Reconstruction and modelling of dynamic contrast-enhanced MRI: from k-space to K^{trans}

Natalia V. Korobova

Propositions belonging to the dissertation

Reconstruction and modelling of dynamic contrast-enhanced MRI: from k-space to K^{trans}

Natalia V. Korobova

1. DCE-MRI can play a substantial role in personalised medicine for cancer patients.
2. Overreliance on simulations risks models that excel at synthetic benchmarks but flop when applied to complex real-world data.
3. The direct correspondence between predicted parameters and ground truth is not a prerequisite, provided that the technique demonstrates clinical benefits for the patient.
4. The future success of DCE-MRI largely depends on its reproducibility and repeatability.
5. For DCE-MRI to be adopted clinically, collaboration among researchers, clinicians, software developers, DevOps engineers, and industry partners is essential.
6. Science must not allow the gap in medical technology between developed countries and the rest of the world to widen further.
7. A four-year, publication-driven PhD program prioritises short-term wins over long-term research impact.
8. Researchers are publicly funded and hence are responsible for their work to the public.
9. Travelling is a privilege.
10. With basic literacy and numeracy skills in decline, increased funding for school education is crucial for a bright future.

Reconstruction and modelling of dynamic contrast-enhanced MRI: from k-space to K^{trans}

Natalia V. Korobova

Copyright © 2026 by Natalia V. Korobova
This thesis is licensed under a Creative Commons “Attribution 4.0 International” license.
Financial support for printing of this thesis was kindly provided by Cancer Center Amsterdam
Cover: Maud van As
ISBN: 978-94-6534-522-2

Reconstruction and modelling of dynamic contrast-enhanced MRI: from k-space to Ktrans

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
ten overstaan van een door het College voor Promoties ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op vrijdag 18 september 2026, te 10.00 uur

door Natalia Vladimirovna Korobova
geboren te Kazan

Promotiecommissie

Promotores:	dr. O.J. Gurney-Champion	AMC-UvA
	prof. dr. ir. A.J. Nederveen	AMC-UvA
Overige leden:	prof. dr. ir. G.J. Strijkers	AMC-UvA
	prof. dr. H.A. Marquering	AMC-UvA
	prof. dr. E.T. van Bavel	AMC-UvA
	prof. dr. J.P.W. Pluim	Technische Universiteit Eindhoven
	dr. ir. P. van Houdt	NKI AvL
	dr. ir. C.M.W. Tax	Universiteit Utrecht

Faculteit der Geneeskunde

Contents

1 General introduction 3
- 1.1 MRI . . . 4
- 1.2 Dynamic contrast-enhanced MRI . . . 5
- 1.3 Applications . . . 9
- 1.4 Challenges of DCE-MRI . . . 12
- 1.5 Outline . . . 13

2 A correction for modeling radial, spiral, and PROPELLER dynamic contrast-enhanced data: Time-averaged extended Tofts 23
- 2.1 Introduction . . . 25
- 2.2 Methods . . . 28
- 2.3 Results . . . 32
- 2.4 Discussion and conclusions . . . 38
- Appendix 2.A Proof of Eq. 2.3 . . . 43
- Appendix 2.B Definition of analytical AIF proposed by Orton et al. . . . 43
- Appendix 2.C Definition of time-averaged AIF . . . 44
- Appendix 2.D Results of repeated Experiments 1 and 2 with time-averaged AIF 44
- Appendix 2.E Comparison of time-averaging implmentations in the concentration and signal domains . . . 50

3 Respiratory motion–corrected reconstruction of abdominal dynamic contrast-enhanced MRI 61
- 3.1 Introduction . . . 63
- 3.2 Materials and methods . . . 64
- 3.3 Results . . . 70
- 3.4 Discussion . . . 75

4 Model-based reconstruction for quantitative DCE-MRI parameters 85
- 4.1 Introduction . . . 87
- 4.2 Materials and methods . . . 89
- 4.3 Results . . . 97
- 4.4 Discussion . . . 101
- Appendix 4.A Configuration settings for PyQMRI . . . 109
- Appendix 4.B Selection of weighting factors . . . 109

5 Uncertainty estimates in pharmacokinetic modelling of DCE-MRI 121
- 5.1 Introduction . . . 123
- 5.2 Theory . . . 125
- 5.3 Methods . . . 127
- 5.4 Results . . . 135
- 5.5 Discussion . . . 144
- Appendix 5.A Signal model . . . 150

Appendix 5.B Arterial input function . . . 150
Appendix 5.C DCE-MRI acquisition protocol . . . 151
Appendix 5.D Physics-Informed Mean-Variance Estimation . . . 152
Appendix 5.E Additional visualisations . . . 154

6 General discussion 165
6.1 Image quality . . . 165
6.2 Repeatability and reproducibility . . . 168
6.3 Clinical translation . . . 170
6.4 Environmental impact . . . 171

Summary 178

Nederlandse samenvatting 181

Portfolio 185

Acknowledgments 187

Acknowledgements 187

Curriculum Vitae 192

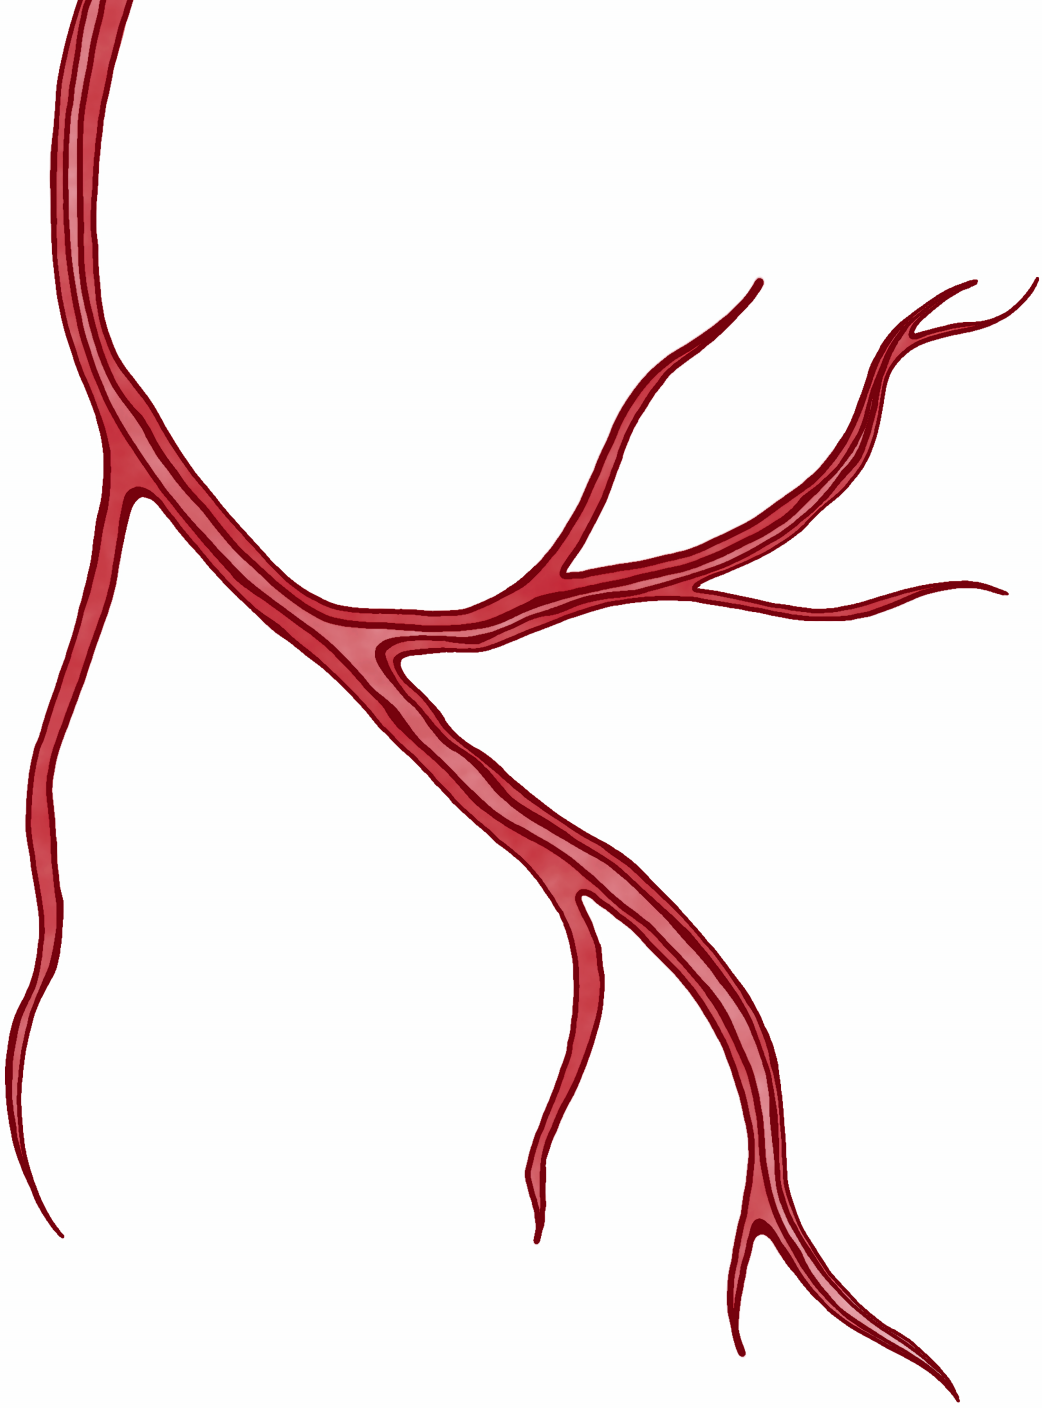

In the beginning, there was nothing, which exploded.

Sir Terry Pratchett

1

General introduction

The vascular network (Figure 1.1), consisting of large arteries, smaller arterioles, and capillaries, plays a fundamental role in maintaining homeostasis by regulating body temperature and transporting oxygen, carbon dioxide, and essential nutrients to tissues and organs. Remarkably, the importance of vessels was already recognized in ancient times - the *Ebers Papyrus*, dating ca. 1500 BCE from Ancient Egypt, contains early descriptions suggesting that vital substances, such as air, were believed to be conveyed through vessels within the body [1].

Both the structure and function of the vascular network are crucial for the health of cells and organs. Abnormalities in vascular organization or perfusion can indicate a wide range of pathological conditions, including hypoxic tumors [2], cerebral small vessel disease [3], and coronary artery disease [4]. Importantly, many of these vascular changes can now be assessed noninvasively using medical imaging techniques [5], reducing the need for invasive diagnostic procedures. Among these techniques, magnetic resonance imaging (MRI) offers particular advantages, as it is free of ionizing radiation and provides excellent soft-tissue contrast along with the ability to probe functional and physiological processes [6].

Among other MRI-based methods, dynamic contrast-enhanced MRI (DCE-MRI) enables the quantitative assessment of tissue perfusion and vascular permeability and has demonstrated considerable potential in oncology and other clinical applications [4, 7–11]. Despite this promise, DCE-MRI remains largely a research-oriented technique, owing to its methodological complexity and sensitivity to acquisition and processing choices [12]. This thesis aims to address some of these challenges by focusing on image reconstruction,

processing, and modeling aspects of DCE-MRI. Progress in these areas is essential for improving the robustness of the technique and supporting its translation into routine clinical practice.

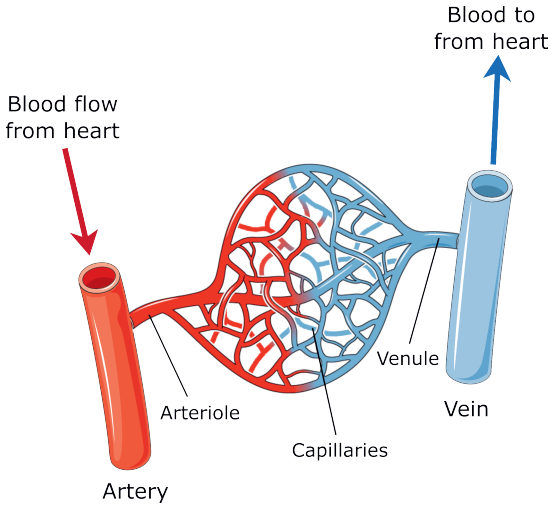

Figure 1.1: Schematic representation of a vascular network illustrating hierarchical branching from major vessels to microvasculature, highlighting structural organization and connectivity.

1.1 MRI

Magnetic Resonance Imaging (MRI) is a tomographic imaging modality for acquiring in vivo images of the human body, widely used for disease detection and diagnosis, treatment planning and monitoring, image-guided interventions, and the investigation of anatomical, functional, and physiological processes. MRI uses strong magnetic fields and radiofrequency pulses to generate high-contrast images of soft tissues. When hydrogen nuclei, present in water, fat, and other macromolecules, are placed in an external magnetic field, their nuclear spins align with the field and precess around its direction. Radiofrequency pulses are then applied to perturb this equilibrium, and the resulting collective spin response is detected and used to generate images of the body [13].

Image contrast in MRI primarily arises from differences in hydrogen proton density and relaxation properties across tissues. In particular, T_1 (longitudinal) and T_2 (transverse) relaxation times describe how quickly the perturbed nuclear magnetization returns to equilibrium along and perpendicular to the main magnetic field, respectively. By adjusting sequence parameters, MRI can be made sensitive to these relaxation mechanisms, yielding T_1-weighted and T_2-weighted images that emphasize different tissue characteristics and

provide complementary diagnostic information.

The spatial localisation of hydrogen nuclei in MRI is achieved through the application of magnetic field gradients that make the frequency and phase of proton precession position-dependent [14, 15]. This principle, discovered by Paul C. Lauterbur and Sir Peter Mansfield (recipients of the 2003 Nobel Prize in Physiology or Medicine [16]), enables the encoding of spatial information into the MR signal. The signal, characterized by its spatially varying frequencies and phases, is recorded in k-space (the spatial frequency domain). The final anatomical image is then reconstructed by applying an inverse Fourier transform to map the data from k-space into physical space.

1.2 Dynamic contrast-enhanced MRI

Dynamic contrast-enhanced (DCE) MRI is a specialized form of MRI involving the intravenous injection of a paramagnetic contrast agent. Its goal is to monitor the dynamic uptake and clearance of the contrast agent in a tissue of interest. The contrast agent creates a local reduction in T_1 relaxation time of nearby water protons, producing signal enhancement on T_1-weighted images. By acquiring multiple imaging volumes over time, typically 3-20 seconds per image [17], the dynamic uptake and subsequent washout of the contrast agent produce a time-varying signal intensity curve, the shape of which depends on tissue perfusion and microvascular permeability (Figure 1.2).

These dynamic signal curves can be analyzed to provide pharmacokinetic imaging biomarkers of tissue perfusion and vascular permeability. The quantitative DCE-MRI analysis involves several steps: estimation of contrast agent concentration, defining arterial input function, and pharmacokinetic modelling. In practice, however, a fully quantitative conversion to contrast agent concentration is not always performed. In some applications, the relative signal enhancement or normalized signal-time curves are analyzed directly to extract biomarkers.

1.2.1 Contrast agent concentration

In T_1-weighted DCE-MRI, the measured signal intensity is commonly converted into a time-dependent contrast agent concentration. This conversion is performed in two steps. First, signal changes over time are related to changes in T_1 time, and second, these T_1 changes are converted into contrast agent concentration over time [18].

To enable this conversion, the relationship between MR signal intensity and T_1 must be explicitly modelled, as it depends on the specific imaging sequence and requires accurate knowledge of acquisition parameters, in particular the flip angle (FA). In DCE-MRI,

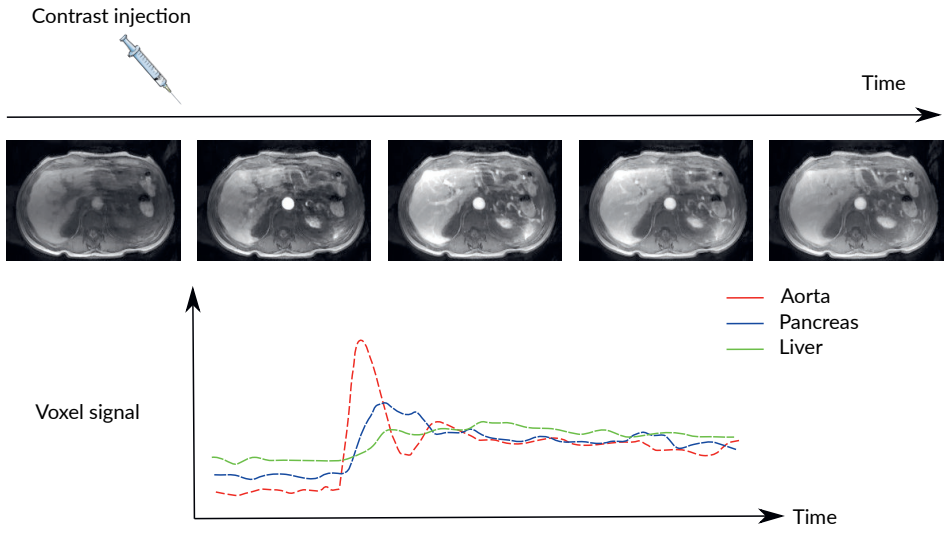

Figure 1.2: Example of an abdominal DCE-MRI scan. Contrast agent injection occurs shortly after the start of the scan. The tissue signal increases as the contrast agent enters the vasculature and extravascular space, followed by a gradual decrease as the agent washes out. Below the images, an example of signal intensity–time curves is shown for three representative voxels: the aorta (red), pancreas (blue), and liver (green).

a spoiled gradient-echo sequence is most commonly used due to its speed and strong T_1 weighting. For this sequence, the signal equation is given by

$$S = S_0 \frac{(1 - e^{-TR/T_1}) \sin FA}{1 - e^{-TR/T_1} \cos FA} \tag{1.1}$$

where S_0 represents the fully relaxed signal ($TR >> T_1$), and can be estimated from the measured pre-contrast signal intensity.

By inverting this signal model, post-contrast T_1 values can be obtained and subsequently used to estimate the contrast agent concentration. In addition to T_1 values, the relaxivity of the contrast agent, r_1, is required, which is typically determined from the literature or prior calibration measurements. The relationship between contrast agent concentration C and the longitudinal relaxation rate $R_1 = 1/T_1$ is given by:

$$\frac{1}{T_1} = \frac{1}{T_{10}} + r_1 C \tag{1.2}$$

where T_{10} denotes the native (pre-contrast) longitudinal relaxation time. Accurate estimation of T_{10} is required prior to DCE acquisition and analysis.

When baseline T_{10} values are available, the contrast agent relaxivity r_1 is known, and acquisition parameters such as the repetition time TR and flip angle FA are specified,

the measured signal intensity can be transformed into a quantitative estimate of contrast agent concentration.

1.2.2 Arterial input function

The arterial input function (AIF) describes the time-dependent concentration of contrast agent within an arterial vessel supplying the tissue of interest. In quantitative DCE-MRI analysis, the AIF is required to model the delivery of contrast agent to the tissue, enabling estimation of physiological parameters. In in vivo studies, the true AIF is generally unknown and difficult to measure directly. This difficulty arises from its characteristic temporal profile, which features rapid contrast uptake, a short and sharp peak, and a gradual washout phase (red line, Figure 1.2), as well as from the small size of the vessels in which it occurs, which can be challenging to segment and, especially in the abdomen, may exhibit motion over time.

Several approaches have been developed to estimate or determine the AIF accurately. The gold standard approach involves direct arterial blood sampling through a catheter inserted into the subject [19]. Although this technique provides highly accurate measurements, achieving the high temporal resolution necessary for precise AIF characterization is technically challenging. Additionally, the invasiveness of this approach might be inconvenient for some patients. Finally, the concentration curve in the vessel being studied may not be identical to the AIF feeding the tissue of interest.

A patient-specific AIF can be estimated from the patient´s imaging data, e. g. DCE-MRI or dynamic CT, by segmenting a region of interest (ROI) in a large feeding artery near the tissue of interest and extracting the signal intensity curve from this ROI. The accuracy of this approach depends on acquisition parameters, including temporal and spatial resolution, as well as signal-to-noise ratio (SNR). Additionally, signal measurements from arteries may be degraded by artifacts, such as partial volume effects, inflow effects, and blood flow pulsatility and turbulence.

A population-based AIF represents a commonly used alternative and is derived by averaging AIF measurements across a cohort of subjects. This method reduces the influence of individual artifacts and noise but sacrifices subject-specific information, potentially introducing bias in the estimation of pharmacokinetic parameters.

1.2.3 Pharmacokinetic modelling

Pharmacokinetic modelling describes the time-dependent concentration of contrast agent within tissue through a set of physiological parameters, often referred to as phar-

macokinetic parameters. This relationship is described as the response of the tissue to an impulse of contrast agent:

$$C(t) = IRF(t) * C_p(t) \tag{1.3}$$

where $C_p(t)$ is the impulse of contrast agent represented by AIF; $IRF(t)$ is the tissue impulse response function, which varies depending on the chosen pharmacokinetic model; and $C(t)$ is the measured tissue contrast agent concentration derived from the DCE-MRI signal [20].

A pharmacokinetic model is used to estimate the pharmacokinetic parameters that define $IRF(t)$. These parameters are obtained by fitting the model-predicted concentration curve to the measured concentration-time curve, such that the predicted tissue response best matches the observed data.

The choice of the pharmacokinetic model depends on the physiology of the tissue of interest and/or on a specific pathology. In the following chapters of this thesis, the extended Tofts-Kety model was extensively used for the DCE-MRI analysis. Therefore, here we will focus on this model.

1.2.3.1 Extended Tofts-Kety model

In the extended Tofts-Kety model, each voxel is described by three compartments: tissue cells, blood plasma, and the extravascular extracellular space (EES) [21, 22] (Figure 1.3). Following intravenous injection, a gadolinium-based contrast agent leaks from the blood plasma into the EES via passive diffusion driven by a concentration gradient. The assumption is that gadolinium-based contrast agents do not enter cells. Hence, their tissue concentration is governed by the fractional blood plasma volume v_p $[a.u.]$ and the fractional EES volume v_e $[a.u.]$, both expressed as dimensionless quantities. The temporal evolution of contrast agent concentration is further influenced by the volume transfer constant K^{trans} $[min^{-1}]$, which characterizes the influx of contrast agent from plasma into the EES, and by the rate constant k_{ep}, $[min^{-1}]$, which describes the reflux of contrast agent from the EES back into the plasma. Importantly, K^{trans} and k_{ep} are not independent parameters, as they are both derived from the same underlying exchange process and depend on concentration differences between compartments governed by diffusion. This coupling introduces parameter redundancy in the model. In practice, only two of the three parameters k_{ep}, K^{trans}, and v_e are typically estimated directly, while the third is derived from their relationship, most commonly $k_{ep} = \frac{K^{trans}}{v_e}$.

The IRF of the extended Tofts-Kety model is given by

$$IRF(t) = K^{trans} e^{-k_{ep}t} + v_p, \tag{1.4}$$

leading to the following expression for the tissue contrast agent concentration

$$C(t) = v_p C_p(t) + K^{trans} \int_0^t C_p(\tau) e^{-k_{ep}(t-\tau)} d\tau. \tag{1.5}$$

The pharmacokinetic parameters are estimated by fitting the model such that the resulting contrast agent concentration-time curve best matches the measured tissue concentration curve (Figure 1.4).

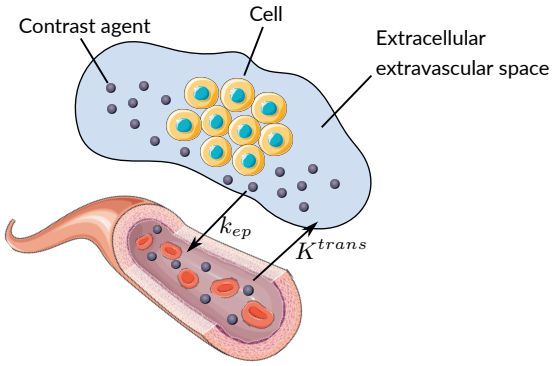

Figure 1.3: Illustration of the extended Tofts–Kety model. Following intravenous injection, the contrast agent (black circles) diffuses between the blood plasma and the extravascular extracellular space (EES, shown in light blue), but does not enter tissue cells (yellow and blue circles). The tissue concentration of the contrast agent is modelled using the fractional volumes of blood plasma (v_p) and EES (v_e), along with the exchange rate constants K^{trans} (plasma $\rightarrow$ EES) and k_{ep} (EES $\rightarrow$ plasma).

