mri quantification of altered vascular permeability in autoimmune encephalitis. *Radiology*, 310(3), March 2024. ISSN 1527-1315. doi: 10.1148/radiol.230701. URL http://dx.doi.org/10.1148/radiol.230701.

J.M. Chi, M. Mackay, A. Hoang, K. Cheng, C. Aranow, J. Ivanidze, B. Volpe, B. Diamond, and P.C. Sanelli. Alterations in blood-brain barrier permeability in patients with systemic lupus erythematosus. *American Journal of Neuroradiology*, February 2019. ISSN 1936-959X. doi: 10.3174/ajnr.a5990. URL http://dx.doi.org/10.3174/ajnr.A5990.

Elisabeth S. Lindland, Anne Marit Solheim, Silje Andreassen, Robin Bugge, Randi Eikeland, Harald Reiso, Aslaug R. Lorentzen, Hanne F. Harbo, Mona K. Beyer, and Atle Bjrnerud. Dynamic contrast-enhanced mri shows altered bloodbrain barrier function of deep gray matter structures in neuroborreliosis: a casecontrol study. *European Radiology Experimental*, 7(1), September 2023. ISSN 2509-9280. doi: 10.1186/s41747-023-00365-6. URL http://dx.doi.org/10.1186/s41747-023-00365-6.

Joga Chaganti, Karthik Marripudi, Lukas P. Staub, Caroline D. Rae, Thomas M. Gates, Kirsten J. Moffat, and Bruce J. Brew. Imaging correlates of the bloodbrain barrier disruption in hiv-associated neurocognitive disorder and therapeutic implications. *AIDS*, 33(12):18431852, October 2019. ISSN 0269-9370. doi: 10.1097/qad.0000000000002300. URL http://dx.doi.org/10.1097/QAD.0000000000002300.

A. Varatharaj, C. Jacob, A. Darekar, B. Yuen, S. Cramer, H. Larsson, and I. Galea. Measurement variability of blood-brain barrier permeability using dynamic contrast-enhanced magnetic resonance imaging. *Imaging Neuroscience*, 2:1–16, oct 2024a. doi: 10.1162/imag_a_00324.

Mark J. Beresford, Anwar R. Padhani, N. Jane Taylor, Mei-Lin Ah-See, J. James Stirling, Andreas Makris, James A. d'Arcy, and David J. Collins. Inter- and intraobserver variability in the evaluation of dynamic breast cancer mri. *Journal of Magnetic Resonance Imaging*, 24(6):1316–1325, October 2006. ISSN 1522-2586. doi: 10.1002/jmri.20768. URL http://dx.doi.org/10.1002/jmri.20768.

Fernando Calamante. Arterial input function in perfusion mri: A comprehensive review. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 74:1–32, October 2013. ISSN 0079-6565. doi: 10.1016/j.pnmrs.2013.04.002. URL http://dx.doi.org/10.1016/j.pnmrs.2013.04.002.

Vera C. Keil, Burkhard Mädler, Jürgen Gieseke, Rolf Fimmers, Elke Hattingen, Hans H. Schild, and Dariusch R. Hadizadeh. Effects of arterial input function selection on kinetic parameters in brain dynamic contrast-enhanced mri. *Magnetic Resonance Imaging*, 40:83–90, July 2017. ISSN 0730-725X. doi: 10.1016/j.mri.2017.04.006. URL http://dx.doi.org/10.1016/j.mri.2017.04.006.

Samuel R. Barnes, Thomas S.C. Ng, Naomi Santa-Maria, Axel Montagne, Berislav V. Zlokovic, and Russell E. Jacobs. ROCKETSHIP: a flexible and modular software tool for the planning, processing and analysis of dynamic mri studies. *BMC Medical Imaging*, 15(1):19, 2015. doi: 10.1186/s12880-015-0062-3.

Jakub Nalepa, Pablo Ribalta Lorenzo, Michal Marcinkiewicz, Barbara Bobek-Billewicz, Pawel Wawrzyniak, Maksym Walczak, Michal Kawulok, Wojciech Dudzik, Krzysztof Kotowski, Izabela Burda, Bartosz Machura, Grzegorz Mrukwa, Pawel Ulrych, and Michael P. Hayball. Fully-automated deep learning-powered system for dce-mri analysis of brain tumors. *Artificial Intelligence in Medicine*, 102:101769, 2020. doi: 10.1016/j.artmed.2019.101769.

Leonie Henschel, Sailesh Conjeti, Santiago Estrada, Kersten Diers, Bruce Fischl, and Martin Reuter. Fastsurfer: A fast and accurate deep learning based neuroimaging pipeline. *NeuroImage*, 219:117012, 2020. doi: 10.1016/j.neuroimage.2020.117012.

Hugo Angleys, Leif Østergaard, and Sune N Jespersen. The effects of capillary transit time heterogeneity (cth) on brain oxygenation. *Journal of Cerebral Blood Flow & Metabolism*, 35(5):806–817, February 2015. ISSN 1559-7016. doi: 10.1038/jcbfm.2014.254. URL http://dx.doi.org/10.1038/jcbfm.2014.254.