1.3 Applications

1.3.1 Oncology

DCE-MRI has a wide-range of applications, with oncology representing possibly the most established and extensively studied domain. Because tumor neovasculature differs from vasculature in healthy tissue, exhibiting hypervascularity, excessive and disorganized branching, irregular vessel morphology, and increased tortuosity, blood flows chaotically,

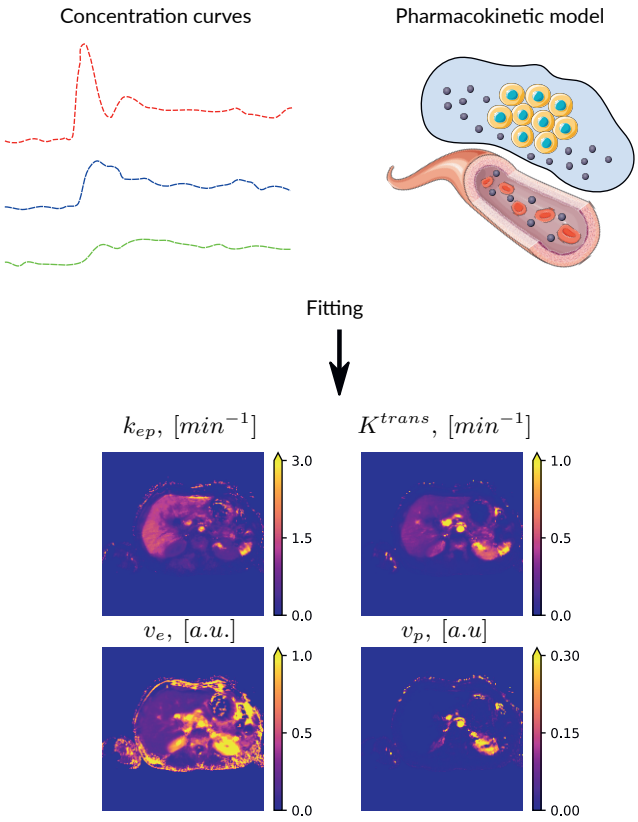

Figure 1.4: Schematic illustration of pharmacokinetic parameter estimation and visualization in DCE-MRI. Contrast agent concentration–time curves are extracted from the dynamic images and used to fit a pharmacokinetic model by optimizing model parameters. This fitting procedure is typically performed on a voxel-wise basis. The resulting pharmacokinetic parameters are then displayed as quantitative parametric maps, in which parameter values are encoded using a color scale.

accompanied by increased vascular permeability and leakiness [2]. Hence, DCE-MRI is well-suited to interrogate tumor microvascular physiology.

DCE-MRI can differentiate malignant from benign lesions by exploiting different patterns in contrast enhancement kinetics. The diagnostic value of DCE-MRI has been demonstrated across multiple tumor types, including breast cancer [7], soft tissue mass (includes muscles, tendons, ligaments, cartilage, fat and other tissues) [23], uterine tumors [24], liver tumors [25], and head and neck tumors [26]. Furthermore, DCE-MRI has shown capacity in assessing tumor grade and aggressiveness [27–29].

Beyond diagnosis, DCE-MRI has demonstrated a substantial value in monitoring therapeutic response. Treatment-induced vascular changes precede measurable reductions in tumor size [30,31]. Hence, functional imaging techniques, such as DCE-MRI, can detect response earlier than conventional imaging. The ability of DCE-MRI to monitor treatment response has been reported in patients receiving chemotherapy [32], radiotherapy [8], and targeted therapy [33].

Finally, DCE-MRI has shown a potential for generating prognostic biomarkers associated with tumor survival outcomes. In the future, such biomarkers may facilitate personalized treatment strategies, enabling patients with poor prognostic indicators to receive intensified therapy while allowing patients with a favourable prognosis to avoid overtreatment and its associated toxicities. Multiple studies investigated the prognostic value of DCE-MRI in breast cancer [34], cervical cancer [35], head and neck cancer [36], and liver cancer [9].

1.3.2 Other applications

DCE-MRI has demonstrated valuable applications in the musculoskeletal system, particularly in the assessment of inflammatory disorders and bone perfusion [37]. In inflammatory conditions such as Charcot foot, a disorder characterized by intense inflammation and microvascular dysfunction, DCE-MRI can distinguish between the acute inflammatory phase and clinical improvement by evaluating dynamic enhancement patterns [38]. Beyond inflammatory disorders, altered bone perfusion plays an important role in degenerative joint disease, as impaired or abnormal subchondral blood flow has been linked to cartilage loss. In this context, DCE-MRI has shown a positive association between quantitative bone perfusion parameters and the progression of osteoarthritis [10].

In neuroimaging, DCE-MRI has been used to assess blood–brain barrier (BBB) integrity and microvascular dysfunction across a range of neurological conditions. By quantifying contrast agent leakage and permeability, DCE-MRI enables the evaluation of BBB disruption in acute ischemic stroke [11], traumatic brain injury [39], and cerebral small vessel

disease [3]. In addition, subtle BBB breakdown detected with DCE-MRI has been reported in neurodegenerative disorders, including Alzheimer's disease [40], where microvascular impairment may contribute to disease pathophysiology.

In cardiovascular imaging, DCE-MRI has demonstrated potential in the assessment of myocardial perfusion, which plays a central role in the pathophysiology of coronary artery disease (CAD) [4]. In particular, myocardial perfusion imaging with DCE-MRI has shown strong prognostic value for risk stratification in patients with known or suspected CAD [41]. Furthermore, the use of DCE-MRI–based perfusion assessment has been shown to reduce the need for unnecessary invasive coronary angiography without an associated increase in adverse clinical outcomes [42]. In patients with multivessel coronary artery disease, quantitative perfusion cardiovascular MR has demonstrated superior diagnostic accuracy at the vessel level compared with qualitative perfusion assessment [43].

1.4 Challenges of DCE-MRI

Although DCE-MRI has demonstrated considerable potential across a wide range of clinically relevant applications, its routine clinical implementation remains challenging. In fact, quantitative DCE-MRI is not yet adopted in clinical practice and remains a research-oriented technique. This work focuses on challenges associated with image reconstruction, signal modeling, and post-processing. While additional limitations exist, such as cross-site repeatability and reproducibility, hardware- and sequence-dependent variability, day-to-day and physiological fluctuations, uncertainties in contrast agent relaxivity, and errors in native T_1 estimation, these factors are considered beyond the scope of this work.

Like all MRI techniques, DCE-MRI involves an inherent trade-off between scan time, spatial resolution, and SNR. However, this trade-off is particularly critical in the context of DCE-MRI. High temporal resolution is essential for accurately capturing the dynamic contrast agent concentration–time curve, especially during the rapid bolus arrival phase in feeding arteries, where signal changes occur over very short time scales. At the same time, high spatial resolution is crucial for applications such as tumor imaging, as tumors are known to be highly heterogeneous and fine spatial detail is required to resolve intratumoral variations in vascularity and permeability. Finally, high SNR plays a key role in the reliable estimation of pharmacokinetic parameters, as increasing noise levels directly degrade the accuracy and stability of quantitative perfusion and permeability estimates. Although acceleration techniques such as parallel imaging [44] and compressed sensing [45] have substantially advanced DCE-MRI, they often come at the cost of reduced SNR

(parallel imaging's G-factor) or temporal blurring (compressed sensing regularisation). Hence, finding algorithms that can deal with poor SNR or poor temporal resolution would allow us to redistribute the balance in favour of spatial resolution, improving overall image quality. Moreover, if we are able to move outside the conventional trade-off constraints, we can further improve the DCE image quality [46].

Due to this trade-off, clinical DCE-MRI protocols typically acquire only a limited number of temporal phases, rather than high temporal resolution acquisitions commonly used in research settings. A standard clinical protocol generally consists of one pre-contrast image followed by three or more post-contrast acquisitions, timed to sample the arterial, venous, and delayed phases of contrast enhancement. For imaging of the abdomen and thorax, 10-20 second breath-holding is frequently required, as respiratory motion can severely degrade image quality. In research-oriented DCE-MRI, however, breath-holding is not feasible because of the high temporal resolution requirement. Hence, alternative motion-handling strategies are necessary. Although motion-robust acquisition schemes for abdominal DCE-MRI have been proposed, such as golden-angle radial stack-of-stars acquisitions, these approaches reduce motion but do not fully eliminate respiratory motion.

The estimation of pharmacokinetic parameters in DCE-MRI is equal finding a solution to an inverse problem. This problem is inherently ill-posed because the solution may be non-unique or unstable, and in certain cases, such as non-perfused tissues, a valid solution may not exist at all. As a result, fitting algorithms can converge to different parameter estimates depending on the choice of initial conditions. Moreover, the presence of noise, which is unavoidable in all experimental data, can cause optimization algorithms to become trapped in local minima or to overfit noise. Consequently, the extraction of pharmacokinetic parameters from accelerated or low-SNR DCE-MRI data requires careful methodological design, robust fitting strategies, and cautious interpretation of results.

1.5 Outline

To meet the high temporal resolution required for accurate pharmacokinetic modeling, many DCE studies rely on undersampling of k-space, often combined with oversampling of the k-space center to improve SNR and mitigate motion-related artifacts (e.g. radial and spiral sampling). This preferential sampling of the k-space center can influence the temporal characteristics of the reconstructed signal and, consequently, affect pharmacokinetic modeling. In **Chapter 2**, we propose a modification of the standard extended Tofts model that incorporates a temporal adjustment to account for the effects of oversampling of the

k-space center. Additionally, we investigate how the choice of temporal resolution impacts the magnitude of the required correction.

Respiratory motion during abdominal DCE-MRI is a major source of complex artifacts. If not properly addressed, voxel displacement caused by breathing can lead to temporal signal mixing, resulting in corrupted contrast-enhancement curves and adversely affecting pharmacokinetic modeling. In **Chapter 3**, we propose a motion-corrected compressed sensing reconstruction that incorporates respiratory motion fields directly into the reconstruction process, thereby producing motion-compensated dynamic image series. This approach, unlike image-based registration techniques, also corrects for intra-scan motion. Using this approach, we demonstrate improved image quality in abdominal DCE-MRI examinations.

As discussed earlier, achieving high temporal resolution in DCE-MRI is inherently constrained by trade-offs with SNR and spatial resolution, both of which are essential for adequate image quality. In **Chapter 4**, we investigate whether this spatiotemporal trade-off can be alleviated through the direct estimation of pharmacokinetic parameters from k-space data. Specifically, we employ a model-based reconstruction framework in which the inverse problem formulation is extended with a pharmacokinetic model, namely the extended Tofts model. This approach is compared with a conventional pipeline consisting of image reconstruction followed by voxel-wise pharmacokinetic parameter fitting. Our results demonstrate that direct parameter estimation can mitigate the spatiotemporal trade-off, enabling faster imaging without compromising the reliability of the estimated pharmacokinetic parameters.

Quantitative DCE-MRI typically involves voxel-wise estimation of pharmacokinetic parameters. Conventional non-linear least-squares fitting is known to be highly sensitive to noise, whereas recent deep learning–based approaches tend to be more robust and often produce smoother, visually appealing parametric maps. However, such visual smoothness can inadvertently create a false sense of confidence if the underlying parameter estimates are inaccurate. To address this concern, **Chapter 5** introduces a deep learning framework that estimates both pharmacokinetic parameters and their associated uncertainties. By explicitly quantifying uncertainty, this approach aims to support more informed interpretation of AI-generated parametric maps and reduce the risk of overconfidence in the presence of noisy or ambiguous data.

Finally, **Chapter 6** provides a general discussion of the thesis, building on the technical contributions in image reconstruction, modelling, and post-processing, and placing them in a broader context. It expands beyond these contributions to discuss key challenges for the clinical translation of quantitative DCE-MRI, including image quality, reproducibility,

and clinical validation, while also addressing limitations and outlining directions for future research.

1

Bibliography

[1] J. Barr, "Vascular medicine and surgery in ancient egypt," *Journal of Vascular Surgery*, vol. 60, no. 1, pp. 260–263, 2014.

[2] P. Carmeliet and R. K. Jain, "Angiogenesis in cancer and other diseases," *nature*, vol. 407, no. 6801, pp. 249–257, 2000.

[3] M. J. Thrippleton, W. H. Backes, S. Sourbron, M. Ingrisch, M. J. van Osch, M. Dichgans, F. Fazekas, S. Ropele, R. Frayne, R. J. van Oostenbrugge *et al.*, "Quantifying blood-brain barrier leakage in small vessel disease: review and consensus recommendations," *Alzheimer's & Dementia*, vol. 15, no. 6, pp. 840–858, 2019.

[4] R. Franks, S. Plein, and A. Chiribiri, "Clinical application of dynamic contrast enhanced perfusion imaging by cardiovascular magnetic resonance," *Frontiers in cardiovascular medicine*, vol. 8, p. 768563, 2021.

[5] M. Günther, "Perfusion imaging," *Journal of Magnetic Resonance Imaging*, vol. 40, no. 2, pp. 269–279, 2014.

[6] D. W. McRobbie, E. A. Moore, M. J. Graves, and M. R. Prince, *MRI from Picture to Proton*. Cambridge university press, 2017.

[7] S. A. Jansen, X. Fan, G. S. Karczmar, H. Abe, R. A. Schmidt, and G. M. Newstead, "Differentiation between benign and malignant breast lesions detected by bilateral dynamic contrast-enhanced mri: a sensitivity and specificity study," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 59, no. 4, pp. 747–754, 2008.

[8] R. Mulyadi, P. P. Putri, R. A. Zairinal, J. Prihartono *et al.*, "Dynamic contrast-enhanced magnetic resonance imaging parameter changes as an early biomarker of tumor responses following radiation therapy in patients with spinal metastases: a systematic review," *Radiation Oncology Journal*, vol. 41, no. 4, p. 225, 2023.

[9] B.-B. Chen, Z.-Z. Lin, Y.-Y. Shao, C. Hsu, C.-H. Hsu, A.-L. Cheng, P.-C. Liang, and T. T.-F. Shih, "Early changes in dce-mri biomarkers may predict survival outcomes in patients with advanced hepatocellular carcinoma after sorafenib failure: two prospective phase ii trials," *Cancers*, vol. 13, no. 19, p. 4962, 2021.

[10] L. T. Muftuler, J. P. Jarman, H. J. Yu, V. O. Gardner, D. J. Maiman, and V. E. Arpinar, "Association between intervertebral disc degeneration and endplate perfusion studied by dce-mri," *European Spine Journal*, vol. 24, no. 4, pp. 679–685, 2015.

[11] K. Villringer, B. E. Sanz Cuesta, A.-C. Ostwaldt, U. Grittner, P. Brunecker, A. A. Khalil, K. Schindler, O. Eisenblätter, H. Audebert, and J. B. Fiebach, "Dce-mri blood–brain barrier assessment in acute ischemic stroke," *Neurology*, vol. 88, no. 5, pp. 433–440, 2017.

[12] X. M. Li, T. Nguyen, H. D. Sparks, K. Sung, and J. Chiang, "Quantitative dynamic contrast-enhanced magnetic resonance imaging (dce-mri) in hepatocellular carcinoma: A review of emerging applications for locoregional therapy," *Bioengineering*, vol. 12, no. 8, p. 870, 2025.

[13] R. W. Brown, Y.-C. N. Cheng, E. M. Haacke, M. R. Thompson, and R. Venkatesan, *Magnetic resonance imaging: physical principles and sequence design*. John Wiley & Sons, 2014.

[14] P. C. Lauterbur, "Image formation by induced local interactions: examples employing nuclear magnetic resonance," *nature*, vol. 242, no. 5394, pp. 190–191, 1973.

[15] P. Mansfield and P. K. Grannell, "Nmr'diffraction'in solids?" *Journal of Physics C: solid state physics*, vol. 6, no. 22, p. L422, 1973.

[16] C. Boesch, "Nobel prizes for nuclear magnetic resonance: 2003 and historical perspectives," *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 20, no. 2, pp. 177–179, 2004.

[17] A. Shukla-Dave, N. A. Obuchowski, T. L. Chenevert, S. Jambawalikar, L. H. Schwartz, D. Malyarenko, W. Huang, S. M. Noworolski, R. J. Young, M. S. Shiroishi *et al.*, "Quantitative imaging biomarkers alliance (qiba) recommendations for improved precision of dwi and dce-mri derived biomarkers in multicenter oncology trials," *Journal of Magnetic Resonance Imaging*, vol. 49, no. 7, pp. e101–e121, 2019.

[18] P. S. Tofts, "T1-weighted dce imaging concepts: modelling, acquisition and analysis," *signal*, vol. 500, no. 450, p. 400, 2010.

[19] F. Khalifa, A. Soliman, A. El-Baz, M. Abou El-Ghar, T. El-Diasty, G. Gimel'farb, R. Ouseph, and A. C. Dwyer, "Models and methods for analyzing dce-mri: A review," *Medical physics*, vol. 41, no. 12, p. 124301, 2014.

[20] B. Dickie and P. J. van Houdt, "Dynamic contrast-enhanced mri," in *Advances in Magnetic Resonance Technology and Applications*. Elsevier, 2023, vol. 11, pp. 17–40.

[21] P. S. Tofts, "Modeling tracer kinetics in dynamic gd-dtpa mr imaging," *Journal of magnetic resonance imaging*, vol. 7, no. 1, pp. 91–101, 1997.

[22] P. S. Tofts, G. Brix, D. L. Buckley, J. L. Evelhoch, E. Henderson, M. V. Knopp, H. B. Larsson, T.-Y. Lee, N. A. Mayr, G. J. Parker *et al.*, "Estimating kinetic parameters from dynamic contrast-enhanced t1-weighted mri of a diffusable tracer: standardized quantities and symbols," *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 10, no. 3, pp. 223–232, 1999.

[23] F. Del Grande, S. Ahlawat, T. Subhangwong, and L. Fayad, "Characterization of indeterminate soft tissue masses referred for biopsy: what is the added value of contrast imaging at 3.0 tesla?" *Journal of Magnetic Resonance Imaging*, vol. 45, no. 2, pp. 390–400, 2017.

[24] A. Goto, S. Takeuchi, K. Sugimura, and T. Maruo, "Usefulness of gd-dtpa contrast-enhanced dynamic mri and serum determination of ldh and its isozymes in the differential diagnosis of leiomyosarcoma from degenerated leiomyoma of the uterus," *International Journal of Gynecological Cancer*, vol. 12, no. 4, pp. 354–361, 2002.

[25] X. Zheng, L. Xiao, X. Fan, N. Huang, Z. Su, and X. Xu, "Free breathing dce-mri with motion correction and its values for benign and malignant liver tumor differentiation," *Radiology of Infectious Diseases*, vol. 2, no. 2, pp. 65–71, 2015.

[26] M. El Backry, M. Shady, A. E. Mousa, and M. M. Zaky, "Role of dynamic contrast enhanced mr perfusion in differentiation between benign and malignant tumors," *The Egyptian Journal of Radiology and Nuclear Medicine*, vol. 46, no. 3, pp. 715–726, 2015.

[27] X. Li, Q. Wang, Y. Dou, Y. Zhang, J. Tao, L. Yang, and S. Wang, "Soft tissue sarcoma: can dynamic contrast-enhanced (dce) mri be used to predict the histological grade?" *Skeletal radiology*, vol. 49, no. 11, pp. 1829–1838, 2020.

[28] J. Liang, D. Liu, P. Gao, D. Zhang, H. Chen, C. Shi, and L. Luo, "Diagnostic values of dce-mri and dsc-mri for differentiation between high-grade and low-grade gliomas: a comprehensive meta-analysis," *Academic Radiology*, vol. 25, no. 3, pp. 338–348, 2018.

[29] E. K. Vos, G. J. Litjens, T. Kobus, T. Hambrock, C. A. Hulsbergen-van de Kaa, J. O. Barentsz, H. J. Huisman, and T. W. Scheenen, "Assessment of prostate cancer aggres-

siveness using dynamic contrast-enhanced magnetic resonance imaging at 3 t," *European urology*, vol. 64, no. 3, pp. 448–455, 2013.

[30] J. P. O'Connor, R. A. Carano, A. R. Clamp, J. Ross, C. C. Ho, A. Jackson, G. J. Parker, C. J. Rose, F. V. Peale, M. Friesenhahn *et al.*, "Quantifying antivascular effects of monoclonal antibodies to vascular endothelial growth factor: insights from imaging," *Clinical Cancer Research*, vol. 15, no. 21, pp. 6674–6682, 2009.

[31] J. D. Rojas, V. Papadopoulou, T. J. Czernuszewicz, R. M. Rajamahendiran, A. Chytil, Y.-C. Chiang, D. C. Chong, V. L. Bautch, W. K. Rathmell, S. Aylward *et al.*, "Ultrasound measurement of vascular density to evaluate response to anti-angiogenic therapy in renal cell carcinoma," *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 3, pp. 873–880, 2018.

[32] C. Hayes, A. R. Padhani, and M. O. Leach, "Assessing changes in tumour vascular function using dynamic contrast-enhanced magnetic resonance imaging," *NMR in Biomedicine: An International Journal Devoted to the Development and Application of Magnetic Resonance In Vivo*, vol. 15, no. 2, pp. 154–163, 2002.

[33] M. A. Rosen and M. D. Schnall, "Dynamic contrast-enhanced magnetic resonance imaging for assessing tumor vascularity and vascular effects of targeted therapies in renal cell carcinoma," *Clinical Cancer Research*, vol. 13, no. 2, pp. 770s–776s, 2007.

[34] M. G. Heldahl, T. F. Bathen, J. Rydland, K. A. Kvistad, S. Lundgren, I. S. Gribbestad, and P. E. Goa, "Prognostic value of pretreatment dynamic contrast-enhanced mr imaging in breast cancer patients receiving neoadjuvant chemotherapy: overall survival predicted from combined time course and volume analysis," *Acta Radiologica*, vol. 51, no. 6, pp. 604–612, 2010.

[35] F. Qin, H. Pang, J. Ma, M. Zhao, X. Jiang, R. Tong, T. Yu, Y. Luo, and Y. Dong, "Combined dynamic contrast enhanced mri parameter with clinical factors predict the survival of concurrent chemo-radiotherapy in patients with 2018 figo iiicr stage cervical cancer," *European journal of radiology*, vol. 141, p. 109787, 2021.

[36] A. Shukla-Dave, N. Y. Lee, J. F. Jansen, H. T. Thaler, H. E. Stambuk, M. G. Fury, S. G. Patel, A. L. Moreira, E. Sherman, S. Karimi *et al.*, "Dynamic contrast-enhanced magnetic resonance imaging as a predictor of outcome in head-and-neck squamous cell carcinoma patients with nodal metastases," *International Journal of Radiation Oncology* Biology* Physics*, vol. 82, no. 5, pp. 1837–1844, 2012.

[37] P. Sujlana, J. Skrok, and L. M. Fayad, "Review of dynamic contrast-enhanced mri: technical aspects and applications in the musculoskeletal system," *Journal of Magnetic Resonance Imaging*, vol. 47, no. 4, pp. 875–890, 2018.

[38] M. McGill, L. Molyneaux, T. Bolton, K. Ioannou, R. Uren, and D. Yue, "Response of charcot's arthropathy to contact casting: assessment by quantitative techniques," *Diabetologia*, vol. 43, no. 4, pp. 481–484, 2000.

[39] J. B. Ware, S. Sinha, J. Morrison, A. E. Walter, J. J. Gugger, A. L. Schneider, C. Dabrowski, H. Zamore, L. Wesley, B. Magdamo *et al.*, "Dynamic contrast enhanced mri for characterization of blood-brain-barrier dysfunction after traumatic brain injury," *NeuroImage: Clinical*, vol. 36, p. 103236, 2022.

[40] A. Montagne, S. R. Barnes, D. A. Nation, K. Kisler, A. W. Toga, and B. V. Zlokovic, "Imaging subtle leaks in the blood–brain barrier in the aging human brain: potential pitfalls, challenges, and possible solutions," *Geroscience*, vol. 44, no. 3, pp. 1339–1351, 2022.

[41] M. J. Lipinski, C. M. McVey, J. S. Berger, C. M. Kramer, and M. Salerno, "Prognostic value of stress cardiac magnetic resonance imaging in patients with known or suspected coronary artery disease: a systematic review and meta-analysis," *Journal of the American College of Cardiology*, vol. 62, no. 9, pp. 826–838, 2013.

[42] J. P. Greenwood, D. P. Ripley, C. Berry, G. P. McCann, S. Plein, C. Bucciarelli-Ducci, E. Dall'Armellina, A. Prasad, P. Bijsterveld, J. R. Foley *et al.*, "Effect of care guided by cardiovascular magnetic resonance, myocardial perfusion scintigraphy, or nice guidelines on subsequent unnecessary angiography rates: the ce-marc 2 randomized clinical trial," *Jama*, vol. 316, no. 10, pp. 1051–1060, 2016.

[43] T. Kotecha, L. Chacko, O. Chehab, N. O'Reilly, A. Martinez-Naharro, J. Lazari, K. D. Knott, J. Brown, D. Knight, V. Muthurangu *et al.*, "Assessment of multivessel coronary artery disease using cardiovascular magnetic resonance pixelwise quantitative perfusion mapping," *Cardiovascular Imaging*, vol. 13, no. 12, pp. 2546–2557, 2020.

[44] A. Deshmane, V. Gulani, M. A. Griswold, and N. Seiberlich, "Parallel mr imaging," *Journal of Magnetic Resonance Imaging*, vol. 36, no. 1, pp. 55–72, 2012.

[45] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing mri," *IEEE signal processing magazine*, vol. 25, no. 2, pp. 72–82, 2008.

[46] N. V. Korobova, S. S. Rauh, M. Wennen, E. M. Schrauben, O. Maier, and O. J. Gurney-Champion, "Model-based reconstruction for quantitative dce-mri parameters," *Magnetic Resonance Materials in Physics, Biology and Medicine*, pp. 1–17, 2026.

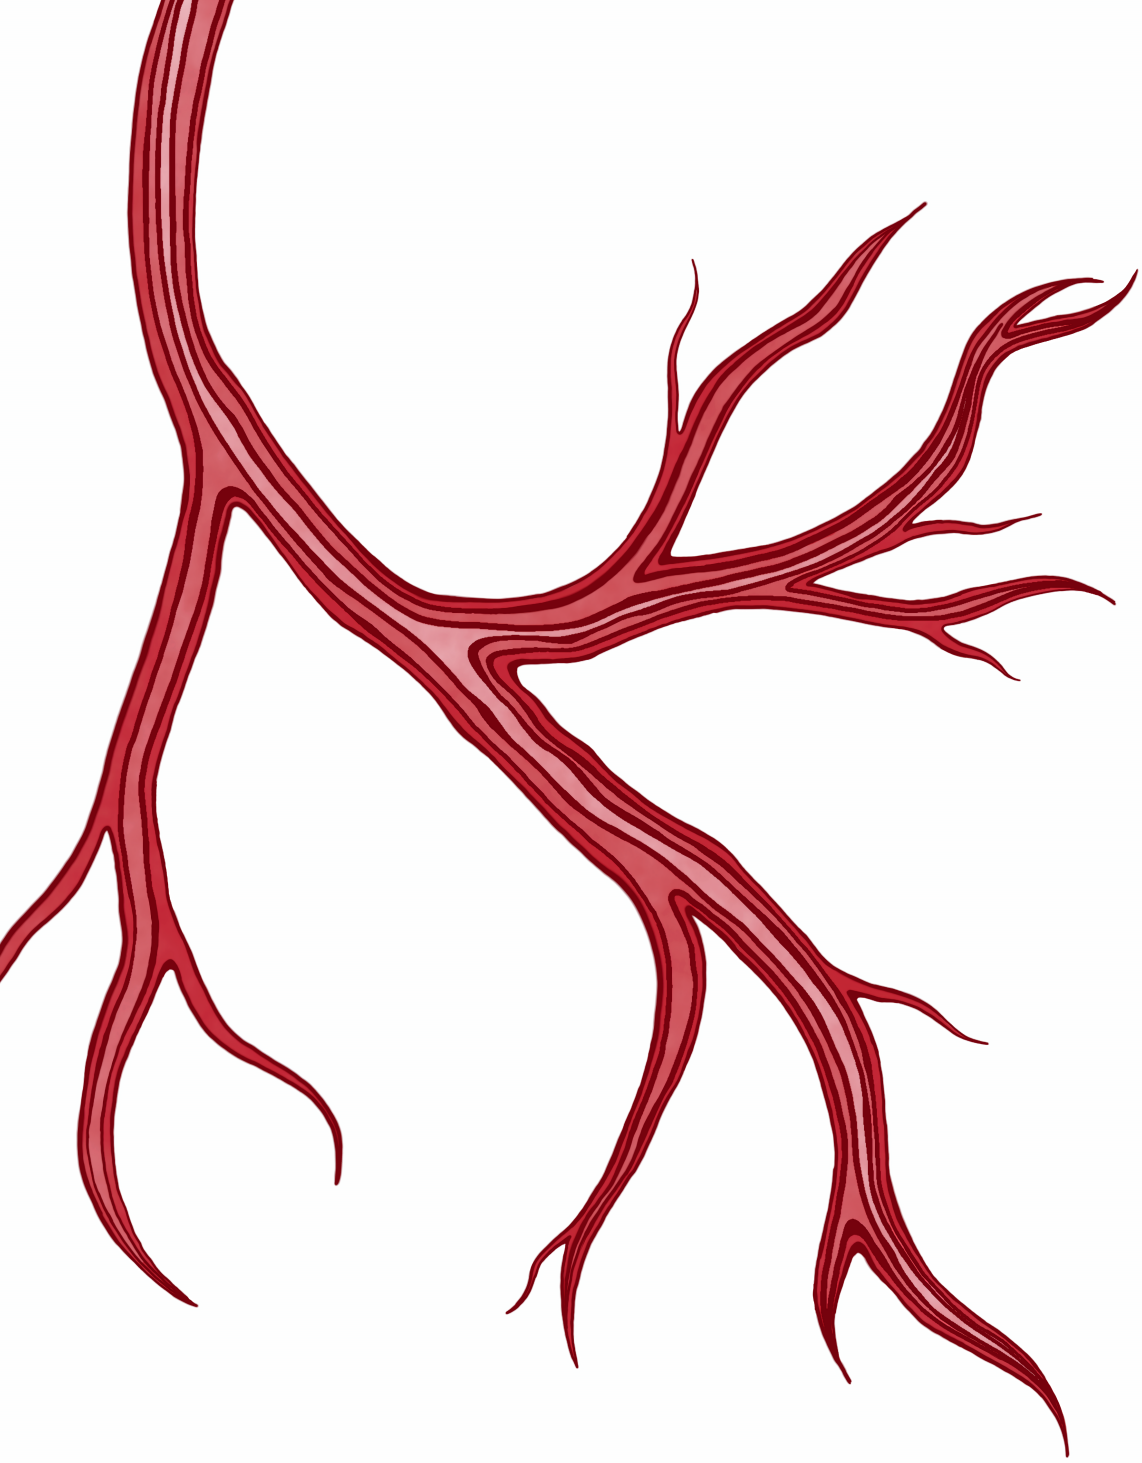

2

A correction for modeling radial, spiral, and PROPELLER dynamic contrast-enhanced data: Time-averaged extended Tofts

Authors: Natalia V. Korobova, Nienke P. M. Wassenaar, Marian A. Troelstra, Eric M. Schrauben, Oliver J. Gurney-Champion

Published in: Magnetic Resonance in Medicine (Volume 94, March 2025)

https://doi.org/10.1002/mrm.30514

Abstract

Purpose: Dynamic contrast-enhanced sequences (e.g. spiral, radial, PROPELLER MRI) often rely on oversampling the center of k-space. Instead of the discrete snapshots obtained by Cartesian sampling, oversampling the k-space center results in time-averaging of the signal. We hypothesize that these time-averaged signals decrease the accuracy of pharmacokinetic modeling and propose a model that accounts for this effect.

Theory and Methods: To test our hypothesis, a modified extended Tofts model tailored to accommodate time-averaged signals is proposed. Simulated Monte Carlo experiments were conducted to compare the performance of the modified model with the conventional model. Additionally, to validate the findings in vivo, models were fitted to pseudo-spiral variable-density dynamic contrast-enhanced MRI scans of pancreatic cancer patients reconstructed at 4, 8, 10, and 15 s/frame.

Results: The simulations demonstrated that for time-averaged acquisitions, our modified extended Tofts model provided more accurate and precise results than conventional models. Additionally, by integrating signals, some information on high temporal behavior was recovered. Particularly, at long acquisitions (15 s/frame), variable-density sampling with the modified model outperformed conventional discrete sampling. In vivo experiments confirmed these findings, as the corrected model showed more consistent estimates of parameters v_p and v_e over the tested sampling frequencies, highlighting its potential to improve accuracy in clinical settings.

Conclusion: Our study demonstrates that time-averaged signals lead to decreased accuracy and precision in pharmacokinetic modeling when ignored. We suggest using our corrected pharmacokinetic model when performing dynamic contrast-enhanced with variable-density acquisitions, especially for dynamic scan times that are 8 s and longer.

2.1 Introduction

Dynamic contrast-enhanced (DCE) MRI offers crucial pathophysiological insight across various areas of clinical research, such as oncology, neurovascular disease, and neurodegenerative conditions [1–5]. DCE-MRI is performed by a dynamic acquisition of T_1-weighted images before and after the intravenous administration of exogenous gadolinium-based contrast agent. The paramagnetic properties of the contrast agent shorten local T_1 relaxation time in blood and tissues; hence, a physiology-dependent signal increase occurs in T_1-weighted images. Through pharmacokinetic modeling, physiological parameters describing tissue microvascular perfusion and vascular permeability can be determined from the DCE signal-time curve [6].

Conventional pharmacokinetic models, including the extended Tofts model [7, 8], describe the contrast agent concentration (C) as a continuous function of time (t) : $C(t)$. These concentration time curves are typically derived from signal time curves using (1) an equation relating changes in T_1 values to changes in the signal [9] and (2) an equation relating changes in T_1 to contrast concentration. Pharmacokinetic parameters are then fitted to discrete concentration samples at time points corresponding to individual dynamic scans in the DCE series. This approach perfectly reflects the conventional Cartesian sampling, in which the center of k-space, which predominantly determines the image contrast, is sampled once per frame (Figure 2.1A).

Accurate quantitative modeling of the DCE signal requires a high temporal resolution [10], which is difficult to achieve over a three-dimensional (3D) volume (e.g., whole liver) using conventional Cartesian sampling. Recently, high-temporal-resolution volumetric DCE has been achieved using variable density sampling patterns, including radial [11], (pseudo-) spiral [12, 13], and PROPELLER [14] techniques. In these acquisitions, the center of k-space is acquired multiple times during each DCE frame. Consequently, the signal is averaged over the acquisition time for each temporal frame (Figure 2.1A). Pharmacokinetic models do not take this into account and therefore may give biased parameter estimates (Figure 2.1B). Hence, these acquisition techniques for DCE modeling need to incorporate time-averaging behavior in signal modeling.

Moreover, such time-averaged signals inherently contain some information about faster processes that occur within their temporal frame, which would have been missed by conventional discrete snapshot acquisitions. For example, the peak of the contrast-agent bolus could be completely missed by the discrete sampling but will influence a time-averaged signal. Time-averaged acquisitions may therefore potentially retain information about fast processes even at slow scan times. This is particularly relevant for

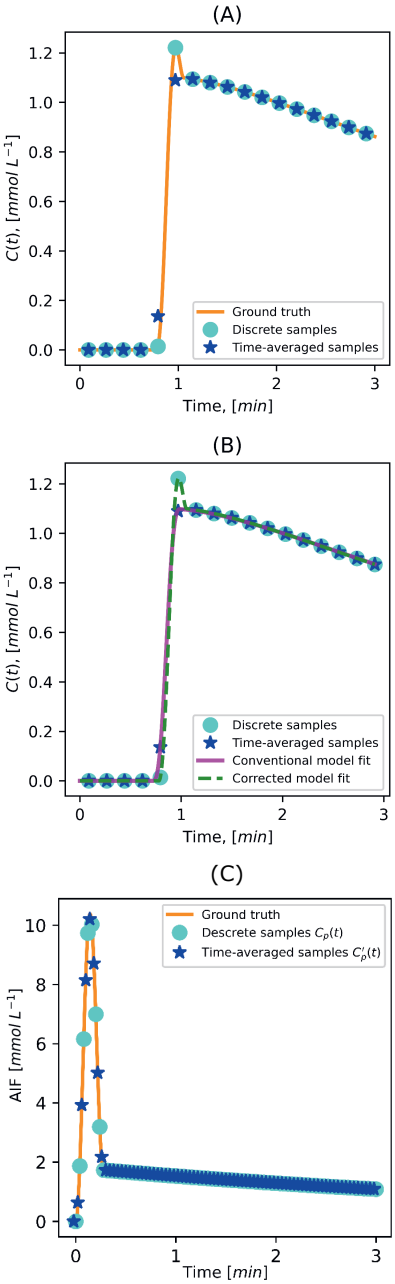

Figure 2.1: (A) Example contrast concentration curve showing the ground-truth curve (orange line), discrete sampling points using a conventional Cartesian acquisition (light blue circles), and time-averaged sampling points obtained with a variable-density trajectory that oversamples the center of k-space (dark blue stars). (B) Example of fitting conventional (pink line) and corrected (green line) pharmacokinetic models to time-averaged samples, highlighting the differences in model accuracy. (C) Example arterial input function (AIF) curve with its discrete sampling (light blue circles) and time-averaged sampling (dark blue stars) shown for comparison.

the supplying artery, where the concentration time curve of the contrast agent, known as the arterial input function (AIF), is crucial for accurate pharmacokinetic modeling [15] (Figure 2.1C).

We hypothesize that these time-averaged signals decrease the accuracy of conventional pharmacokinetic modeling. We further hypothesize that they can result in more accurate parameter estimates when properly modeled. Here we introduce a modified extended Tofts model tailored to account for time-averaged signals. We then use Monte Carlo experiments to explore to what extent time-averaged signals influence DCE modeling when ignored, and whether this can be improved by our new corrected model. Finally, we demonstrate its performance in vivo in pancreatic cancer patients.

2.1.1 Theory

2.1.1.1 DCE model

In this work, the extended Tofts model [7, 8] is used to quantify pharmacokinetic tissue properties. This model assumes that the contrast agent diffuses between two compartments: blood plasma and extracellular-extravascular space. The concentration of the contrast agent at every time point, t, is expressed through pharmacokinetic parameters, $u = [v_p, v_e, K^{trans}]$, as follows:

$$\begin{aligned} C(t, u, C_p(t)) &= v_p C_p(t) + K^{trans} C_p(t) * e^{-\frac{K^{trans}}{v_e} t} \\ &= v_p C_p(t) + K^{trans} \int_0^t C_p(\tau) e^{-\frac{K^{trans}}{v_e}(t-\tau)} d\tau, \end{aligned} \tag{2.1}$$

where v_p [fraction] is the volume of plasma; v_e [fraction] is the volume of extracellular-extravascular space; K^{trans} [min^{-1}] is the influx mass transfer rate; C_p [mmolL^{-1}] is the concentration of the contrast agent in the blood plasma, expressed as the AIF; and $*$ represents a convolution operation. The reflux rate of contrast k_{ep} [min^{-1}] is defined as the ratio between K^{trans} over v_e.

2.1.1.2 Modification of the model

To describe the time-averaged signal per acquisition, we propose to modify the extended Tofts model. To achieve signal averaging over time, we analytically convolve the pharmacokinetic model with a rectangular function with a $rect$-width equivalent to the acquisition time of a DCE frame, Δt, as follows:

$$C'(t,u,C_p(t)) = \frac{1}{\Delta t}C(t,u,C_p) * rect(\frac{t}{\Delta t}) = \frac{1}{\Delta t}\int_{t-\frac{1}{2}\Delta t}^{t+\frac{1}{2}\Delta t} C(\tau,u,C_p)d\tau, \tag{2.2}$$

where $C'(t,u,C_p)$ represents the time-averaged extended Tofts model. Although $C(t,u,C_p)$ is defined as the extended Tofts model in Eq. 2.1 for this work, in principle, Eq. 2.2 is general and holds for all pharmacokinetic models. Mathematically, the convolution of the extended Tofts model (Eq. 2.1) with the rectangular function as described in Eq. 2.2 is equivalent to only convolving the AIF with the following rectangular function:

$$C'(t,u,C_p(t)) = C(t,u,C'_p(t)), \tag{2.3}$$

with

$$C'_p(t) = \frac{1}{\Delta t}C_p(t) * rect(\frac{t}{\Delta t}), \tag{2.4}$$

where $C'_p(t)$ represents the time-averaged AIF. The proof is provided in Appendix 2.A. Hence, the convolution of the AIF with the rectangular function is theoretically sufficient for the correction; measuring the AIF with the same time-averaged acquisition as the DCE scan (e.g., patient-specific AIF), and using conventional pharmacokinetic modeling (Eq. 2.1), theoretically should not result in a bias.

2.2 Methods

All analyses were done in *Python* (v3.9.16; Python Software Foundation). We used the extended Tofts model implementation and fit routines from the *Python* implementation from OG_MO_AUMC_ICR_RMH_NL_UK [16] from OSIPI GitHub [17] as standard DCE Equation 2.1. The chosen model contains an analytical representation of AIF, $C_p(t,\theta_{AIF})$, dependent on a set of four parameters θ_{AIF}, that describes the curve as an initial pass, followed by wash-out (defined in Appendix 2.B). Parameters θ_{AIF} define the shape of AIF. This analytical form of AIF simplifies the mathematical integration in Eq.2.1 with an analytical solution.

To implement the time-averaged variant (Eq.2.2), we analytically convolved this implementation with a block function, for which the solution can be found in Appendix 2.C and on the OSIPI repository: `https://github.com/OSIPI/DCE-DSC-MRI_CodeCollection/tree/develop/src/original/NK_OG_MO_AUMC_ICR_RMH_NL_UK`.

All models were fitted with the nonlinear least squares algorithm with the following constraints: $v_e\ [0-1],\ k_{ep}\ [0-3],\ v_p\ [0-1]$.

We first address our hypothesis in a controlled simulated setting. Subsequently, we show the performance in vivo.

2.2.1 Simulations

2.2.1.1 Data

To evaluate the effect of signal averaging on the pharmacokinetic modeling in a controlled environment, we performed several Monte Carlo simulations. Each simulation contained 50000 contrast-agent concentration curves $C(t, u, \theta_{AIF})$ with a broad range of different pharmacokinetic parameters. The parameters v_e [0–1] and k_{ep} [0–3] were sampled from a uniform distribution and v_p from an exponential distribution with a rate parameter $\lambda = 40$, corresponding to a mean of 0.025, within the range [0–0.37]. To examine the performance over the entire range of parameter v_p, we simulated a similar data set, in which v_p was sampled uniformly [0–1]. Boundaries of simulated parameters included typically observed values in the pancreas, muscle, spleen, and liver [18]. Random Gaussian noise was added to the curves. The standard deviation of the noise was the same for both $C(t, u, \theta_{AIF})$ and $C'(t, u, \theta_{AIF})$. For a given curve, the standard deviation of noise was constant, meaning that signal-to-noise ratio (SNR) changed over time. The standard deviation of the noise was chosen such that the ratio of the mean signal in the time-averaged curve $C'(t, u, \theta_{AIF})$ divided by this standard deviation gave an SNR of 20. We simulated 5 min of $C(t, u, \theta_{AIF})$ at various temporal resolutions Δt: 4, 8, 10, and 15 s/frame. Simulated signal curves were sampled using both of the following (Figure 2.2):

A A temporally discrete Cartesian acquisition $C(t, u, \theta_{AIF})$ (Eq. 2.1)

B A temporally averaged variable density acquisition $C'(t, u, \theta_{AIF})$ (Eq. 2.2)

2.2.1.2 Experiments

First, we answered our first hypothesis and investigated how time-averaged signals decreased the accuracy of conventional pharmacokinetic modeling. We fitted the conventional extended Tofts model to (A) temporally averaged signal (Figure 2.2: Experiment 1) and, as a reference, to (B) discrete signal (Figure 2.2: Reference experiment).

Then, we answered our second hypothesis and investigated whether using our corrected equation allows for correct estimates. This was achieved using our proposed fit

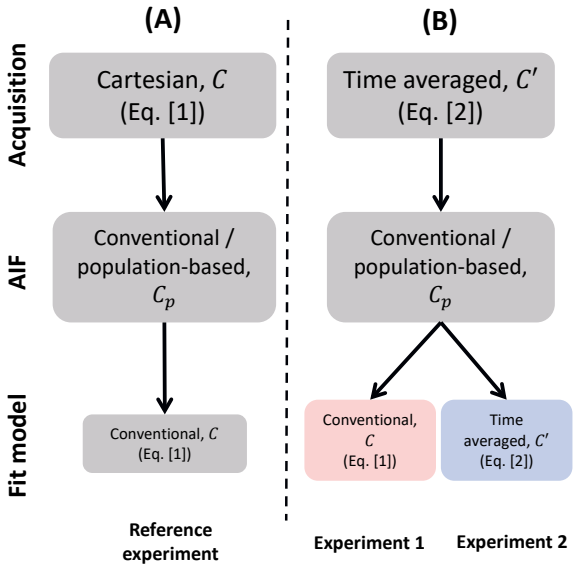

Figure 2.2: Overview of experiments conducted in simulations. (A) Reference experiment replicated temporally discrete data acquired with Cartesian trajectory. The conventional extended Tofts model was used in the fitting of pharmacokinetic parameters. (B) Alternative experiments simulated temporally averaged concentration curves resembling acquisition trajectories that oversample the center of k-space. Both conventional and corrected extended Tofts models were used for fitting. A population-based AIF was used across all experiments.

(Figure 2.2: Experiment 2) and comparing it with Experiment 1. To investigate whether time averaging indeed better retains pharmacokinetic information, Experiment 2 was also compared with the reference experiment. For all experiments, θ_{AIF} was assumed to be known, reflecting a population-based AIF (e.g., from literature [19], obtained with discrete acquisitions reflecting ground-truth curves). Parameters θ_{AIF} were then incorporated into analytically expressed AIF, which further was used in both conventional and corrected extended Tofts models.

Additionally, we extended Experiments 1 and 2 by changing a population-based AIF with a convolved AIF C_p' (with the equations from Appendix 2.C), reflecting the measurement of a patient-specific AIF with a variable density acquisition. The results are not included in the main text but are added to Appendix 2.D.

2.2.1.3 Evaluation

The performance of models was estimated in terms of accuracy (measured by the mean error in the fitted pharmacokinetic parameters) and precision (represented by the standard deviation of these errors). Differences in accuracy across various experiments were tested with a paired Wilcoxon signed-rank test, whereas the difference in precision was assessed

with a paired Levene test. A p-value below 0.05 was considered significant.

2.2.2 In vivo

We validated our methods in vivo, where no ground-truth reference is available. Instead, we assessed whether our proposed method keeps pharmacokinetic parameters consistent across different reconstructed dynamic times. In contrast, the conventional approach may introduce bias when the time per dynamic is too long. To test this, we reconstructed the data at 4, 8, 10, and 15 s per frame.

2.2.2.1 Data acquisition

Eight patients diagnosed with pancreatic ductal adenocarcinoma underwent an abdominal DCE-MRI examination. The study was approved by the ethics committee of the Amsterdam UMC location AMC (NL73810.018.20), and all participants gave written consent before the start of the study. MRI scans were performed on a 3T system (Ingenia; Philips, Best, The Netherlands), using a 16-channel phased-array anterior coil and a 16-channel phased-array posterior receive coil. After 30 s, 0.2 mL/kg of 0.5 mmol/mL gadoterate meglumine (Dotarem; Guerbet, Villepinte, France) followed by a 15 mL saline flush were administered at a rate of 5 mL/s. The DCE protocol consisted of a dynamic series of a 3D pseudo-spiral turbo field echo–spoiled gradient-echo sequence with spectral-attenuated inversion-recovery fat suppresion [13]. Pseudo-spiral sampling was achieved using a modified version of the free-running Amsterdam UMC ("Prospective Undersampling in Multiple Dimensions" [PROUD]) software patch `https://mriresearch.amsterdam/software/aumcproudpatch/` [20, 21]. The imaging parameters were as follows: 3D axial (right–left readout) orientation with field of view = 374 × 258 × 105 mm, flip angle = 25°, turbo field echo shot length = 50 readouts, spectral-attenuated inversion-recovery prepulse duration = 61.1 ms, repetition time/echo time = 3.8/1.8 ms, number of slices = 35, matrix size = 220 × 152, slice thickness = 3 mm, in-plane resolution = 1.7 mm, and acquisition time = 7 min. The dynamic series were reconstructed in MATLAB (R2023a; MathWorks) with the BART Toolbox (v0.5.00; `https://mrirecon.github.io/bart/`) using compressed sensing [22, 23] with regularization in the temporal (total variation $\lambda = 0.05$) and spatial (total variation $\lambda = 0.05$) domain at temporal resolutions of 4, 8, 10, and 15 s/frame. Participants were instructed to breathe freely during the entire DCE image acquisition. Post hoc breathing motion correction was performed based on an auto-focus approach [24] as described by Wassenaar et al [13].

2.2.2.2 Data postprocessing

Before pharmacokinetic modeling, the dynamic signal was converted to the concentration of the contrast agent with the spoiled gradient-echo signal model. The baseline T_1 time was set to 725 ms for the pancreas [25].

Similar to the simulations, pharmacokinetic parameters were calculated using the conventional (Eq. 2.1) and corrected (Eq. 2.2) models. A population-based AIF was used and was adjusted for an average hematocrit fraction of a healthy adult of 42% [26].