Henrik B. W. Larsson, Mark B. Vestergaard, Ulrich Lindberg, Stig P. Cramer, and Helle K. Iversen. Brain capillary transit time heterogeneity in healthy volunteers measured by dynamic contrast-enhanced t1-weighted perfusion mri. *Journal of Magnetic Resonance Imaging*, 45(6):1809–1820, 2017a. doi: 10.1002/jmri.25488.

Mark B. Vestergaard, Helle K. Iversen, Sofie Amalie Simonsen, Ulrich Lindberg, Stig P. Cramer, Ulla B. Andersen, and Henrik B. W. Larsson. Capillary transit time heterogeneity inhibits cerebral oxygen metabolism in patients with reduced cerebrovascular reserve capacity from steno-occlusive disease. *Journal of Cerebral Blood Flow & Metabolism*, 43(3):460–475, 2023a. doi: 10.1177/0271678X221139084.

Henrik B. W. Larsson, Mark B. Vestergaard, Ulrich Lindberg, Hanne K. Iversen, and Stig P. Cramer. Brain capillary transit time heterogeneity in healthy volunteers measured by dynamic contrast-enhanced t1-weighted perfusion mri. *Journal of Magnetic Resonance Imaging*, 45(6):1809–1820, June 2017b. doi: 10.1002/jmri.25488. URL https://doi.org/10.1002/jmri.25488.

Mark B. Vestergaard, Hanne K. Iversen, Stine A. Simonsen, Ulrich Lindberg, Stig P. Cramer, Ulla B. Andersen, and Henrik B. Larsson. Capillary transit time heterogeneity inhibits cerebral oxygen metabolism in patients with reduced cerebrovascular reserve capacity from steno-occlusive disease. *Journal of Cerebral Blood Flow & Metabolism*, 43(3):460–475, March 2023b. doi: 10.1177/0271678X221139084. URL https://doi.org/10.1177/0271678X221139084.

Edis Devin Tireli. Trained cnn models for slice classification and roi segmentation in *p*-brain dce-mri pipeline. Zenodo, 2025. URL https://doi.org/10.5281/zenodo.15655348.

Iris C.M. Verheggen, Jaap J.A. de Jong, Martin P.J. van Boxtel, Erik H.B.M. Gronenschild, Wim M. Palm, Albert A. Postma, Jeroen F.A. Jansen, Frans R.J. Verhey, and Walter H. Backes. Increase in blood-brain barrier leakage in healthy, older adults. *Geroscience*, 42(4):1183–1193, 2020. doi: 10.1007/s11357-020-00211-2.

Michael A. Branch, Thomas F. Coleman, and Yuying Li. A subspace, interior, and conjugate gradient method for large-scale bound-constrained minimization problems. *SIAM Journal on Scientific Computing*, 21(1):1–23, 1999. doi: 10.1137/S1064827595289108.

Thomas F. Coleman and Yuying Li. An interior trust region approach for nonlinear minimization subject to bounds. *SIAM Journal on Optimization*, 6(2):418–445, 1996. doi: 10.1137/0806023.

Stig P Cramer, Henrik B W Larsson, Maria H Knudsen, Helle J Simonsen, Mark B Vestergaard, and Ulrich Lindberg. Reproducibility and optimal arterial input function selection in dynamic contrast-enhanced perfusion MRI in the healthy brain. *J. Magn. Reson. Imaging*, 57(4):1229–1240, April 2023b.

Cameron Manning, Michael Stringer, Ben Dickie, Una Clancy, Maria C Valdés Hernandez, Stewart J Wiseman, Daniela Jaime Garcia, Eleni Sakka, Walter H Backes, Michael Ingrisch, Francesca Chappell, Fergus Doubal, Craig Buckley, Laura M Parkes, Geoff J M Parker, Ian Marshall, Joanna M Wardlaw, and Michael J Thrippleton. Sources of systematic error in DCE-MRI estimation of low-level blood-brain barrier leakage. *Magn. Reson. Med.*, 86(4):1888–1903, October 2021.

Michael J Thrippleton, Walter H Backes, Steven Sourbron, Michael Ingrisch, Matthias J P van Osch, Martin Dichgans, Franz Fazekas, Stefan Ropele, Richard Frayne, Robert J van Oostenbrugge, Eric E Smith, and Joanna M Wardlaw. Quantifying blood-brain barrier leakage in small vessel disease: Review and consensus recommendations. *Alzheimers. Dement.*, 15(6):840–858, June 2019.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *MICCAI*, 2015. URL http://arxiv.org/abs/1505.04597.

Mark Jenkinson, Christian F. Beckmann, Timothy E. J. Behrens, Mark W. Woolrich, and Stephen M. Smith. Fsl. *NeuroImage*, 62(2):782–790, 2012. doi: 10.1016/j.neuroimage.2011.09.015.

Aravinthan Varatharaj, Carmen Jacob, Angela Darekar, Brian Yuen, Stig P. Cramer, Henrik B. W. Larsson, and Ian Galea. Measurement variability of blood–brain barrier permeability using dynamic contrast–enhanced mri. *Imaging Neuroscience*, 2:1–16, 2024b. doi: 10.1162/imag_a_00324.

S. Butterworth. On the theory of filter amplifiers. *Experimental Wireless & the Radio Engineer*, 7:536–541, October 1930.