2.2.2.3 Evaluation

For all 8 patients, the pharmacokinetic parameters were calculated within a manually segmented region of interest (ROI) in the tumor. Additionally, these parameters were measured in non-tumorous pancreatic tissue in 6 patients — 2 patients were excluded due to atrophy of the pancreas. Delineation of ROIs was performed by a researcher with 4 years of experience in pancreatic imaging on the DCE-MRI scan with a 15-s temporal resolution. The arterial phase frame (between 20 s and 35 s following contrast injection) was manually selected for ROI placement. The clinically performed contrast-enhanced CT scan and a T_2-weighted turbo spin-echo MRI scan made during the research MR protocol were used for anatomical guidance. The performance of the models was considered to be robust when predictions of pharmacokinetic parameters were consistent across different temporal resolutions. Therefore, to estimate the accuracy of both models, we calculated the mean absolute difference between predictions derived from the fastest dynamic series (4 s/frame) and the remaining series for each subject. We used a paired signed-rank Wilcoxon test to assess whether both methods significantly differ.

2.3 Results

2.3.1 Simulations

Figures 2.3, 2.4 and 2.5 present estimated parameters K^{trans}, v_p, v_e obtained from fitting the extended Tofts model with and without correction for discretely sampled (Figure 2.2A) and time-averaged (Figure 2.2B) signals acquired at different temporal resolutions. As anticipated, all experiments showed superior performance at the highest temporal resolution, where averaging is almost identical to discrete sampling, and decreased performance at lower temporal resolutions.

When averaging is not accounted for, measuring with a time-averaged signal (Experiment 1) had a substantial decrease in performance compared with conventional Cartesian

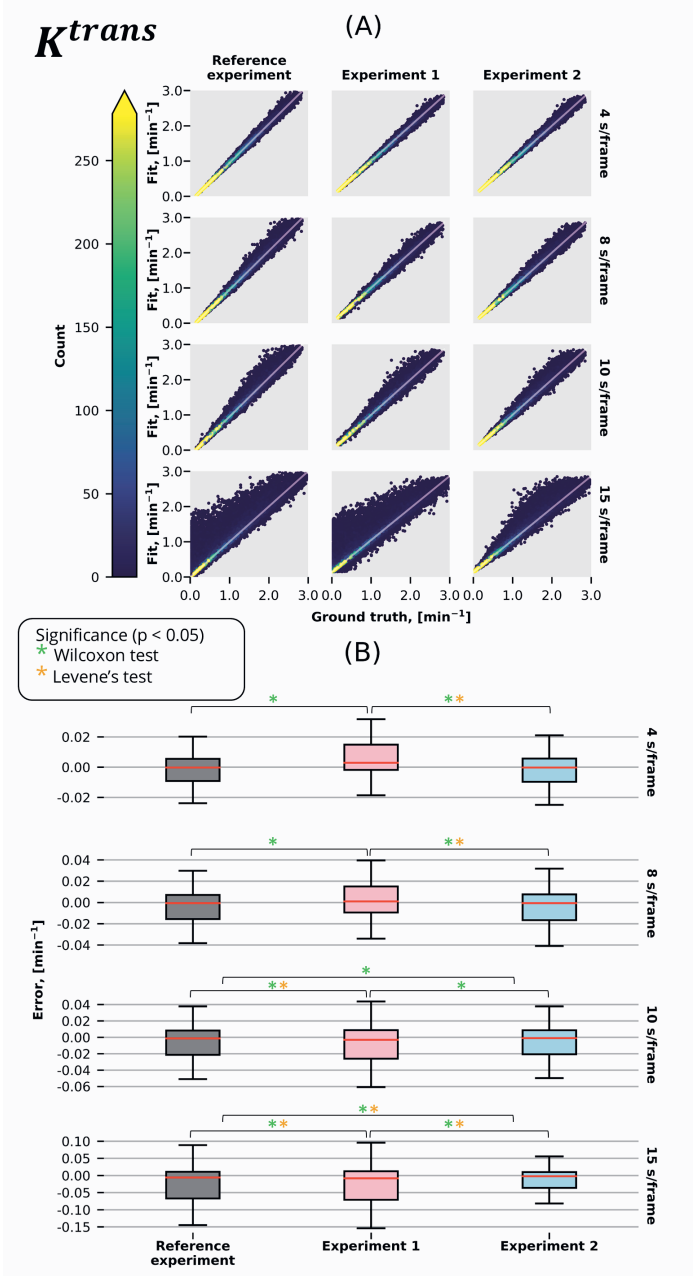

Figure 2.3: The K^{trans} results of the experiments described in Figure 2.2 were conducted at various temporal resolutions. (A) The graph displays the fitted value of the parameter K^{trans} plotted against the ground truth. The color gradient represents the density of points, whereas the white diagonal line represents the function Fit = Ground Truth. (B) Distribution of errors in the predicted parameter K^{trans} for the same experiments and temporal resolutions. Significant differences, as identified by the Wilcoxon (accuracy) and Levene (precision) tests (p < 0.05), are denoted by green and orange stars, respectively.

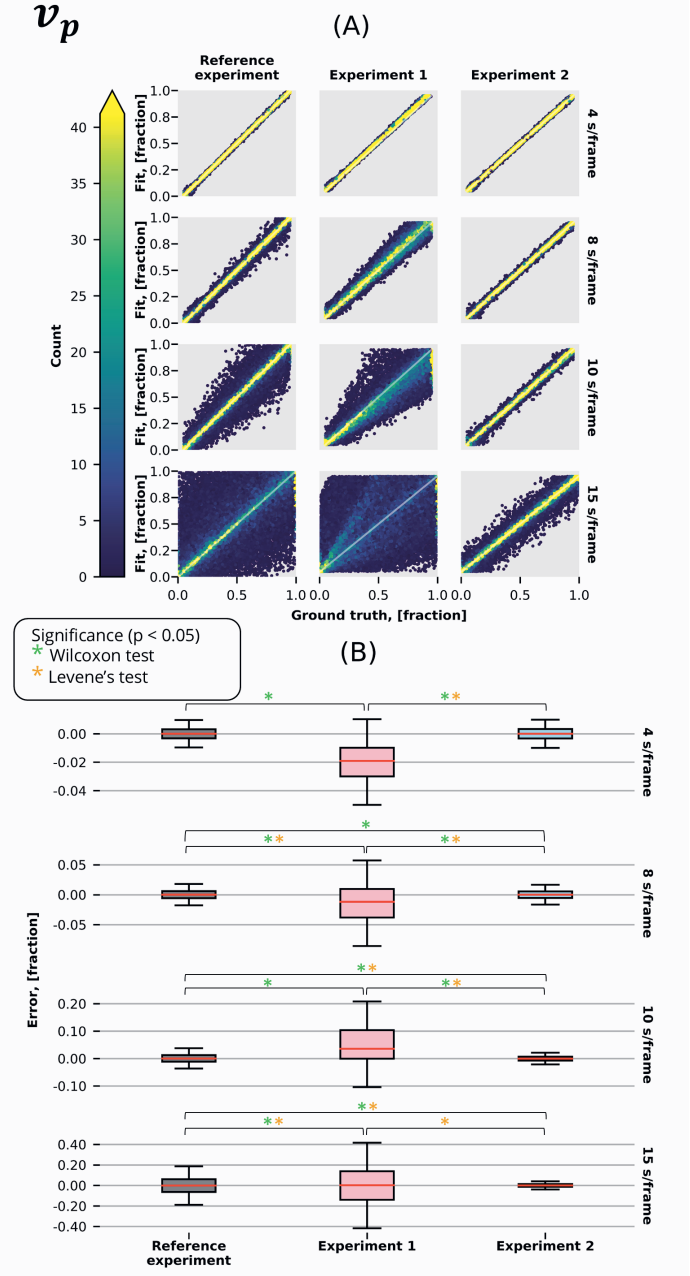

Figure 2.4: The v_p results of the experiments described in Figure 2.2 were conducted at various temporal resolutions. (A) The graph displays the fitted value of the parameter v_p plotted against the ground truth. The color gradient represents the density of points, whereas the white diagonal line represents the function Fit = Ground Truth. (B) Distribution of errors in the predicted parameter v_p for the same experiments and temporal resolutions. Significant differences, as identified by the Wilcoxon (accuracy) and Levene (precision) tests (p < 0.05), are denoted by green and orange stars, respectively.

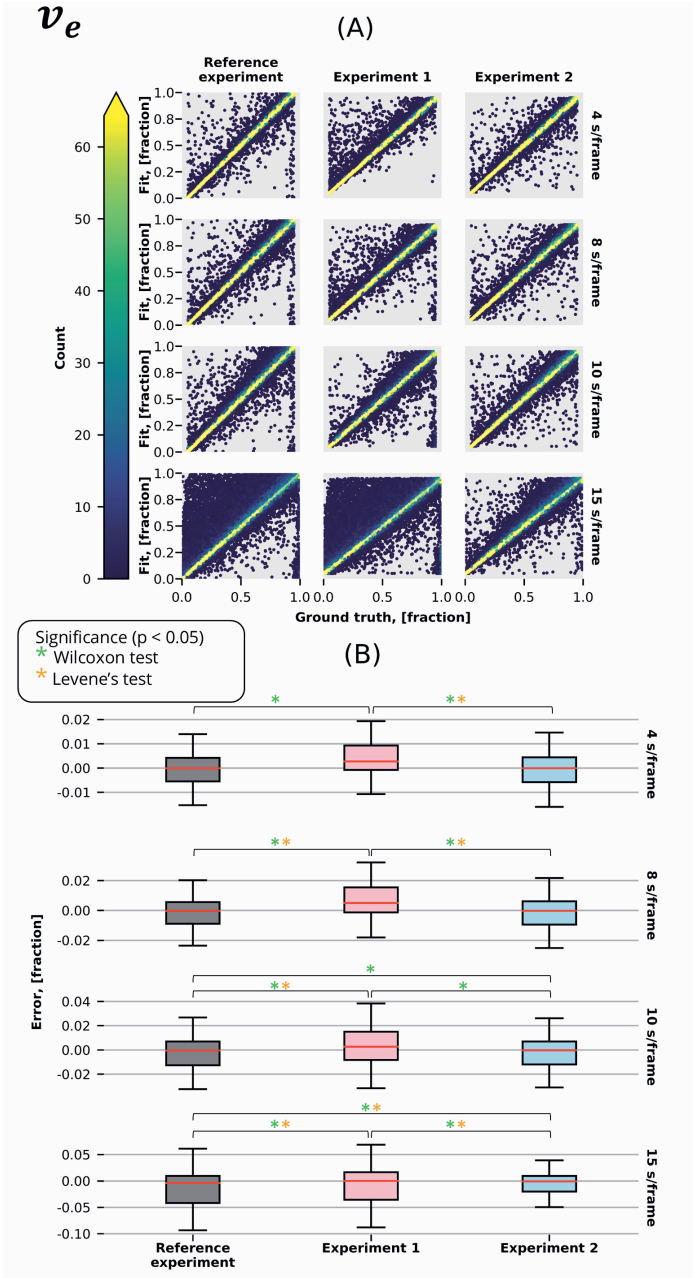

Figure 2.5: The v_e results of the experiments described in Figure 2.2 were conducted at various temporal resolutions. (A) The graph displays the fitted value of the parameter v_e plotted against the ground truth. The color gradient represents the density of points, whereas the white diagonal line represents the function Fit = Ground Truth. (B) Distribution of errors in the predicted parameter v_e for the same experiments and temporal resolutions. Significant differences, as identified by the Wilcoxon (accuracy) and Levene (precision) tests (p < 0.05), are denoted by green and orange stars, respectively.

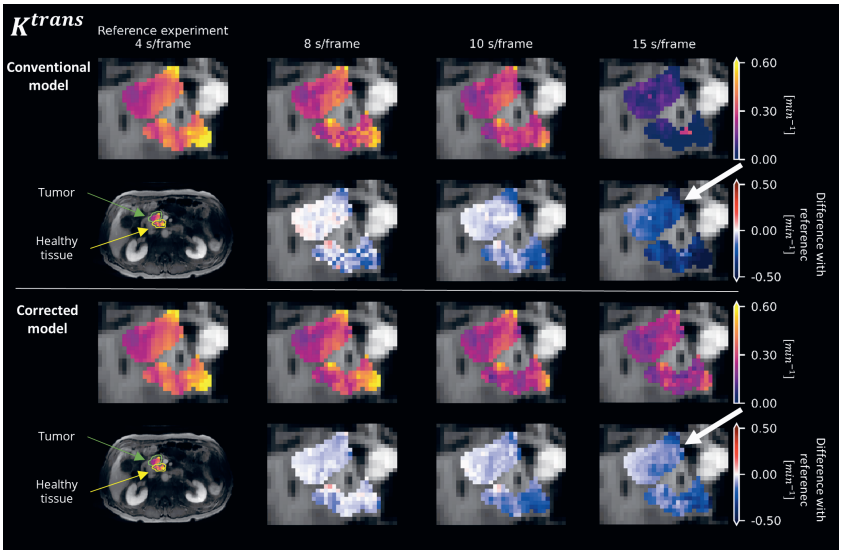

Figure 2.6: Estimation of the parameter K^{trans} in the tumor (*green arrow*) and healthy tissue (*yellow arrow*) of a pancreatic ductal adenocarcinoma patient. The results were calculated using both the conventional model (*Row 1*) and the corrected model (*Row 3*) at different temporal resolutions. Additionally, the corresponding differences in predictions (*Rows 2* and *4*) across various temporal resolutions compared with the reference (*Column 1*), which correspond to the fastest acquisition, are shown.

measurements (reference experiment). Particularly, we observed significantly better precision in 7 of 12 comparisons (3 parameters, 4 temporal resolutions) and accuracy in 12 of 12 comparisons.

When measuring a time-averaged signal, the correction in pharmacokinetic modeling (Experiment 2) produced significantly more accurate (12 of 12 comparisons) and precise (9 of 12 comparisons) estimates in comparison to the conventional model.

Interestingly, at longer temporal resolutions (10 and 15 s/frame), the conventional Cartesian measurements resulted in substantially poorer estimates when compared with the time-averaged model. Particularly, accuracy was significantly improved in 4 of 6 comparisons (3 parameters, 2 temporal resolutions) and precision in 6 of 6.

2.3.2 In vivo

Overall, the in vivo experiments (Figures 2.6,2.7,2.8) provided empirical evidence supporting the effectiveness of the corrected extended Tofts model in accurately estimating pharmacokinetic parameters from averaged signals.

The corrected model demonstrated greater consistency over different temporal resolutions compared with the conventional model, suggesting its ability to handle averaging

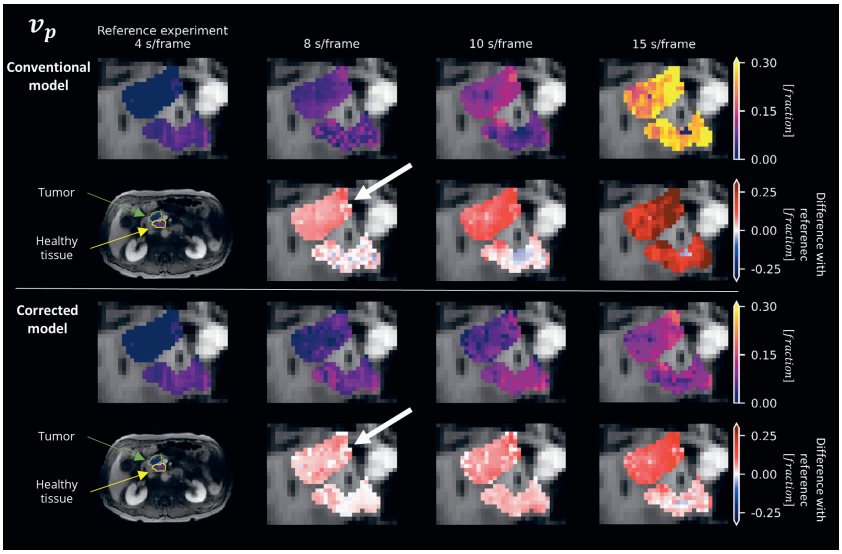

Figure 2.7: Estimation of the parameter v_p in the tumor (*green arrow*) and healthy tissue (*yellow arrow*) of a pancreatic ductal adenocarcinoma patient. The results were calculated using both the conventional model (*Row 1*) and the corrected model (*Row 3*) at different temporal resolutions. Additionally, the corresponding differences in predictions (*Rows 2* and *4*) across various temporal resolutions compared with the reference (*Column 1*), which correspond to the fastest acquisition, are shown.

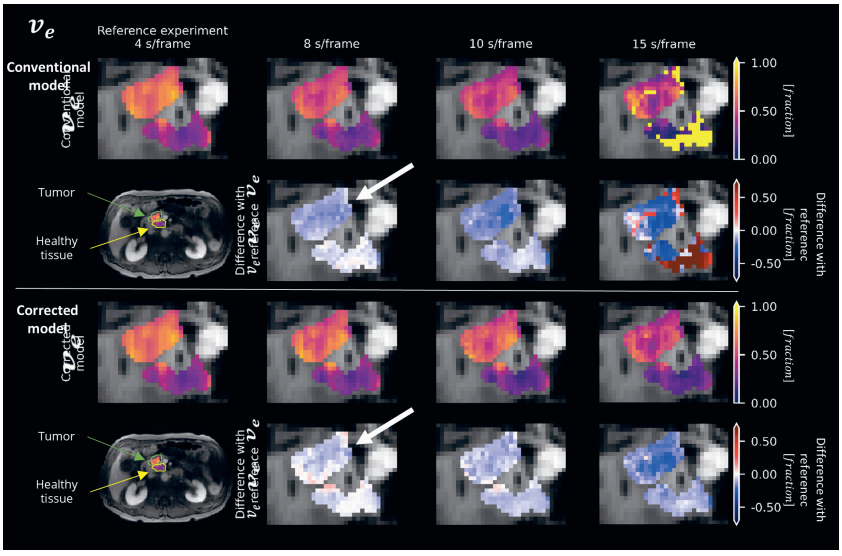

Figure 2.8: Estimation of the parameter v_e in the tumor (*green arrow*) and healthy tissue (*yellow arrow*) of a pancreatic ductal adenocarcinoma patient. The results were calculated using both the conventional model (*Row 1*) and the corrected model (*Row 3*) at different temporal resolutions. Additionally, the corresponding differences in predictions (*Rows 2* and *4*) across various temporal resolutions compared with the reference (*Column 1*), which correspond to the fastest acquisition, are shown.

signals more effectively. Despite the small number of patients, the corrected model was significantly more consistent in its parameter estimates in 13 of 18 comparisons (3 parameters, 3 temporal resolutions, 2 ROIs). Moreover, at a temporal resolution of 4 s/frame, the parameter v_p shows a clear distinction between tumor and healthy tissue, which is better preserved at worse temporal resolutions (>4 s/frame) when using the corrected model (Figure 2.9). A similar, although less pronounced, effect was observed in parameters K^{trans} and v_e.

2.4 Discussion and conclusions

This work demonstrates that modern acquisitions that oversample the center of k-space introduce a systematic and random error in the parameter estimation if not properly considered in the modeling. Our correction of the extended Tofts model accounts for this and results in accurate and precise parameter estimates. In simulations, we showed that the proposed model outperforms conventional modeling for estimating all pharmacokinetic parameters, with the largest effect in v_p. Moreover, our simulations showed that, when properly considered, variable density sampling outperforms conventional discrete sampling, especially for lower temporal resolutions. This improved behavior can be attributed to the additional temporal information inherent in time-averaged signals — a feature absent in discrete acquisition. Our in vivo findings support the conclusions drawn from the simulations. Our new model will help the introduction of novel acquisition techniques to perform accurate and precise DCE modeling. Moreover, it suggests that such new acquisition techniques, when properly modeled, may have unanticipated added benefits of better retaining fast (subtemporal resolution) processes, as compared with conventional acquisitions. This behavior was observed in simulations and in vivo, in which parameter maps remained more consistent between reconstruction times.

Building on this thought, it may be hypothesized that for variable-density sampling trajectories, the temporal resolution becomes less important. Ironically, these acquisitions were partially developed to enable higher temporal resolutions. However, if temporal resolution becomes less crucial for accurate pharmacokinetic modeling, we could consider increasing spatial resolution, reducing acceleration factor, and/or including motion robustness (e.g., soft-gating [27–29] typically requires dynamic times that are longer than a respiratory cycle). Longer temporal frames may also loosen the need for temporal regularization; such regularization can affect the accuracy of pharmacokinetic modeling [30].

Additionally, simulated experiments have revealed that the choice of AIF, temporal resolution, and the implementation of the extended Tofts model might influence the esti-

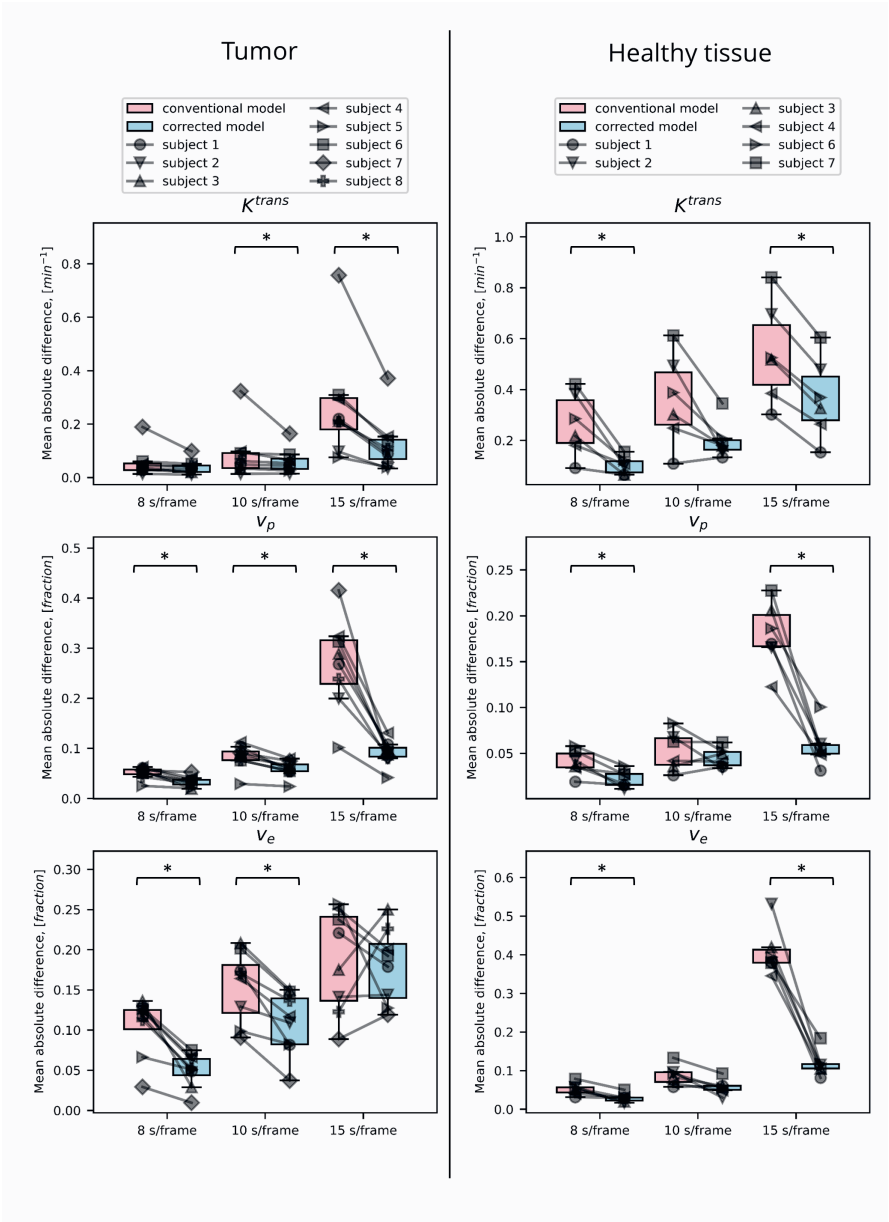

Figure 2.9: The mean absolute difference (MAD) of predicted parameters was obtained at various temporal resolutions compared with the reference acquisition, which had the fastest temporal resolution of 4 s/frame. MAD values are presented for both the tumor (*left*) and healthy tissue (*right*). A lower MAD value indicates a closer agreement between the predicted parameters and the reference, reflecting improved accuracy in parameter estimation.

mation of pharmacokinetic parameters. When using a patient-specific AIF derived from a time-averaged acquisition, Eq. 2.3 suggests that the correction of a pharmacokinetic model is unnecessary. However, Appendix 2.D shows that this only works when taking care of the implementation. In our naïve implementation, discrepancy was caused by two factors: the first being a result of discrete sampling of the AIF, and the second being related to the choice of model implementation. Our implementation has an analytical description of the AIF, which was designed with non-convolved AIF signals in mind. The analytical equation does not have the degrees of freedom to describe the convolved AIF accurately.

Another important point is that, in vivo, a fair comparison of approaches is challenging and would require administering two doses of contrast agent. With no direct patient benefit, this was considered unethical. Even if two doses are administered, the second DCE acquisition would be affected by the first, given that complete contrast-agent elimination takes approximately 48 h. Hence, our only metric in vivo was model consistency. One challenge in this approach is that the "best" modeling is done on high temporal resolutions, which typically contain the most noise. This noise adds to additional uncertainty that may mask some of the benefits. That said, the conducted in vivo experiments agreed with the simulations and supported our equation.

In this work, we chose to build on the computationally convenient analytical representation of the extended Tofts model and managed to implement additional convolution analytically, too. This does come at the cost of working in the concentration domain. One might argue that temporal averaging happens in the signal space and that modeling the averaging in the concentration space might be affected by the nonlinear relation between signal and concentration. Note that in our approach, the global nonlinearity (e.g., nonlinear effect occurring between precontrast and max contrast) was corrected for by converting the dynamic signal to the concentration curve with the spoiled gradient-echo signal model and considering the specific relaxivities of the paramagnetic agent. However, local nonlinear effects, occurring during the averaging of one dynamic, may still play a role. In Appendix 2.E, following Quantitative Imaging Biomarkers Alliance–recommended guidelines [31], we show that these nonlinear effects do not significantly affect the accuracy when accounting for temporal averaging in concentration space. That said, an alternative approach may be to fit directly in the signal domain, where the time averaging is applied by numerical convolution of the signal function with a rectangular function, as follows:

$$S'(t) = \frac{1}{\Delta t} rect(\frac{t}{\Delta t}) * S(t), \tag{2.5}$$

where $S'(t)$ is the convolved signal; Δt is the measuring time per dynamic; and $S(t)$ is the signal equation for data following the extended Tofts model (see Appendix 2.E for details). This numerical convolution can be made arbitrarily accurate by numerically sampling $S(t)$ with small steps (e.g., we chose 0.1 s) and fitted to the actual signal. Although this equation better represents the imaging physics, it loses its analytical character. In Appendix 2.E, we see that this approach reaches very comparable performance as convolving in the concentration domain.

Finally, in vivo DCE-MRI signals typically exhibit Rician noise, which when converted to concentration space can be more complex. With the lack of a noise model in the concentration domain, and to avoid additional biases of Rician noise, we implemented Gaussian noise in our simulations. This choice allows us to focus on the models' accuracy without additional noise-related biases skewing parameter estimations and obscuring the models' performances. Additionally, for SNR values above 3, Rician noise can be well-approximated by Gaussian noise [32]. In our simulations, the average SNR in the concentration domain was 20, and we do not expect Rician noise to cause a bias in that case. Although the SNR precontrast would be 0 in the contrast-concentration domain, for most tissues it is nonzero in the signal domain (in which noise is Rician).

2.4.1 Limitations

One drawback is the limited amount of controlled experimental validation of the new equation. Unfortunately, we do not have an established DCE phantom that models the dynamic microvasculature behavior, and we had no ground-truth in vivo references. However, the experiments that we performed do support the equation. First, the equations are derived using basic imaging concepts and hence should be correct. Second, Eq. 2.3 was validated using open source–tested software. Third, we validated it in vivo, where the model was consistent over sampling rates.

Another shortcoming lies in our use of compressed sensing - a technique designed for the reconstruction of undersampled k-space. Its temporal regularization may introduce blurring to the temporal signal, resulting in some inaccuracy in our suggested equation. It would be interesting to explore the possibility of accounting for this effect by incorporating it as an additional blurring factor in Eq. 2.2 in future work. However, as we can now model longer dynamic scan times more accurately, potentially, temporal regularization is less necessary.

2.4.2 Conclusions

In conclusion, temporally averaged signals acquired with variable density trajectories oversampling the k-space center require modification of pharmacokinetic modeling to obtain accurate and precise parameter estimates, especially at low temporal resolutions. Our proposed correction showed improved estimates of pharmacokinetic parameters in simulations and in vivo and can be directly adapted. The correction will help with the implementation of modern variable density acquisition techniques as a tool for DCE-MRI.

Chapter appendix

2.A Proof of Eq. 2.3

$$C'(t,u,C_p(t)) = \frac{1}{\Delta t}C(t,u,C_p(t)) * rect(\frac{t}{\Delta t}) \\ = C(t,u,\frac{1}{\Delta t}C_p(t) * rect(\frac{t}{\Delta t})) = C(t,u,C'_p(t)) \quad (2.6)$$

If we consider:

$$f = rect(\frac{t}{\Delta t}), \quad (2.7)$$

$$g = C_p(t,\theta_{AIF}), \quad (2.8)$$

$$h = e^{-\frac{K^{trans}}{v_e}t}, \quad (2.9)$$

then,

$$C'(t,u,C_p) = \frac{1}{\Delta t}C(t,u,C_p(t)) * rect(\frac{t}{\Delta t}) \\ = \frac{1}{\Delta t}(v_p g) * f + \frac{1}{\Delta t}K^{trans}(g*h)*f \\ = v_p(\frac{1}{\Delta t}g*f) + K^{trans}(\frac{1}{\Delta t}g*f)*h \\ = C(t,u,C_p(t) * rect(\frac{t}{\Delta t})) = C(t,u,C'_p(t)) \quad (2.10)$$

2.B Definition of analytical AIF proposed by Orton et al.

The analytical arterial input function earlier proposed by M. R. Orton et al [16]:

$$C_p(t,\theta_{AIF}) = \begin{cases} \alpha_B(1-\cos(\mu_B t)) + \alpha_B\alpha_G f(t,\mu_G) & \text{for } 0 \le t \le t_B \\ \alpha_B\alpha_G f(t_B,\mu_G)e^{-\mu_G(t-t_B)} & \text{for } t > t_B \end{cases}, \quad (2.11)$$

where $\theta_{AIF} = \{\alpha_B, \alpha_G, \mu_B, \mu_G\}, t_B = 2\pi\mu_B^{-1}$ and

$$f(t,\alpha) = \frac{1}{\alpha}(1 - e^{-\alpha t}) - \frac{1}{\alpha^2 + \mu_B^2}(\alpha\cos(\mu_B t) + \mu_B \sin(\mu_B t) - \alpha e^{-\alpha t}). \tag{2.12}$$

Parameters $\theta_{AIF} = \{\alpha_B, \alpha_G, \mu_B, \mu_G\}$ define the shape of arterial input function and are usually obtained by fitting this analytical model $C_p(t, \theta_{AIF})$ to the DCE signal.

2.C Definition of time-averaged AIF

Using Eqs. 2.4, 2.11, the analytical input function corrected for the time-averaging is

$$\begin{aligned} C'_p(t,\theta_{AIF}) = {} & \frac{\alpha_B}{\Delta t}(f(t+\frac{\Delta t}{2},0) - f(t-\frac{\Delta t}{2},0)) + \\ & \frac{\alpha_B\alpha_G}{\Delta t\mu_G}(f(t+\frac{\Delta t}{2},0) - f(t-\frac{\Delta t}{2},0) - f(t+\frac{\Delta t}{2},\mu_G) + f(t-\frac{\Delta t}{2},\mu_G)) \end{aligned} \tag{2.13}$$

for $0 \leq t \leq t_B$, and it is

$$\begin{aligned} C'_p(t,\theta_{AIF}) = {} & \frac{\alpha_B}{\Delta t}(f(t_B+\frac{\Delta t}{2},0) - f(t_B-\frac{\Delta t}{2},0)) + \\ & \frac{\alpha_B\alpha_G}{\Delta t\mu_G}(f(t_B+\frac{\Delta t}{2},0) - f(t_B-\frac{\Delta t}{2},0) - \\ & f(t_B+\frac{\Delta t}{2},\mu_G)e^{-\mu_G(t+\frac{\Delta t}{2}-t_B)} + f(t_B-\frac{\Delta t}{2},\mu_G)e^{-\mu_G(t-\frac{\Delta t}{2}-t_B)}) \end{aligned} \tag{2.14}$$

for $t > t_B$.

The derivation of this equation relies on two key principles. First. the rectangular function in Eq. 2.4 can be represented as the sum of two Heaviside step functions, each shifted relative to the other along the temporal axis by Δt. Second, a Heaviside function can be expressed as an exponential function with the coefficient of zero.

2.D Results of repeated Experiments 1 and 2 with time-averaged AIF

Our intention in the following experiments was to put Eq. 2.3 into practical use and show that one does not need to consider the convolution if a patient-specific AIF is measured. Therefore, Experiments 1 and 2 were modified by using a temporally averaged (patient-specific) AIF C'_p instead of discretely sampled AIF. We conducted experiments by perform-

ing two fits: one with the conventional model (Experiment 1 with time-averaged AIF) and one with the corrected model (Experiment 2 with time-averaged AIF), both fitting a convolved model to the AIF as using a convolved model for the signal.

Figures 2.D.1, 2.D.2, 2.D.3 present estimated parameters K^{trans}, v_p, v_e obtained from fitting both models. Despite the initial expectation based on Eq. 2.3 that, with a patient-specific AIF C'_p, the conventional and corrected models would yield similar results, our calculations showed the opposite: parameters estimated with the conventional model were less accurate and precise.

There were two factors contributing to a poor translation to practice. The first factor is that we were solving an inverse problem where parameters were fitted to discretely measured data. While the direct computation of Eq. 2.3 is straightforward, solving its inverse problem is more complex due to the discretization, the choice of sampling intervals, and the fitting process itself, all of which can introduce errors. The second factor is that in these experiments, we are using an analytical description of the AIF [16], which was developed with a non-convolved signal in mind. The analytical equation does not have the degrees of freedom to describe the convolved AIF accurately.

An alternative implementation of the model such as LEK_UoEdinburgh from the OSIPI GitHub [17], that does not parametrize the AIF could be the solution. However, in preliminary experiments (data not shown), this model still showed a discrepancy at low temporal resolutions. The numerical integration required for unparameterized AIF is less accurate at lower temporal resolutions. This leads to underestimation of the concentration curves and, consequently, a poor prediction of the pharmacokinetic parameters.

To further verify Eq. 2.3 as accurately as possible, we mimicked the view-sharing reconstruction approach, where the virtually acquired temporal resolution of 15 s/frame (to mimic temporal blurring) was much lower than the virtual temporal sampling of 0.1 s/frame (to ensure sufficient sampling). The simulated data $C'(t, u, C_p)$ (with the equations from Appendix 2.C) represented the left side of Eq. 2.3, while the fitted model $C(t, u, C'_p(t, \theta_{AIF}))$ represented the right side of Eq. 2.3. Therefore, Eq. 2.3 was considered to be valid if the fit adequately described the data. To remove any assumptions on the shape of the AIF, we used the extended Tofts implementation from LEK_UoEdinburgh from the OSIPI GitHub [17]. Figure 2.D.4 shows the results of this experiment. From these findings, we established a good resemblance between predicted parameters and the ground truth, validating Eq. 2.3.

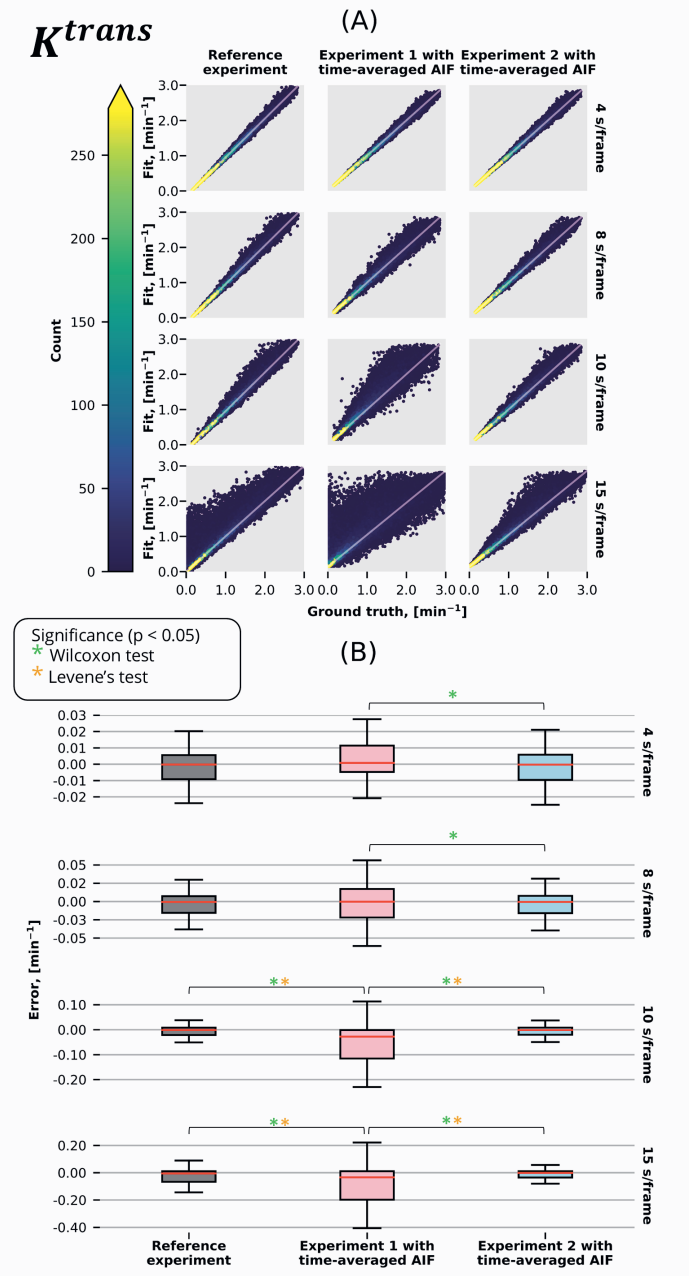

Figure 2.D.1: The K^{trans} results of the experiments at various temporal resolutions. A. The graph displays the fitted value of the parameter K^{trans} plotted against the ground truth. The colour gradient represents the density of points, while the white diagonal line represents the function Fit=Ground Truth. B. The distribution of errors in the predicted parameter K^{trans} for the same experiments and temporal resolutions. Significant differences, as identified by the Wilcoxon (accuracy) and Lavene (precision) tests ($p < 0.05$), are denoted by green and orange stars, respectively.

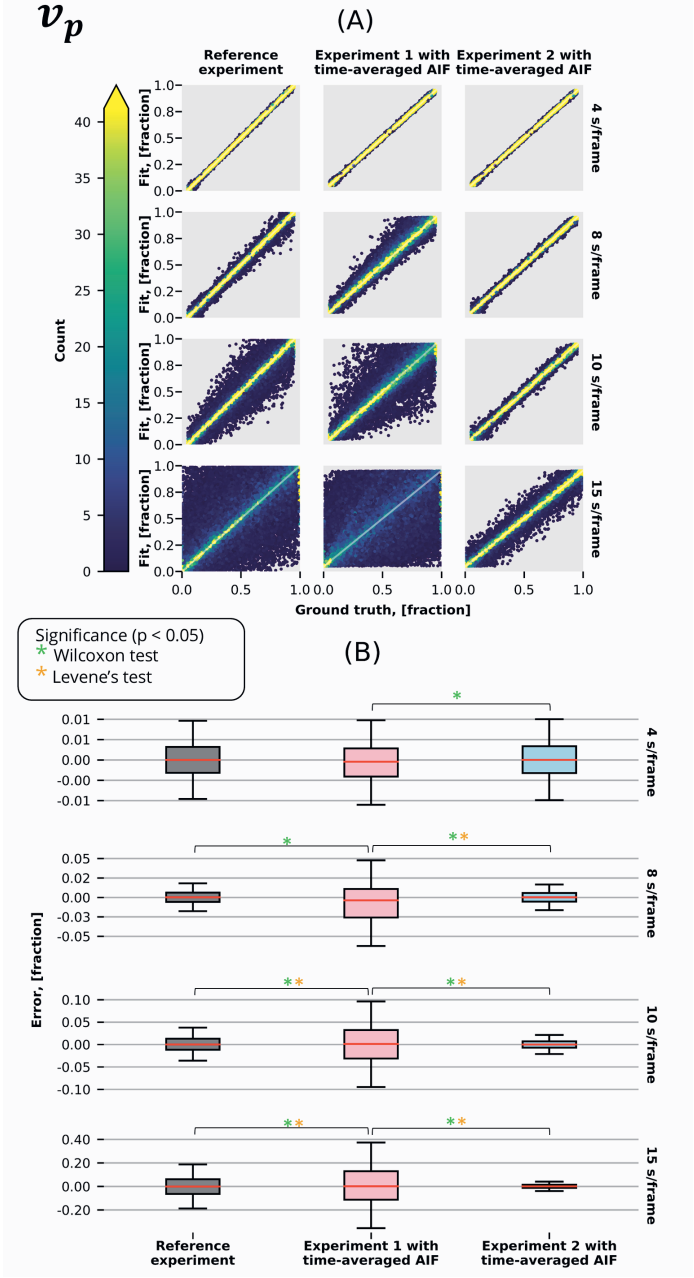

Figure 2.D.2: The v_p results of the experiments at various temporal resolutions. A. The graph displays the fitted value of the parameter v_p plotted against the ground truth. The colour gradient represents the density of points, while the white diagonal line represents the function Fit=Ground Truth. B. The distribution of errors in the predicted parameter v_p for the same experiments and temporal resolutions. Significant differences, as identified by the Wilcoxon (accuracy) and Lavene (precision) tests (p < 0.05), are denoted by green and orange stars, respectively.

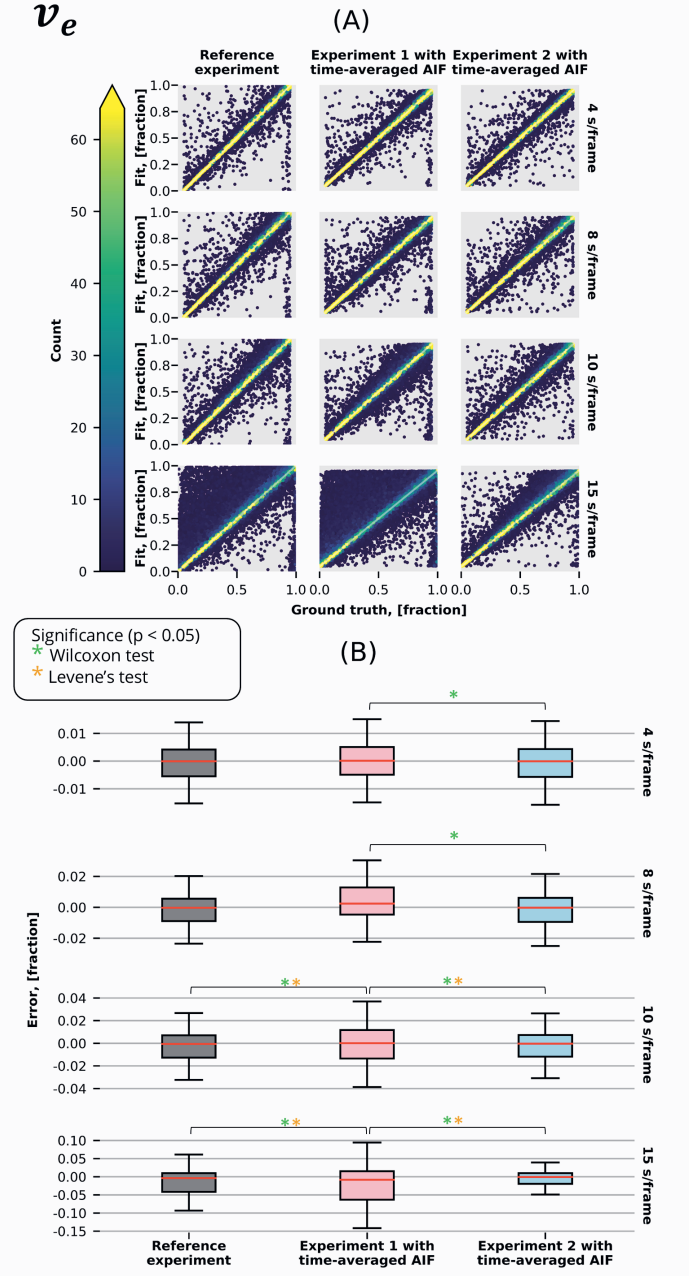

Figure 2.D.3: The v_e results of the experiments at various temporal resolutions. A. The graph displays the fitted value of the parameter v_e plotted against the ground truth. The colour gradient represents the density of points, while the white diagonal line represents the function Fit=Ground Truth. B. The distribution of errors in the predicted parameter v_e for the same experiments and temporal resolutions. Significant differences, as identified by the Wilcoxon (accuracy) and Lavene (precision) tests (p < 0.05), are denoted by green and orange stars, respectively.

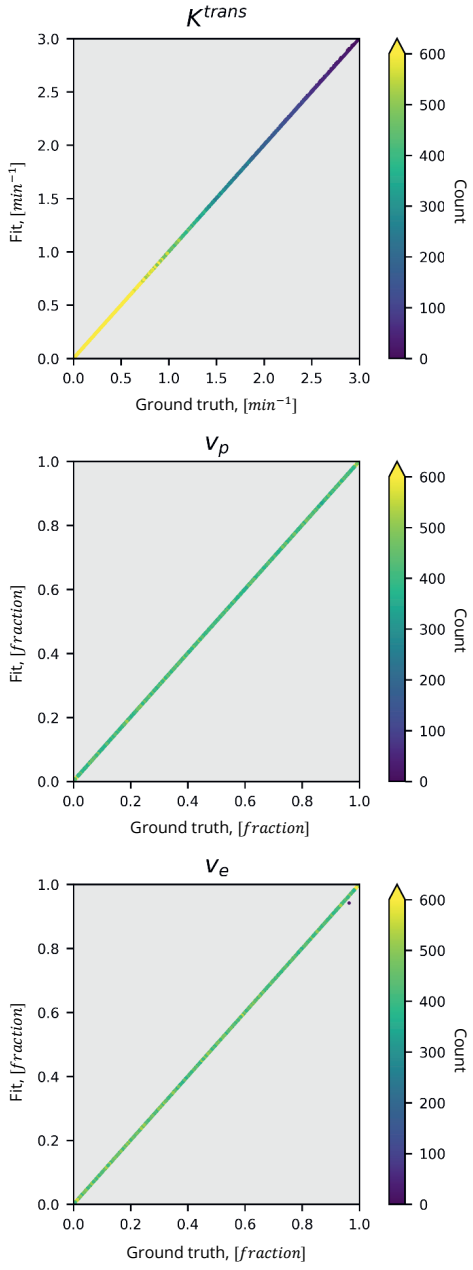

Figure 2.D.4: The results of the experiment validating Equation 2.3. The graphs display the fitted value of the parameters K^{trans}, v_p, v_e plotted against the ground truth. The color gradient represents the density of points at each location.

2.E Comparison of time-averaging implmentations in the concentration and signal domains

To investigate the potential impact of the non-linearities in the signal-concentration dependency, we performed additional simulations. Generally, a linear relation between signal and concentration happens either at larger flip angles or low values of concentration. According to the most recent recommendations from the Quantitative Imaging Biomarker Alliance (QIBA), the recommended flip angle for DCE imaging at 3T ranges between 10-15° for brain and prostate imaging and 10-30° for breast imaging. In alignment with these guidelines, our in vivo data for pancreas imaging was acquired at a flip angle of 25°. This is at the high-end of the flip angle ranges and hence if there is no impact of non-linearity here, one can conclude it is safe to work in the concentration domain.

We examined whether the implementation of time-averaging fit in concentration space results in less accurate and/or precise estimates due to local non-linearities. To test this, we simulated 50,000 signal curves for flip angles of 10°, 15°, and 25° (with T_1 =725 ms, corresponding to pancreas tissue). To achieve this, we first generated concentration curves $C(t, u, C_p)$ every 0.1 seconds using the conventional extended Tofts model and converted these to the effective relaxation time $T_1(t, u, C_p, T_{10})$:

$$T_1(t, u, C_p, T_{10}) = \frac{1}{rC(t, u, C_p) + \frac{1}{T_{10}}}, \tag{2.15}$$

where r $[L/mmol\ s]$ is the relaxivity of the MR contrast agent, and T_{10} is the baseline T_1 value. We then used a signal equation to calculate signal curves. Although, in general, the signal equation depends on the DCE acquisition, here we used the fast gradient echo equation:

$$S(t, u, C_p, T_{10}, M_o, \alpha, TR) = M_0 \frac{sin\alpha(1 - e^{-\frac{TR}{T_1(t,u,C_p,T_{10})}})}{1 - cos\alpha\, e^{-\frac{TR}{T_1(t,u,C_p,T_{10})}}}, \tag{2.16}$$

where M_0 corresponds to spin density, α is the flip angle, and TR is the repetition time. Note that when fitting u to the signal curve, an additional parameter M_0 must also be fitted.

We subsequently performed signal averaging $S'(t, u, C_p, T_{10}, M_o, \alpha, TR)$ by convolving it with a rectangular function:

$$S'(t, u, C_p, T_{10}, M_o, \alpha, TR) = \frac{1}{\Delta t} rect(\frac{t}{\Delta t}) * S(t, u, C_p, T_{10}, M_o, \alpha, TR). \tag{2.17}$$

Using this approach, we simulated different temporal resolutions consistent with those used in the manuscript: 4, 8, 10, and 15 s/frame. Additionally, to ensure accurate estimation of M_0 from the fitting, we prolonged the pre-contrast phase to 10 acquisitions (0.7 - 2.5 min). This was especially important for the 15 s/frame acquisition. In practice, longer acquisition times in vivo will have higher SNR, reducing the potential issues with M_0 estimation. We added noise to the signal curves that corresponded to an average SNR of 60 in the signal domain (note noise levels in the manuscript were in the contrast concentration domain).

We performed fitting using three models:

1. Conventional fitting in the signal domain S.
2. Time-averaged fitting in the signal domain S'.
3. Time-averaged fitting in the concentration domain C'.
    - Here, concentration curves were calculated from the simulated signal curves prior to the fit by inverting of Eqs. 2.15 and 2.16.

We compared these three fit approaches by analyzing their resulting error distributions.

The results (Figures 2.E.1, 2.E.2, 2.E.3) clearly demonstrate that time-averaging fitting in the signal domain and time-averaging in concentration domain perform similarly across all three flip angles. Furthermore, both time-averaged models produced more accurate and more precise pharmacokinetic parameters than the conventional model across all flip angles. The low flip angle (10°) did result in similar performance between conventional and time-averaged models at lower temporal resolutions (up to 10 s/frame). This could be a result of overall lower sensitivity due to signal saturation for low flip angles masking additional effect of averaging.

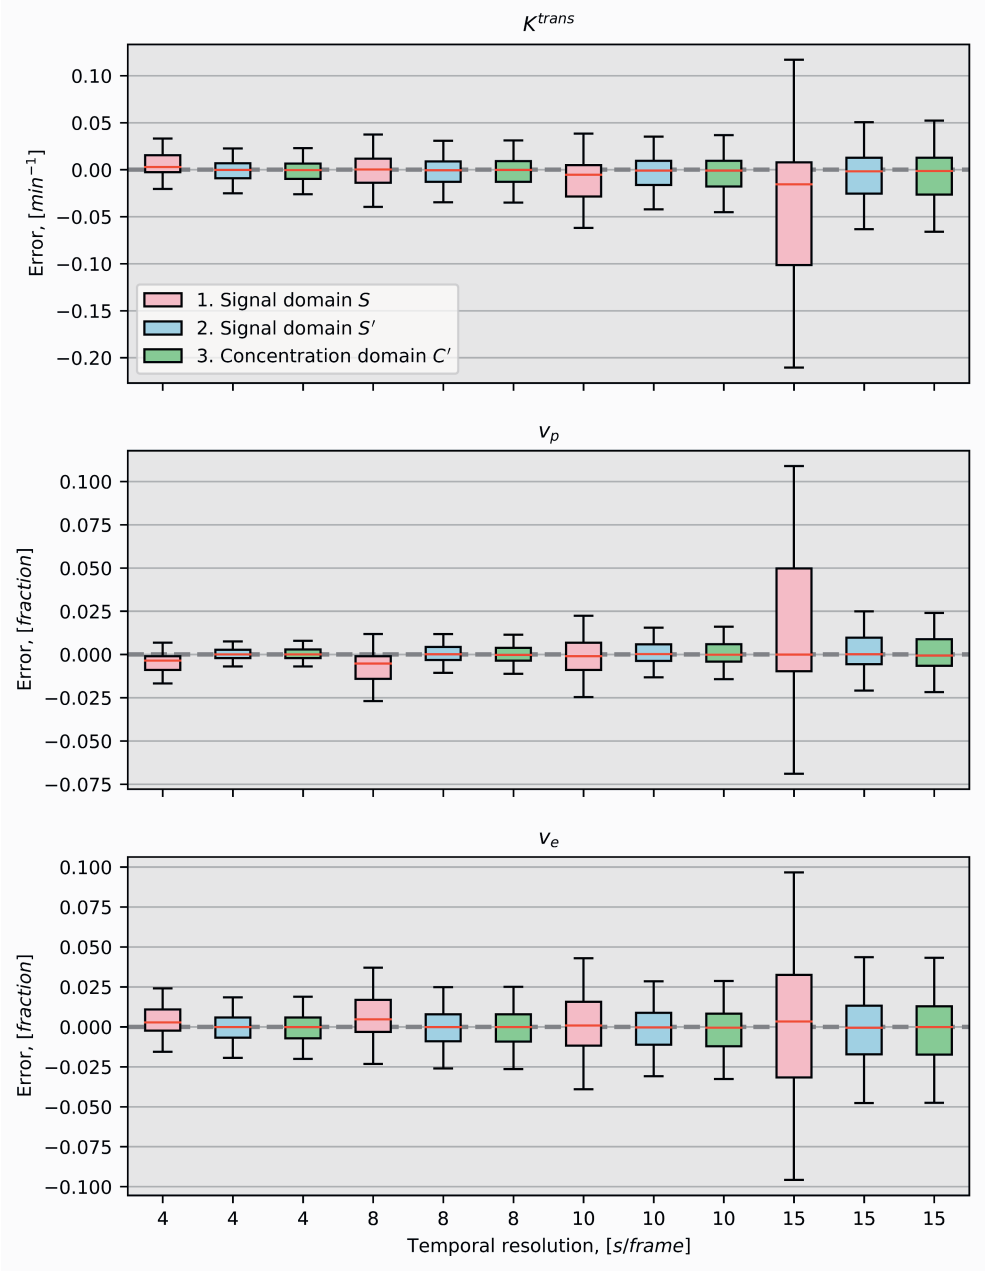

Figure 2.E.1: Error in the estimation of parameters K^{trans} (top), v_p (middle), v_e (bottom) with the conventional fitting in the signal domain (pink), with the time-averaged fitting in the signal domain (blue), and with the time-averaged fitting in the concentration domain. The simulated signal was calculated with the **flip angle of 25°** and temporal resolutions of 4, 8, 10, and 15 s/frame.

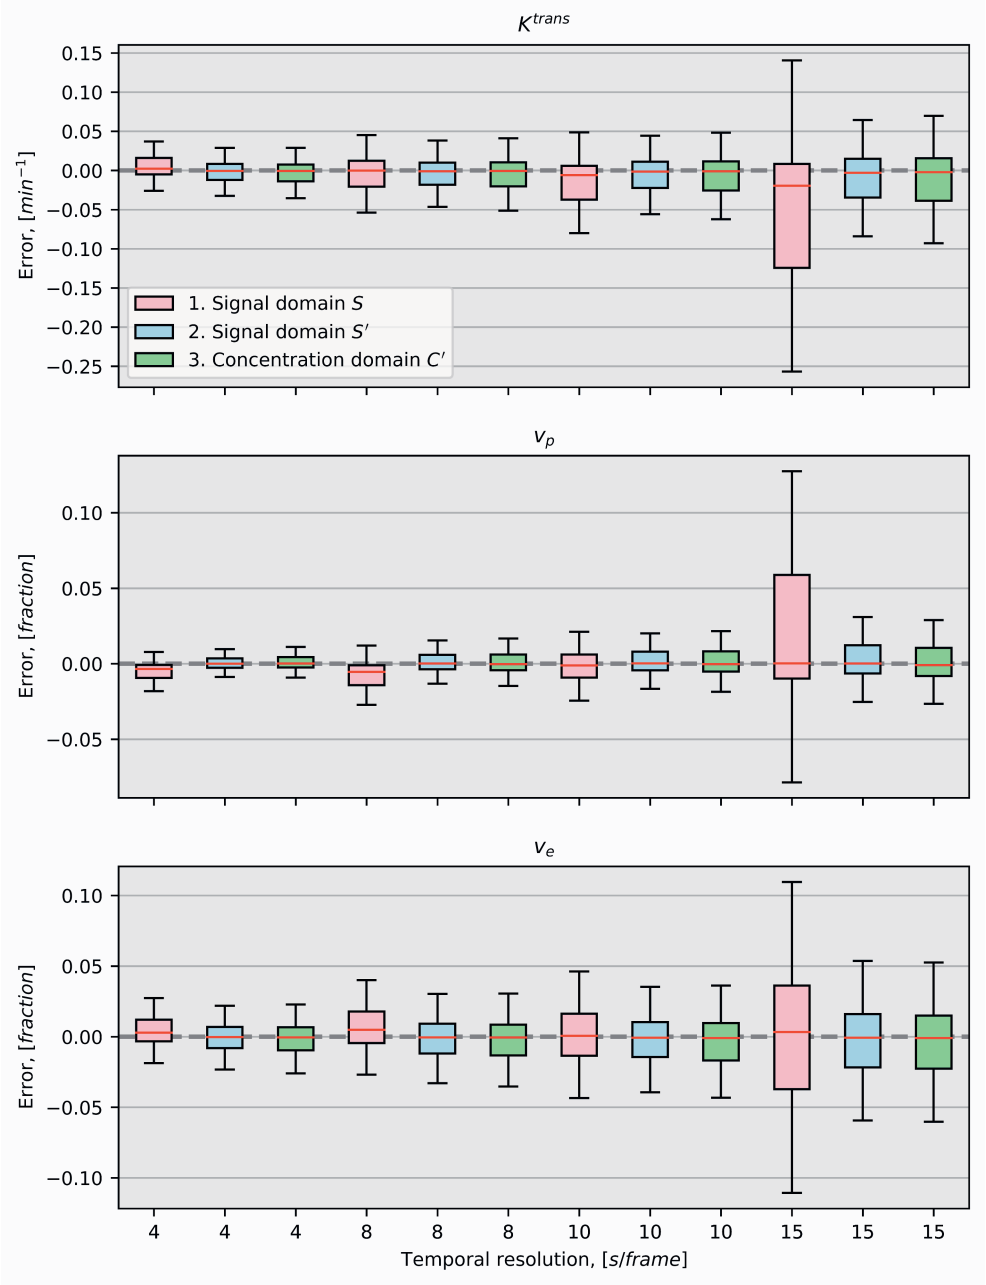

Figure 2.E.2: Error in the estimation of parameters K^{trans} (top), v_p (middle), v_e (bottom) with the conventional fitting in the signal domain (pink), with the time-averaged fitting in the signal domain (blue), and with the time-averaged fitting in the concentration domain. The simulated signal was calculated with the **flip angle of 15°** and temporal resolutions of 4, 8, 10, and 15 s/frame.

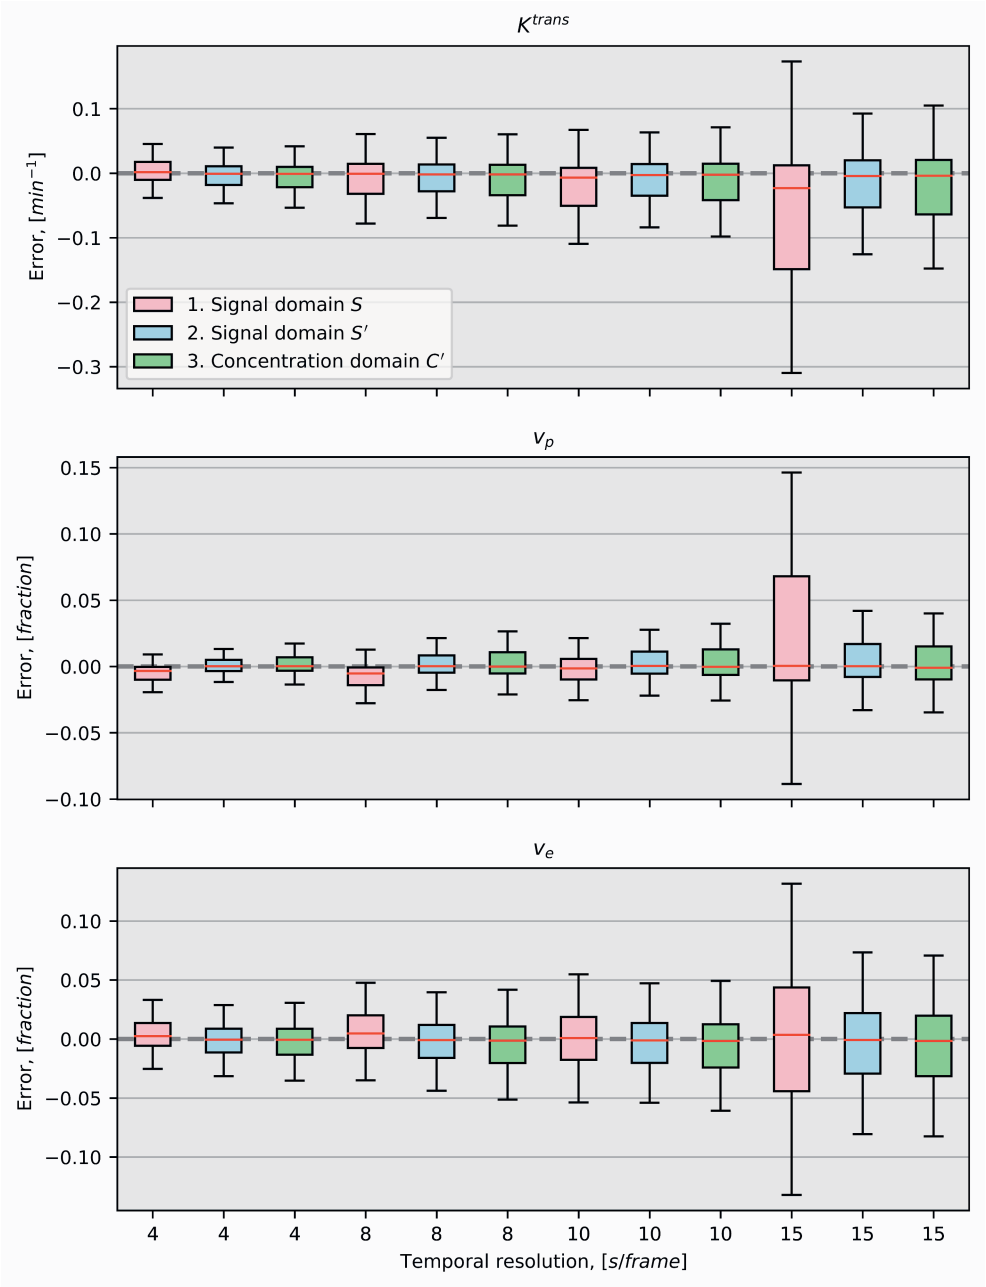

Figure 2.E.3: Error in the estimation of parameters K^{trans} (top), v_p (middle), v_e (bottom) with the conventional fitting in the signal domain (pink), with the time-averaged fitting in the signal domain (blue), and with the time-averaged fitting in the concentration domain. The simulated signal was calculated with the **flip angle of 10°** and temporal resolutions of 4, 8, 10, and 15 s/frame.

Bibliography

[1] Y. Cao, "The promise of dynamic contrast-enhanced imaging in radiation therapy," in *Seminars in radiation oncology*, vol. 21, no. 2. Elsevier, 2011, pp. 147–156.

[2] Y. S. Sung, B. Park, Y. Choi, H.-S. Lim, D.-C. Woo, K. W. Kim, and J. K. Kim, "Dynamic contrast-enhanced mri for oncology drug development," *Journal of Magnetic Resonance Imaging*, vol. 44, no. 2, pp. 251–264, 2016.

[3] R. Raja, G. A. Rosenberg, and A. Caprihan, "Mri measurements of blood-brain barrier function in dementia: a review of recent studies," *Neuropharmacology*, vol. 134, pp. 259–271, 2018.

[4] M. J. Thrippleton, W. H. Backes, S. Sourbron, M. Ingrisch, M. J. van Osch, M. Dichgans, F. Fazekas, S. Ropele, R. Frayne, R. J. van Oostenbrugge *et al.*, "Quantifying blood-brain barrier leakage in small vessel disease: review and consensus recommendations," *Alzheimer's & Dementia*, vol. 15, no. 6, pp. 840–858, 2019.

[5] M. Duering, G. J. Biessels, A. Brodtmann, C. Chen, C. Cordonnier, F.-E. de Leeuw, S. Debette, R. Frayne, E. Jouvent, N. S. Rost *et al.*, "Neuroimaging standards for research into small vessel disease—advances since 2013," *The Lancet Neurology*, vol. 22, no. 7, pp. 602–618, 2023.

[6] S. P. Sourbron and D. L. Buckley, "Classic models for dynamic contrast-enhanced mri," *NMR in Biomedicine*, vol. 26, no. 8, pp. 1004–1027, 2013.

[7] P. S. Tofts, G. Brix, D. L. Buckley, J. L. Evelhoch, E. Henderson, M. V. Knopp, H. B. Larsson, T.-Y. Lee, N. A. Mayr, G. J. Parker *et al.*, "Estimating kinetic parameters from dynamic contrast-enhanced t1-weighted mri of a diffusable tracer: standardized quantities and symbols," *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 10, no. 3, pp. 223–232, 1999.

[8] P. S. Tofts, "Modeling tracer kinetics in dynamic gd-dtpa mr imaging," *Journal of magnetic resonance imaging*, vol. 7, no. 1, pp. 91–101, 1997.

[9] M. Heilmann, F. Kiessling, M. Enderlin, and L. R. Schad, "Determination of pharmacokinetic parameters in dce mri: consequence of nonlinearity between contrast

agent concentration and signal intensity," *Investigative radiology*, vol. 41, no. 6, pp. 536–543, 2006.

[10] E. Henderson, B. K. Rutt, and T.-Y. Lee, "Temporal sampling requirements for the tracer kinetics modeling of breast disease," *Magnetic resonance imaging*, vol. 16, no. 9, pp. 1057–1073, 1998.

[11] L. Feng, L. Axel, H. Chandarana, K. T. Block, D. K. Sodickson, and R. Otazo, "Xd-grasp: golden-angle radial mri with reconstruction of extra motion-state dimensions using compressed sensing," *Magnetic resonance in medicine*, vol. 75, no. 2, pp. 775–788, 2016.

[12] R. Merton, D. Bosshardt, G. J. Strijkers, A. J. Nederveen, E. M. Schrauben, and P. van Ooij, "Reproducibility of 3d thoracic aortic displacement from 3d cine balanced ssfp at 3 t without contrast enhancement," *Magnetic Resonance in Medicine*, vol. 91, no. 2, pp. 466–480, 2024.

[13] N. P. Wassenaar, O. J. Gurney-Champion, A.-S. van Schelt, T. Bruijnen, H. W. van Laarhoven, J. Stoker, A. J. Nederveen, J. H. Runge, and E. M. Schrauben, "Optimizing pseudo-spiral sampling for abdominal dce mri using a digital anthropomorphic phantom," *Magnetic Resonance in Medicine*, vol. 92, no. 5, pp. 2051–2064, 2024.

[14] J. G. Pipe, "Motion correction with propeller mri: application to head motion and free-breathing cardiac imaging," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 963–969, 1999.

[15] M. Rata, D. J. Collins, J. Darcy, C. Messiou, N. Tunariu, N. Desouza, H. Young, M. O. Leach, and M. R. Orton, "Assessment of repeatability and treatment response in early phase clinical trials using dce-mri: comparison of parametric analysis using mr-and ct-derived arterial input functions," *European radiology*, vol. 26, no. 7, pp. 1991–1998, 2016.

[16] M. R. Orton, J. A. d'Arcy, S. Walker-Samuel, D. J. Hawkes, D. Atkinson, D. J. Collins, and M. O. Leach, "Computationally efficient vascular input function models for quantitative kinetic modelling using dce-mri," *Physics in Medicine & Biology*, vol. 53, no. 5, pp. 1225–1239, 2008.

[17] P. J. van Houdt, S. Ragunathan, M. Berks, Z. Ahmed, L. E. Kershaw, O. J. Gurney-Champion, S. Tadimalla, J. Arvidsson, Y. Sun, J. Kallehauge *et al.*, "Contrast-agent-based perfusion mri code repository and testing framework: Ismrm open science

initiative for perfusion imaging (osipi)," *Magnetic resonance in medicine*, vol. 91, no. 5, pp. 1774–1786, 2024.

[18] M. D. Holland, A. Morales, S. Simmons, B. Smith, S. R. Misko, X. Jiang, D. A. Hormuth, C. Christenson, R. P. Koomullil, D. E. Morgan *et al.*, "Disposable point-of-care portable perfusion phantom for quantitative dce-mri," *Medical physics*, vol. 49, no. 1, pp. 271–281, 2022.

[19] G. J. Parker, C. Roberts, A. Macdonald, G. A. Buonaccorsi, S. Cheung, D. L. Buckley, A. Jackson, Y. Watson, K. Davies, and G. C. Jayson, "Experimentally-derived functional form for a population-averaged high-temporal-resolution arterial input function for dynamic contrast-enhanced mri," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 56, no. 5, pp. 993–1000, 2006.

[20] L. M. Gottwald, E. S. Peper, Q. Zhang, B. F. Coolen, G. J. Strijkers, A. J. Nederveen, and P. van Ooij, "Pseudo-spiral sampling and compressed sensing reconstruction provides flexibility of temporal resolution in accelerated aortic 4d flow mri: A comparison with k-t principal component analysis," *NMR in Biomedicine*, vol. 33, no. 4, p. e4255, 2020.

[21] J. Schoormans, G. J. Strijkers, A. C. Hansen, A. J. Nederveen, and B. F. Coolen, "Compressed sensing mri with variable density averaging (cs-vda) outperforms full sampling at low snr," *Physics in Medicine & Biology*, vol. 65, no. 4, p. 045004, 2020.

[22] L. Feng, "Golden-angle radial mri: basics, advances, and applications," *Journal of Magnetic Resonance Imaging*, vol. 56, no. 1, pp. 45–62, 2022.

[23] K. L. Wright, J. I. Hamilton, M. A. Griswold, V. Gulani, and N. Seiberlich, "Non-cartesian parallel imaging reconstruction," *Journal of Magnetic Resonance Imaging*, vol. 40, no. 5, pp. 1022–1040, 2014.

[24] D. Atkinson, D. L. Hill, P. N. Stoyle, P. E. Summers, and S. F. Keevil, "Automatic correction of motion artifacts in magnetic resonance images using an entropy focus criterion," *IEEE Transactions on Medical imaging*, vol. 16, no. 6, pp. 903–910, 1997.

[25] C. M. De Bazelaire, G. D. Duhamel, N. M. Rofsky, and D. C. Alsop, "Mr imaging relaxation times of abdominal and pelvic tissues measured in vivo at 3.0 t: preliminary results," *Radiology*, vol. 230, no. 3, pp. 652–659, 2004.

[26] H. Lu, C. Clingman, X. Golay, and P. C. Van Zijl, "Determining the longitudinal relaxation time (t1) of blood at 3.0 tesla," *Magnetic Resonance in Medicine: an Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 52, no. 3, pp. 679–682, 2004.

[27] J. Y. Cheng, T. Zhang, N. Ruangwattanapaisarn, M. T. Alley, M. Uecker, J. M. Pauly, M. Lustig, and S. S. Vasanawala, "Free-breathing pediatric mri with nonrigid motion correction and acceleration," *Journal of Magnetic Resonance Imaging*, vol. 42, no. 2, pp. 407–420, 2015.

[28] K. M. Johnson, W. F. Block, S. B. Reeder, and A. Samsonov, "Improved least squares mr image reconstruction using estimates of k-space data consistency," *Magnetic resonance in medicine*, vol. 67, no. 6, pp. 1600–1608, 2012.

[29] T. Zhang, J. Y. Cheng, A. G. Potnick, R. A. Barth, M. T. Alley, M. Uecker, M. Lustig, J. M. Pauly, and S. S. Vasanawala, "Fast pediatric 3d free-breathing abdominal dynamic contrast enhanced mri with high spatiotemporal resolution," *Journal of Magnetic Resonance Imaging*, vol. 41, no. 2, pp. 460–473, 2015.

[30] G. Dong, M. Flaschel, M. Hintermüller, K. Papafitsoros, C. Sirotenko, and K. Tabelow, "Data-driven methods for quantitative imaging," *GAMM-Mitteilungen*, vol. 48, no. 1, p. e202470014, 2025.

[31] A. Shukla-Dave, N. A. Obuchowski, T. L. Chenevert, S. Jambawalikar, L. H. Schwartz, D. Malyarenko, W. Huang, S. M. Noworolski, R. J. Young, M. S. Shiroishi *et al.*, "Quantitative imaging biomarkers alliance (qiba) recommendations for improved precision of dwi and dce-mri derived biomarkers in multicenter oncology trials," *Journal of Magnetic Resonance Imaging*, vol. 49, no. 7, pp. e101–e121, 2019.

[32] H. Gudbjartsson and S. Patz, "The rician distribution of noisy mri data," *Magnetic resonance in medicine*, vol. 34, no. 6, pp. 910–914, 1995.

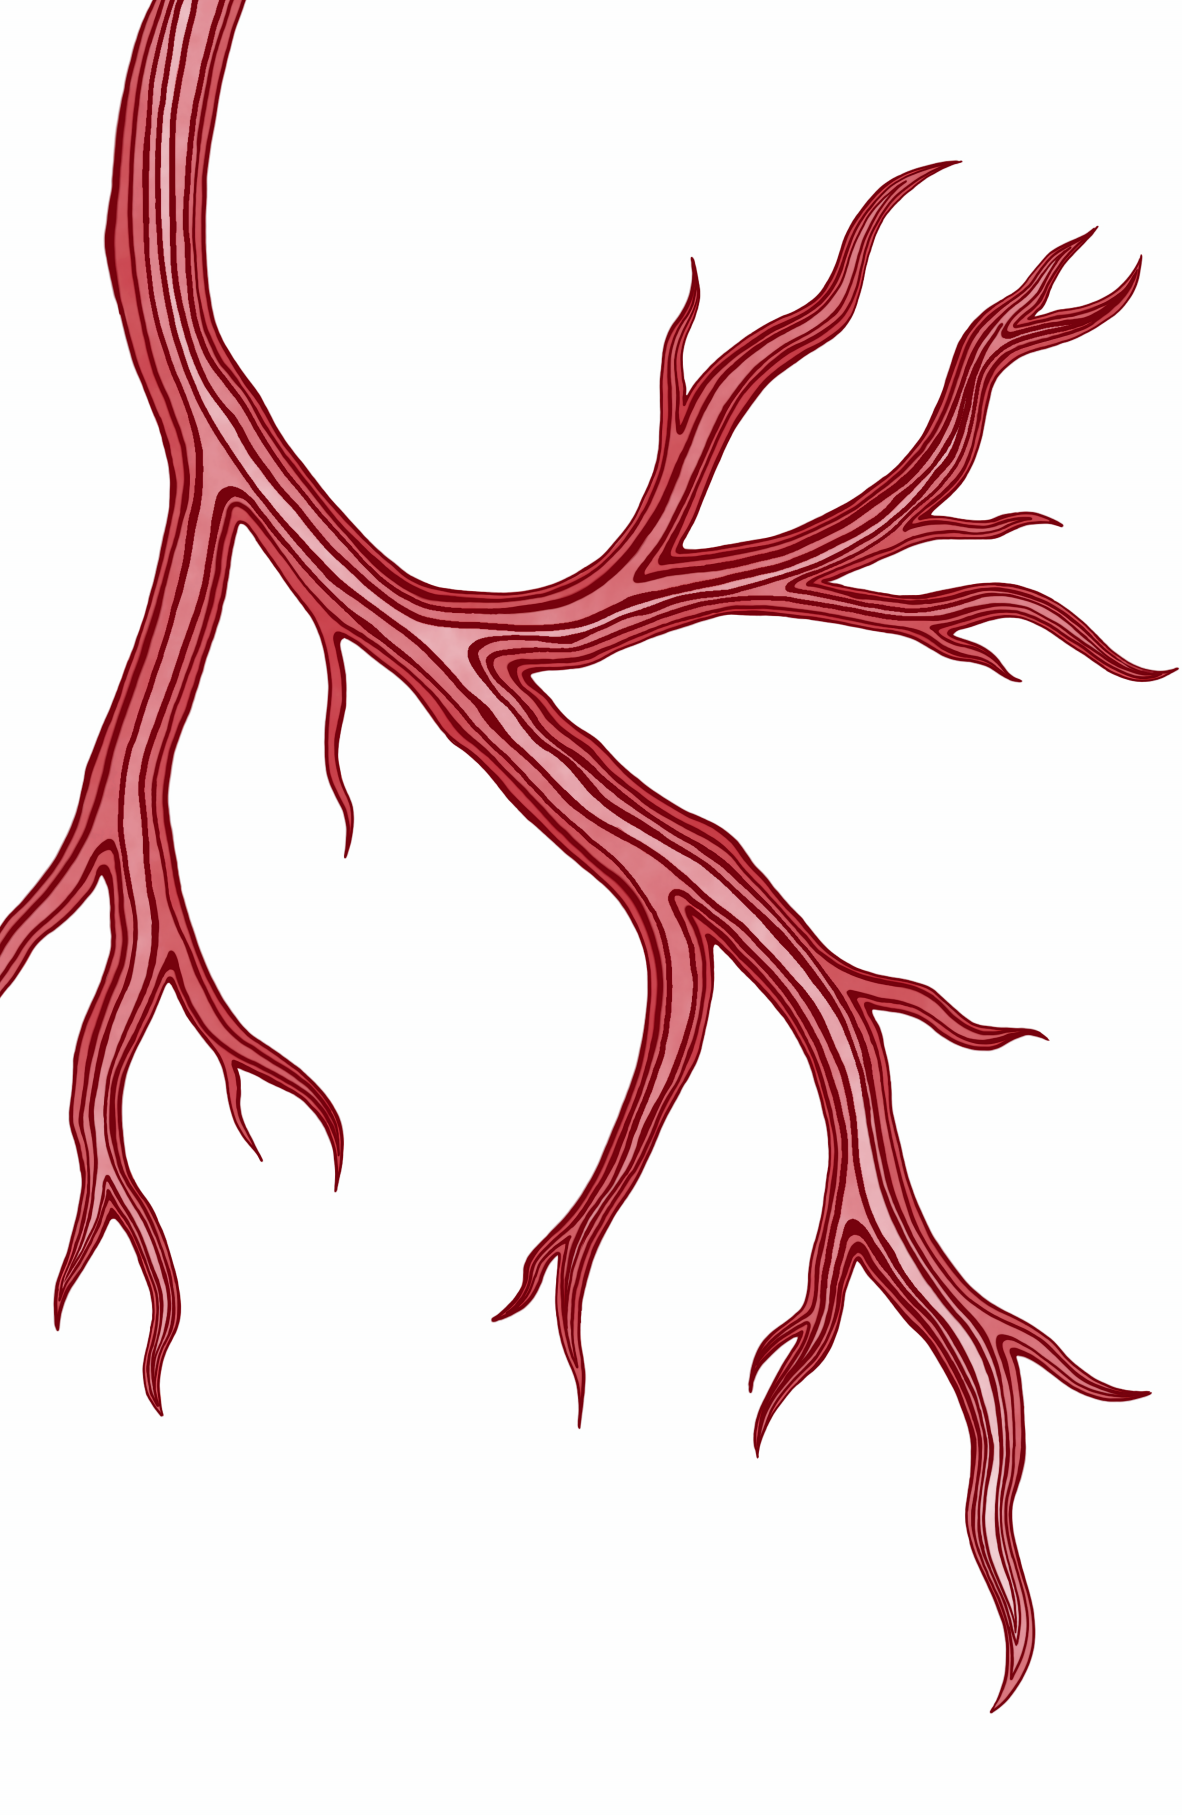

3

Respiratory motion–corrected reconstruction of abdominal dynamic contrast-enhanced MRI

Authors: Natalia V. Korobova, Nienke P. M. Wassenaar, Eric M. Schrauben, Oliver J. Gurney-Champion

Submitted

Abstract

Purpose: Respiratory motion remains a major challenge in abdominal dynamic contrast-enhanced MRI (DCE-MRI), leading to spatial blurring, temporal signal distortion, and reduced accuracy of quantitative pharmacokinetic analysis. This work proposes and evaluates a respiratory motion–corrected reconstruction framework for free-breathing abdominal DCE-MRI.

Methods: The proposed framework consists of three main steps. First, continuous free-breathing DCE-MRI data were acquired using a prospectively undersampled pseudo-spiral k-space trajectory. Then, the acquired data were retrospectively sorted into three respiratory motion states based on a pressure-belt respiratory signal. For each motion state, three-dimensional (3D) images were reconstructed and deformably registered to estimate 3D motion vector fields. Finally, these motion fields were incorporated directly into a compressed sensing (CS) reconstruction to generate a motion-corrected DCE-MRI time series. Weighted view sharing was applied to mitigate undersampling while preserving temporal signal fidelity. The proposed method was evaluated in four patients and compared with standard CS reconstructions without motion correction. Both dynamic images and parameter maps were evaluated. Image sharpness was assessed using the variance of the Laplacian operator. Temporal signal behavior was assessed visually and compared to standard CS recon. Pharmacokinetic parameters derived using the extended Tofts model were visually assessed.

Results: Motion-corrected reconstruction consistently improved image sharpness across all patients, with median dynamic sharpness increasing by a factor of 3.6–7.8 compared to standard CS reconstruction. Small anatomical structures, including the renal medulla and dilated pancreatic duct, were more clearly delineated. Temporal signal profiles showed improved stability and smoothness without substantial alteration of peak enhancement in the aorta. Quantitative pharmacokinetic maps demonstrated enhanced spatial definition while maintaining comparable parameter values relative to standard CS reconstructions.

Conclusion: The proposed respiratory motion–corrected reconstruction improves spatial and temporal fidelity in abdominal DCE-MRI without biasing quantitative analysis. By integrating motion correction directly into reconstruction, this approach offers a flexible and effective solution for free-breathing DCE-MRI, with potential benefits for both clinical assessment and quantitative research applications.

3.1 Introduction

Dynamic contrast-enhanced (DCE) MRI has potential in assessing treatment response to (chemo-) radiotherapy across various cancer types [1–3]. It involves sequentially acquiring T_1-weighted images during intravenously injected gadolinium-based contrast agent, which locally shortens tissue T_1 and T_2 relaxation times. As the agent passes through the capillaries and diffuses into tissues, it causes time-varying changes in MR signal intensity.

These temporal signal variations reflect underlying physiological properties such as tissue perfusion and capillary permeability. To quantify these parameters, the dynamic signal is first converted into contrast agent concentration, after which a pharmacokinetic model is fitted to the resulting concentration-time curves. Accurate fitting typically requires high temporal resolution (3-5 s per frame [4]) to capture the rapid first-pass of the contrast bolus, and low noise levels, which are often compromised by accelerated imaging and motion of internal organs.

In abdominal DCE-MRI, motion poses a substantial challenge. Respiratory motion is the primary contributor, as the proximity of the lungs to abdominal organs causes periodic displacement of the liver, kidneys, pancreas, and surrounding tissues. This motion occurs intra- and inter-frame, and hence degrades image quality when occurring within a dynamic frame, and causes misalignment between frames when occurring between frames. A common clinical motion mitigation strategy is breath-holding during image acquisition. However, this approach compromises temporal resolution and does not guarantee consistent alignment between frames, as the end-expiratory position may vary between breath-holds [5]. Free-breathing acquisitions combined with retrospective image registration represent an alternative. Yet, motion occurring during acquisition time of one frame can lead to image degradation with blurring and/or streaking artifacts. Moreover, in DCE-MRI, the strong signal enhancement caused by contrast uptake can bias registration algorithms, leading to misalignment or distortion of temporal signal patterns. As a result, even after registration, residual motion artifacts often remain in DCE-MRI datasets [6].

A more robust approach is to account for motion directly within the image reconstruction process. The main advantage of the approach is the compensation of intra-scan motion, which is impossible to achieve with image-based techniques. Rigid motion can be corrected in k-space by adjusting the phase of the acquired signal (known as autofocus) [7]. However, incorporating the deformable motion, typically seen in abdominal organs, into the reconstruction is considerably more complex, as it requires embedding a motion vector field operator into the inverse problem formulation. Several methods have aimed to

address this: [8] proposed a low-rank motion-correction reconstruction for first-pass 2D myocardial perfusion imaging; [9] extended this idea with a simultaneous estimation of motion fields and motion resolved images; and [10] introduced a deep learning approach to estimate motion fields from accelerated cardiac cine k-space data. A fully 3D motion-resolved reconstruction for liver DCE-MRI was demonstrated in [11]. However, this approach employed a longer temporal resolution of 6 s/frame and a Cartesian frequency-encoding scheme with pseudo-radial sampling.

In this study, we advance the field by developing and implementing a motion-resolved reconstruction framework tailored for DCE-MRI of the pancreas. Our method enables faster temporal resolution of 3 s/frame using a pseudo-spiral k-space sampling scheme that, unlike radial sampling, supports more computationally efficient inverse Fourier transforms. We demonstrate substantial improvements in both anatomical image quality and quantitative pharmacokinetic maps in patients with pancreatic cancer.

3.2 Materials and methods

3.2.1 Motion correction algorithm

Our approach consists of three main steps (Figure 3.1). First, data were continuously acquired using a pseudo-spiral k-space trajectory implemented with the Amsterdam UMC software patch, [12, 13] (Figure 3.1A). This incoherent sampling strategy oversamples the center of k-space, resulting in improved signal-to-noise ratio (SNR), increased robustness to motion, and flexible retrospective temporal binning. Second, all k-space data were retrospectively sorted into three respiratory motion states based on the pressure belt respiratory waveform, such that each bin contained an equal amount of data (Figure 3.1B). This binning combined samples acquired over the entire scan, ignoring signal intensity variations due to contrast agent inflow. Next, 3D images were reconstructed for each motion state and registered to the reference image to compute the corresponding motion vector fields (Figure 3.1B). Finally, these motion vector fields were incorporated into the reconstruction to generate motion-corrected DCE image series (Figure 3.1C).

3.2.1.1 Respiratory motion estimation

After k-space data were binned into three respiratory motion states, images for each state were reconstructed in MATLAB (R2021b, MathWorks) using ReconFrame (Gyrotools, Zurich, Switzerland) in combination with the BART toolbox (v0.8.00 BART; https://mrirecon.github.io/bart/). Specifically, compressed sensing (CS) reconstruc-

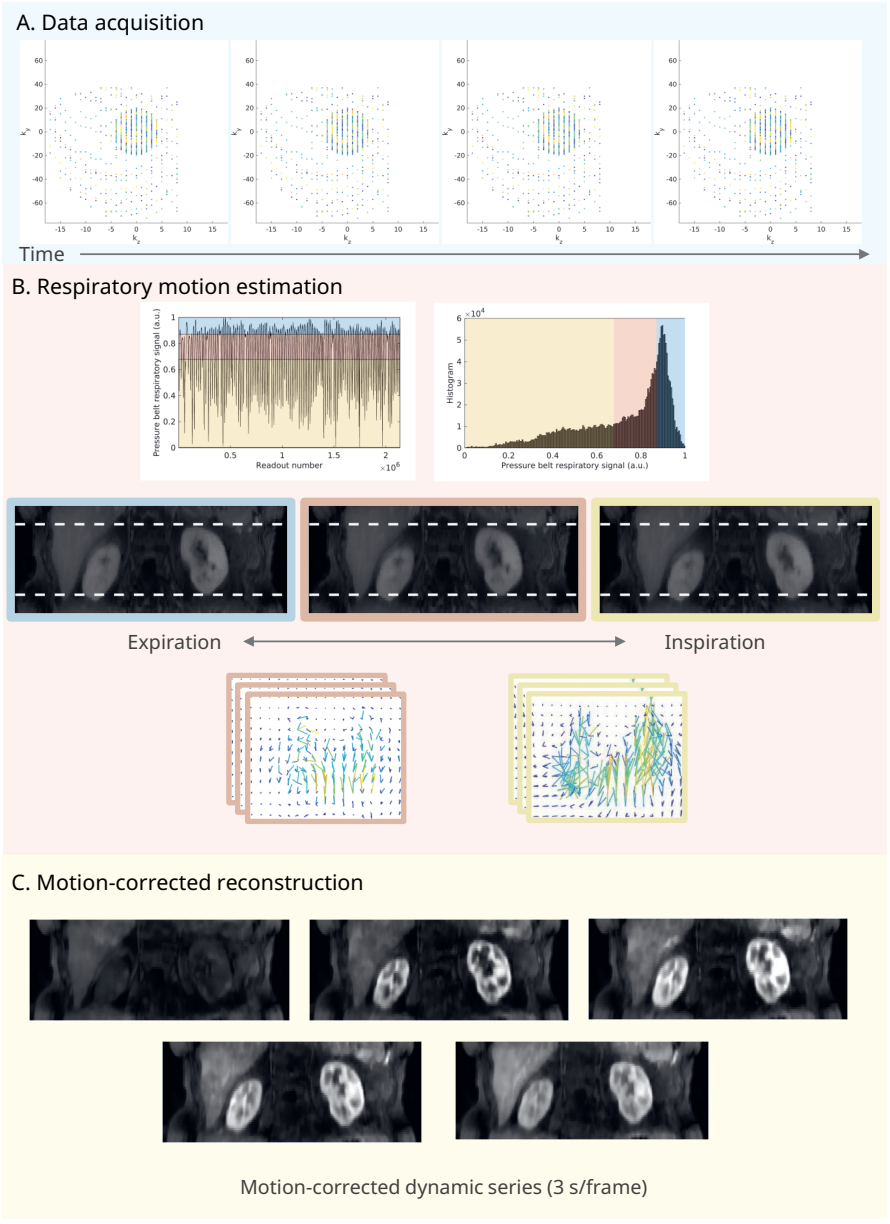

Figure 3.1: Schematic overview of the motion-corrected DCE-MRI reconstruction framework. (A) k-space sampling using a pseudo-spiral acquisition trajectory. (B) Acquired k-space data were retrospectively sorted into three respiratory motion states based on the respiratory belt signal, with bins chosen to contain an equal number of samples. Images corresponding to each motion state were reconstructed and registered to the expiration reference state. (C) Final motion-corrected DCE-MRI images were reconstructed by incorporating the estimated motion vector fields into the reconstruction process.

tion [14] with total-variation regularization applied in both the temporal (λ = 0.01) and spatial (λ = 0.01) domains was used. The expiration state was chosen as the reference. The remaining two motion states were registered to this reference using the MATLAB imregdemons function, which minimizes a sum-of-squared-differences cost function with gradient-descent optimization. Registration was performed with a maximum of 100 iterations, Gaussian smoothing (σ = 2) applied at each iteration, and a four-level multiresolution pyramid. This registration results in motion vector fields $\mathcal{M}$.

3.2.1.2 Motion-corrected image reconstruction

The motion-corrected CS reconstruction of DCE image series incorporated pre-calculated motion-vector fields directly into the forward operator that maps images to k-space [15]:

$$\begin{cases} x^*(t_i) = \text{argmin}_{x(t_i)} \left|\left|y(t_i, m) - Ax(t_i)\right|\right|_2^2 + \mathcal{R}(x(t_i)) \\ Ax(t_i) = \mathcal{P}\mathcal{F}\mathcal{C}\mathcal{M}x(t_i) \end{cases}, \tag{3.1}$$

where, $x(t_i)$ is the MR image to be reconstructed at time t_i, $y(t_i, m)$ is the acquired k-space data at time t_i and respiratory motion state m, $\mathcal{R}$ is the regularization term, and A is the forward model with motion vector fields $\mathcal{M}$, coil sensitivity maps $\mathcal{C}$, Fourier transform $\mathcal{F}$, and a sampling pattern $\mathcal{P}$.

The regularization $\mathcal{R}$ term was represented by total variation in spatial (λ = 0.05) and temporal domains (λ = 0.5).

Here, for each dynamic point, k-space data were sorted into respiratory motion states based on the respiratory belt signal. During reconstruction, all three morion states were corrected to the reference state using $\mathcal{M}$. As a typical respiratory cycle is longer than the acquisition time of a single DCE dynamic, a certain motion state could either be unrepresented or under-represented by a limited number of samples for a given dynamic. In preliminary reconstructions, we found that such data insufficiency led to image undersampling artifacts. To avoid this, we used a view sharing technique [16,17], in which missing k-space samples were supplemented with data from adjacent time points corresponding to the same motion state. Specifically, for a given time point t_i, sampled data from the five preceding dynamics ($y(t_{i-5}, m)$ to $y_(t_{i-1}, m)$) and the five subsequent dynamics ($y(t_{i+1}, m)$ to $y(t_{i+5}, m)$) were incorporated. Data from further time frames were only used if data from closer time frames were unavailable.

However, k-space samples sharing, especially in the central region of k-space, can change image contrast and, consequently, distort the dynamic signal and the derived concentration-time curves. To prevent contrast bias introduced by view-sharing, the data

consistency term was in Eq.3.1 was modified with a weighting scheme. Specifically, each k-space sample contributes to the objective function with a weight $w_{t_i,m}$

$$\left|\left|w_{t_i,m}\, y(t_i,m) - Ax(t_i)\right|\right|_2^2, \tag{3.2}$$

where $w_{t_i,m}$ was defined as a product of a temporal view-sharing penalty $w_{time}(\Delta)$ and a motion state representation penalty $w_{motion}(m)$:

$$\begin{cases} w_{time}(\Delta) = \frac{1}{(\Delta+1)^3} \\ w_{motion}(m) = min(R(t_i))/R(t_i,m) \end{cases}, \tag{3.3}$$

where Δ denotes the temporal distance between the target time point t_i and the time point $t_{i\pm\Delta}$ from which the shared k-space sample $y(t_{i\pm\Delta}, m)$ was added, and $R(t_i, m)$ is an acceleration factor that reflects the sampling density of motion state m at time t_i:

$$R(t_i,m) = \frac{1}{4\pi}\frac{N_{\text{full}}}{N(t_i,m)}. \tag{3.4}$$

Here N_{full} denotes the number of k-space samples required for a fully sampled acquisition (Nyquist sampling), and $N(t_i, m)$ is the number of acquired k-space samples for motion state m at time t_i. $w_{motion}(m)$ prioritizes motion states that have more data samples.

3.2.2 Data

Four patients with pancreatic ductal adenocarcinoma were recruited at the Amsterdam University Medical Center (AMC) between August and November 2023. The study was approved by the institutional review board of the Amsterdam University Medical Center, location AMC (NL73810.018.20). Written informed consent was obtained from all participants prior to MRI scanning. All in vivo imaging was performed on a 3 T scanner (Ingenia, Philips; software release 5.7.1) using a T1-weighted spoiled gradient-echo acquisition (with RF and gradient spoiling) and SPAIR fat suppression. K-space filling followed a pseudo-spiral trajectory, with parameters described in [18]. Acquisition parameters are provided in Table 3.1.

3.2.3 Evaluation

We evaluated our proposed approach against a standard CS reconstruction implemented in the BART toolbox, which did not contain any motion correction techniques. Regularization settings for the standard CS reconstruction were chosen based on previously op-

timized results in [18] with total variation in spatial (λ = 0.05) and temporal domains (λ = 0.05), which was identical to the motion compensated regularisation factors.

We evaluated how motion correction performed in terms of anatomical image quality. For this, we investigated both spatial image quality and temporal signal fidelity. Furthermore, we looked into how motion correction influences quantitative imaging when pharmacokinetic modeling is applied.

3.2.3.1 Anatomical image quality

First, we visually compared reconstructed images with and without motion correction using coronal slices, as respiratory motion predominantly occurs along the foot–head direction. For each subject, we selected representative voxels exhibiting appreciable signal variation and extracted their temporal signal profiles along a line oriented in the foot–head direction to assess the impact of motion correction on temporal consistency.

Image quality was further evaluated by computing a sharpness metric over the entire 3D image volume for all post-contrast time points. Sharpness was quantified using the variance of the Laplacian operator, which measures the second-order spatial derivatives of the image and is sensitive to high-frequency content [19,20]. Regions with pronounced intensity transitions, such as tissue boundaries, yield higher Laplacian variance, whereas blurred images exhibit reduced variance. Consequently, higher sharpness values indicate improved image definition and reduced motion-induced blurring. The resulting sharpness values were compared between reconstructions with and without motion correction.

Additionally, we tested whether temporal signal fidelity was preserved following view sharing and temporal regularization. For this, a region of interest (ROI) was manually delineated within the abdominal aorta on an axial slice located just above the renal arter-

Table 3.1: DCE-MRI acquisition settings.

Contrast	Spoiled gradient echo (RF and gradient spoiling)
Field strength	3 T
Flip angle	25°
TE / TR	1.8 ms / 3.8 ms
Acquisition orientation	Axial (AP readout)
Acquired spatial resolution	1.7 × 1.7 × 3 mm
Reconstructed spatial resolution	1.1 × 1.1 × 3 mm
Acquired matrix size	220 × 152 × 35
Reconstructed matrix size	336 × 232 × 49
TFE shot length	50
Pre-pulse duration	66.1 ms
Reconstructed temporal resolution	3 s
Total reconstructed time	3 min

ies, and the mean dynamic signal was extracted from the enclosed voxels. The aorta was selected because it exhibits the most rapid signal changes among abdominal structures and serves as the primary source for arterial input function (AIF) estimation in pharmacokinetic modeling. Because the standard compressed sensing (CS) reconstruction was optimized for accurate pharmacokinetic parameter estimation rather than for preserving high-frequency temporal dynamics, an additional CS reconstruction was performed using reduced regularization weights (spatial TV λ = 0.01 and temporal TV λ = 0.01). This ensured that potential temporal smoothing effects introduced by stronger regularization were minimized. The resulting time–intensity curves were then compared across three reconstructions: motion-corrected reconstruction, standard CS reconstruction without motion correction, and low-regularization CS reconstruction without motion correction.

3.2.3.2 Quantitative image quality

Quantitative analysis was performed for all subjects by selecting a single axial slice containing healthy renal tissue. Axial plane was chosen because its in-plane voxel dimensions were isotropic, facilitating consistent spatial analysis.

Prior to pharmacokinetic modeling, the dynamic signal intensity in each voxel was converted into a concentration–time curve using a gradient-echo signal model, assuming a short echo time such that T2* effects could be neglected and only T1 relaxation was considered. This conversion required prior knowledge of the native longitudinal relaxation time and the contrast agent relaxivity. A uniform T1 value of 1168 ms was assumed across the slice, corresponding to reported values for renal cortex at 3 T [21], and a relaxivity of 3.5 $mM^{-}1s^{-1}$ was used.

Pharmacokinetic analysis was performed using the extended Tofts model [22, 23] with an analytical representation of the arterial input function (AIF) as proposed in [24]. The population-based AIF was used for all patients [25] with a correction to an average hematocrit fraction of 42% in healthy adults [26]. The extended Tofts model estimates the blood plasma volume fraction (v_p $(fraction)$), the extracellular extravascular space (EES) volume fraction (v_e $(fraction)$), the rate constant for contrast agent transfer from plasma to the EES (K^{trans} (min^{-1})), and back from plasma to EES (k_e (min^{-1})). K^{trans} was computed as the product between k^e over v_e. Model fitting was carried out using a nonlinear least-squares optimization with initial parameter values of $k_e = 0.6$, $v_e = 0.03$, $v_p = 0.0025$, and parameter bounds of $v_e \in [0, 1]$, $v_p \in [0, 1]$, $k_e \in [0, 5]$.

Quantitative parametric maps were visualized for two patients. To exclude background regions and poorly perfused tissue, a mask was applied that removed voxels with a mean signal intensity below twice the median value of the temporally averaged signal.

3.3 Results

3.3.0.1 Anatomical image quality

An example of the reconstructed dynamic series with and without motion correction is shown in Figure 3.2. Without motion correction, respiratory motion degrades the visibility of the gallbladder across time, leading to blurring and inconsistent depiction throughout the dynamic frames. In contrast, the motion-corrected reconstruction substantially improves the delineation of the gallbladder, maintaining its visibility and spatial consistency throughout the dynamic series. Beyond this, the motion-corrected images appear noticeably sharper overall. Improved definition can be observed in vascular structures within the liver, as well as in the spleen. Interestingly, even relatively static regions, such as the paraspinal muscles, exhibit increased sharpness, likely due to reduced motion-induced inconsistencies in k-space that affect the entire image.

Nonetheless, motion correction does not fully eliminate all artifacts, as residual undersampling artifacts remained in the liver and were consistently observed across subjects. These effects may be further mitigated through optimization of reconstruction parameters, such as the choice of regularization weights.

Reconstruction with motion correction resulted in more stable temporal signal profiles for a manually selected pencil beam of voxels along the head–foot direction (Figure 3.3). The temporal profiles of the motion corrected images show calm horizontal profiles, with signal changes over time being dominated by contrast agent arrival. In the absence of motion correction, the profiles reveal that there is substantial motion over time (signal moving up and down), which will lead to poor signal modeling.

Across all four patients, unitless image sharpness was consistently higher for the motion-corrected reconstructions compared to those without motion correction (Figure 3.4). The median improvement in sharpness ranged from 2.39 in Patient 3 (motion-corrected: 2.74 vs. uncorrected: 0.35) to 1.02 in Patient 2 (motion-corrected: 1.41 vs. uncorrected: 0.39), demonstrating a consistent enhancement in image definition following motion correction.

Overall, the intensity of the temporal signal within the aorta was comparable across all three reconstruction approaches (Figure 3.3). For Patients 1 and 4, the motion-corrected reconstruction and the CS reconstruction with lower regularization weights exhibited similar peak signal amplitudes. In contrast, for Patients 2, the peak amplitude in the motion-corrected reconstruction was higher than that obtained with standard CS reconstructions, whereas for Patient 3, the peak was slightly lower in the motion-corrected case. Notably, in Patient 2, the CS reconstruction with higher regularization weights also resulted in a lower

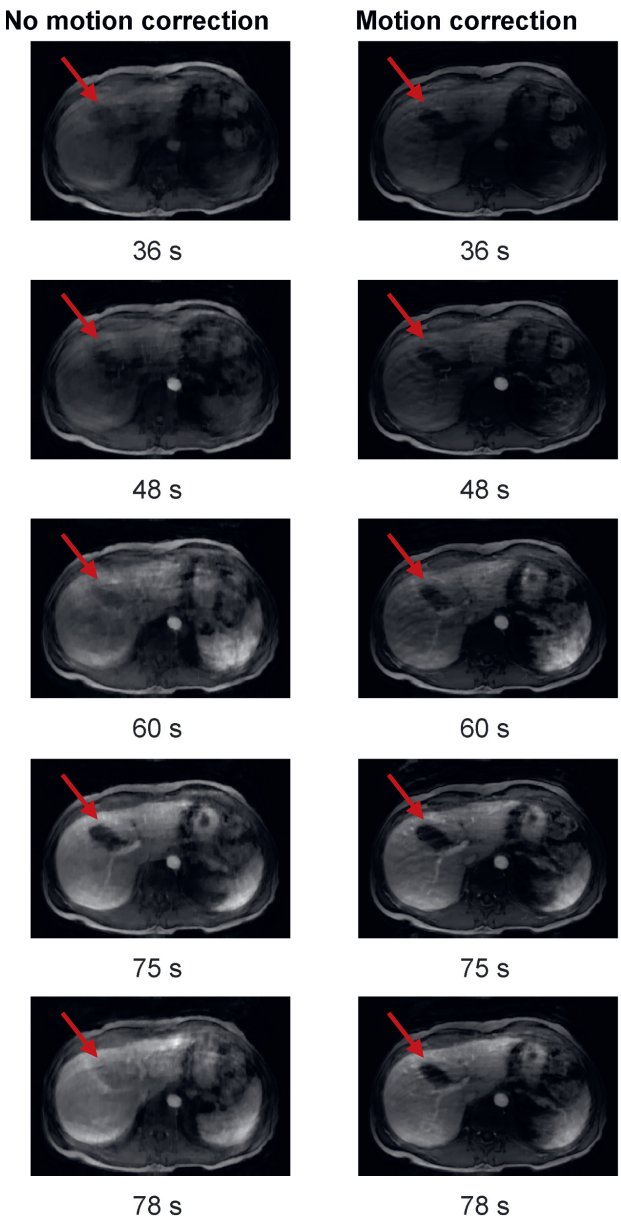

Figure 3.2: Example of reconstructed DCE-MRI from Patient 3 with and without motion correction. The red arrow indicates the gallbladder, which exhibits through-plane motion and reduced visibility across dynamic frames when motion correction is not applied. In contrast, the motion-corrected reconstruction preserves the spatial consistency and improves the visibility of the gallbladder throughout the dynamic series. Note that the window-leveling was kept the same over the dynamics to visualize contrast inflow, which resulted in low signals in the pre-contrast images.

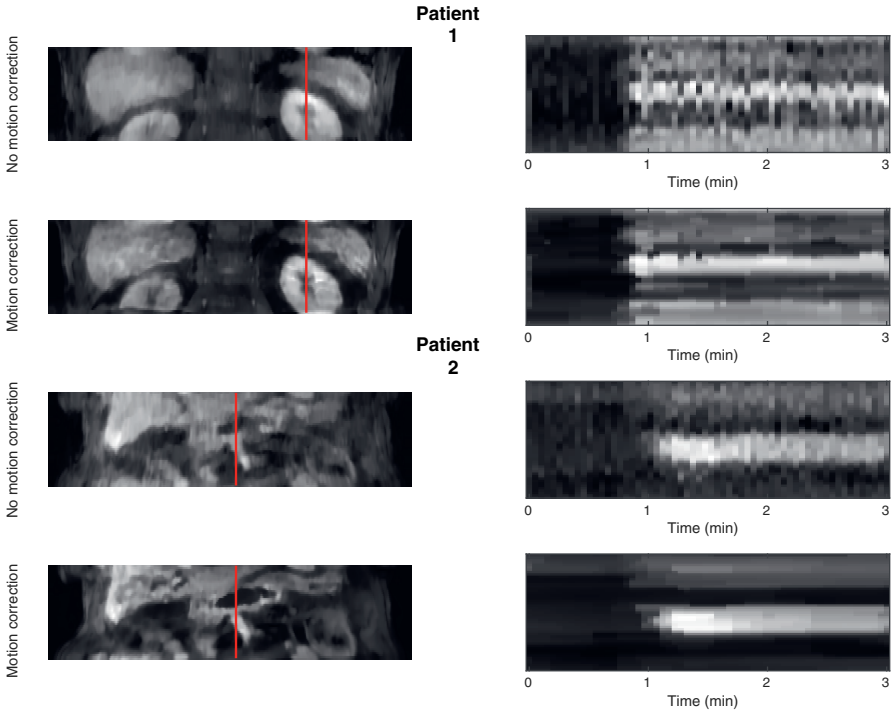

Figure 3.3: Reconstructed images and temporal signal profiles over a pencil beam for Patients 1 and 2, with and without motion correction. First column: Reconstructed images. The red line indicates a manually selected pencil beam for which temporal signal profiles were extracted. Second column: Corresponding temporal signal profiles. The motion corrected signal profiles show calm horizontal profiles, with signal changes over time being dominated by contrast agent arrival. The standard CS results show a messier evolution of the signal over time, with signal profiles moving vertically, indicating motion. Patient 4: Voxels were selected along the spleen–kidney interface. Patient 2: Voxels traversed the liver, the dilated pancreatic duct caused by pancreatic cancer, and the aorta. Note no contrast uptake occurs in the dilated pancreatic duct. The temporal profile of this patient clearly shows that, without motion correction, signal contamination between the duct and surrounding tissues occurs, whereas the motion-corrected reconstruction preserves a sharper transition between structures, resulting in more distinct and physiologically plausible temporal curves.

peak amplitude compared with the low-regularization CS reconstruction. Conversely, for Patient 4, the opposite behavior was observed, with the higher-regularization CS reconstruction yielding a higher peak amplitude than the lower-regularization case. These observations suggest that the accurate estimation of true peak signal intensity could be misrepresented with both standard CS approaches. Finally, the motion-corrected reconstruction consistently produced smoother temporal curves, which is beneficial for quantitative analysis.

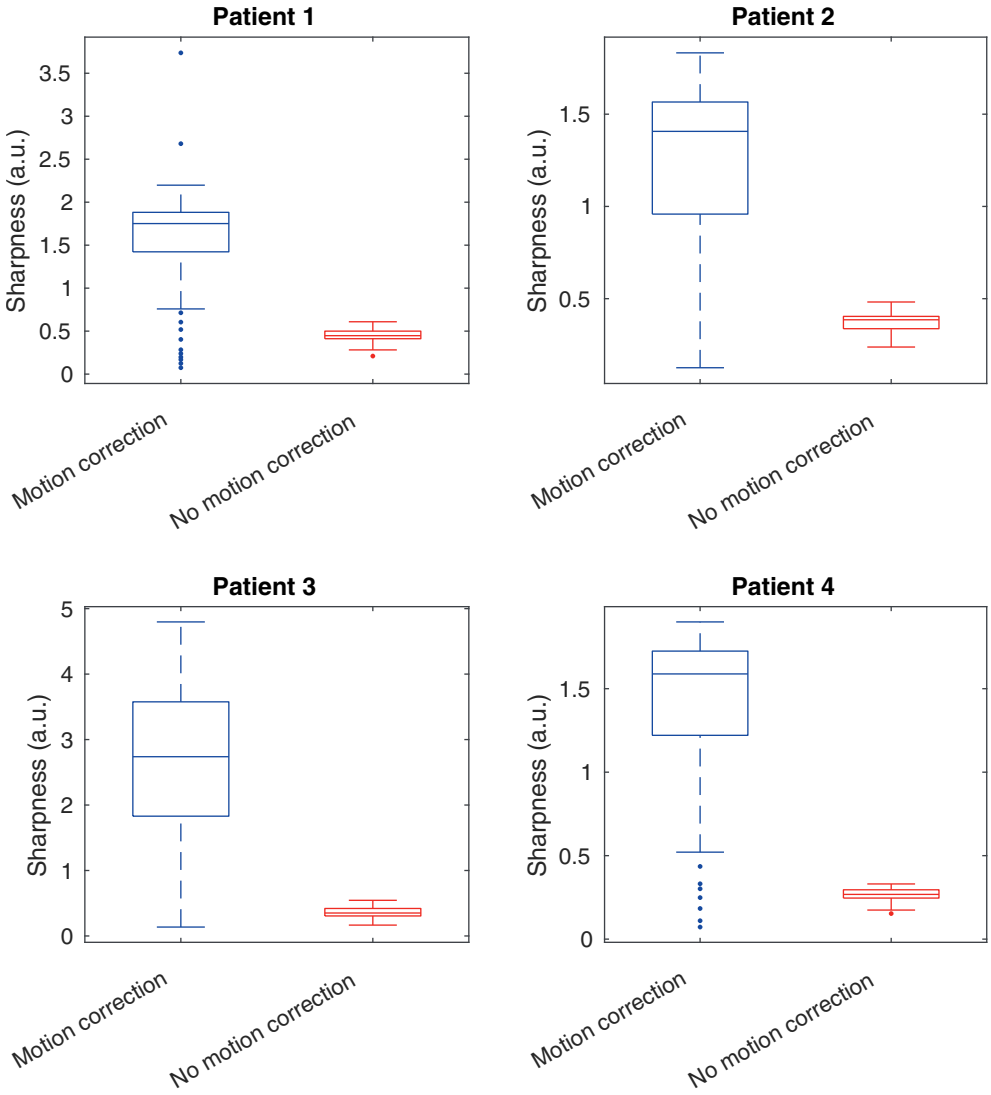

Figure 3.4: Distribution of the sharpness metric, calculated as the variance of the Laplacian operator, across all post-contrast dynamics. A higher value of the Laplacian represents sharper images. Four panels show results for each patient. Red boxplots represent reconstructions without motion correction, while blue boxplots correspond to reconstructions with motion correction, illustrating consistently higher image sharpness when motion correction is applied.

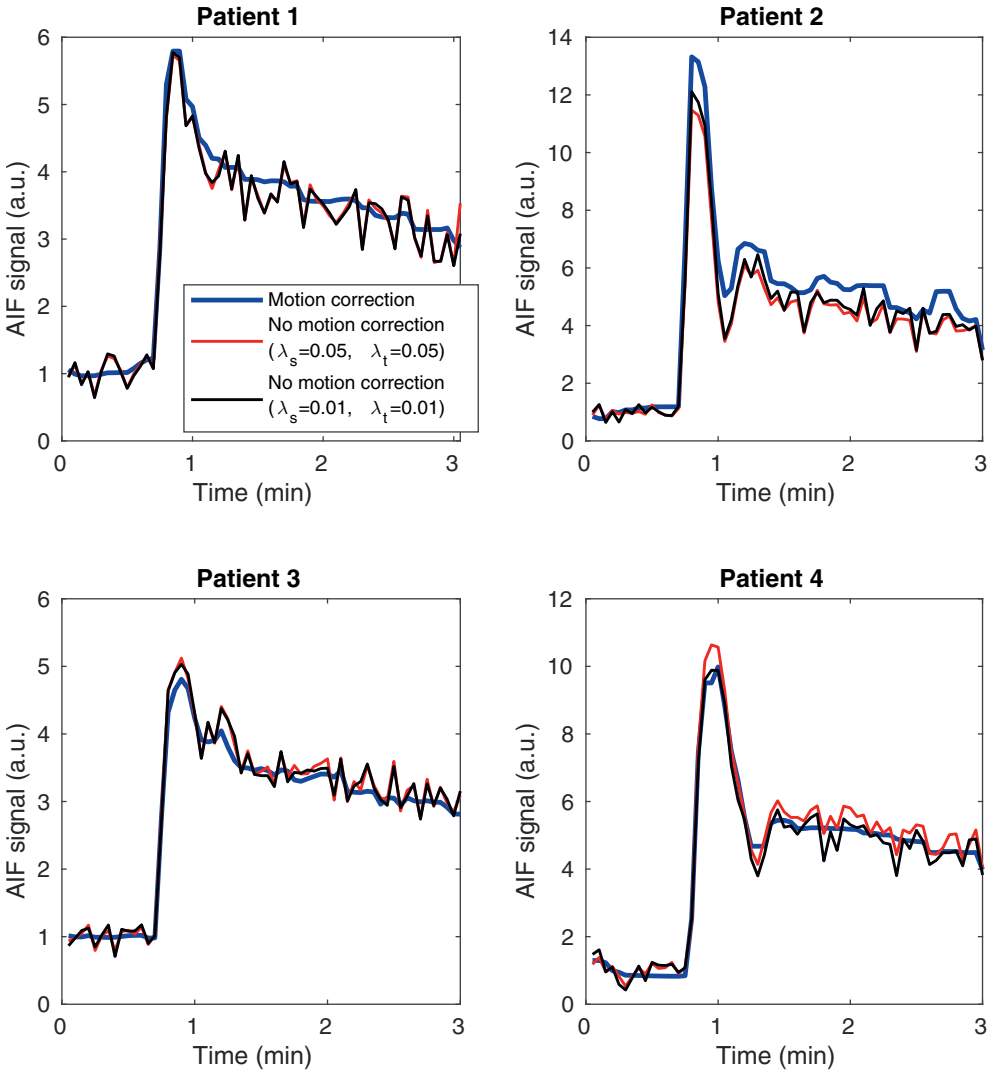

Figure 3.5: Mean dynamic signal extracted from an ROI drawn in the abdominal aorta above the renal arteries. The signal is shown for three reconstruction methods: motion-corrected reconstruction (blue), CS reconstruction without motion correction using spatial and temporal regularization weights of 0.05 (red), and CS reconstruction without motion correction using lower regularization weights of 0.01 (black). For all methods, signal curves were normalized to the mean baseline signal intensity prior to contrast arrival.

3.3.0.2 Quantitative image quality

Overall, the motion-corrected reconstruction produced sharper quantitative maps compared with the conventional CS reconstruction without motion correction (Figures 3.6 and 3.7). This improvement is particularly evident in the contrast between the renal cortex (white arrows) and medulla (green arrows) on the k_e, v_e, and K^{trans} maps. Although the

v_p maps exhibited high sharpness in motion-corrected reconstruction, they also seem to show increased noise. The masks applied for each reconstruction method differed: the mask for motion-corrected reconstruction was more stringent, whereas the uncorrected CS reconstruction showed less restrictive masking, reflecting the reduced contrast between tissues. Elevated noise in the back muscles and intestines was expected, as these tissues are poorly perfused and, in the case of the intestines, subject to uncorrected peristaltic motion.

3.4 Discussion

Respiratory motion presents a major challenge in abdominal DCE-MRI acquisition and processing, as uncorrected motion can introduce complex artifacts, blurring, and misalignment of dynamic images, and hence distorts temporal signal dynamics. In this work, we demonstrated how respiratory motion can effectively be mitigated to improve image quality. The proposed respiratory motion–correction approach operates directly during image reconstruction and, unlike conventional post hoc image registration, explicitly accounts for intra-scan motion, removing blurring and misalignment of dynamic images. We showed that applying this motion correction leads to improved image quality in terms of both spatial fidelity and temporal signal consistency, which are critical for accurate pharmacokinetic modeling.

Motion-corrected reconstructions consistently increased overall image sharpness, enabling clearer visualization of fine anatomical structures such as the dilated pancreatic duct and the renal medulla, as well as improved delineation of tissue boundaries. Despite the view-sharing, which could cause temporal blurring, we show that rapidly varying signal dynamics were well preserved. This was demonstrated by the aortic signal, where peak intensities were comparable across reconstruction methods. The motion-corrected reconstruction also produced smoother temporal curves, which is advantageous for quantitative analysis, as excessive noise can lead to unstable parameter estimation and convergence to local minima during model fitting.

Although the aorta is a key region for arterial input function estimation, the measured signal intensities remained sensitive to ROI placement. To mitigate this effect, ROIs were positioned consistently across patients. Nevertheless, accurate estimation of the true AIF remains a challenge in quantitative DCE-MRI, as evidenced by the variability in peak signal intensities observed when varying regularization strength in standard CS reconstructions without motion correction. If accurate AIF estimation were available, its use would be beneficial for quantitative analysis. However, in this study, incorporating patient-specific

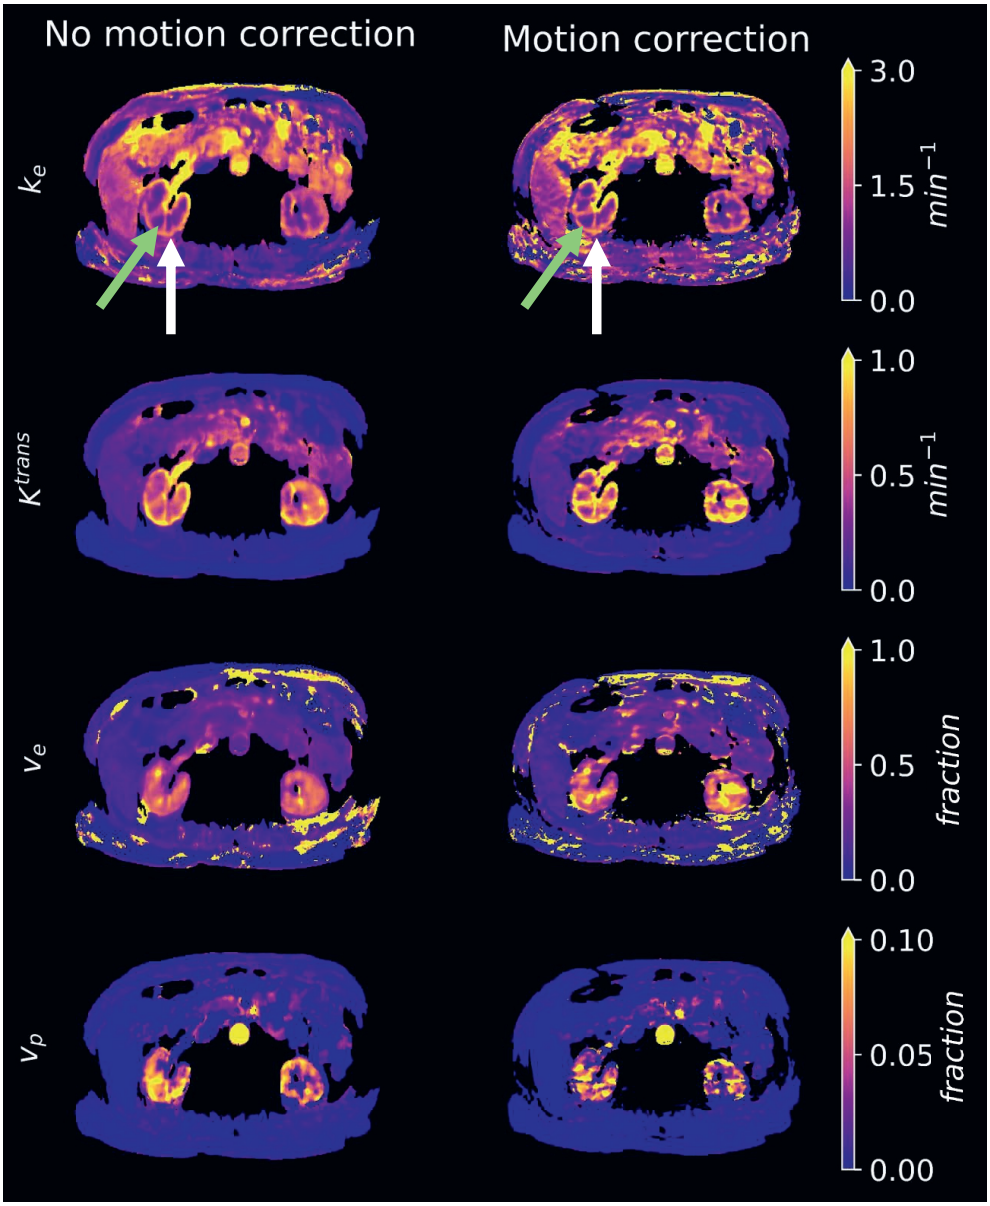

Figure 3.6: Quantitative pharmacokinetic maps for Patient 2, comparing reconstructions with and without motion correction. Parameter values are represented using the color scale shown next to each map. Green arrows on the k_e maps indicate the renal medulla, while white arrows indicate the renal cortex.

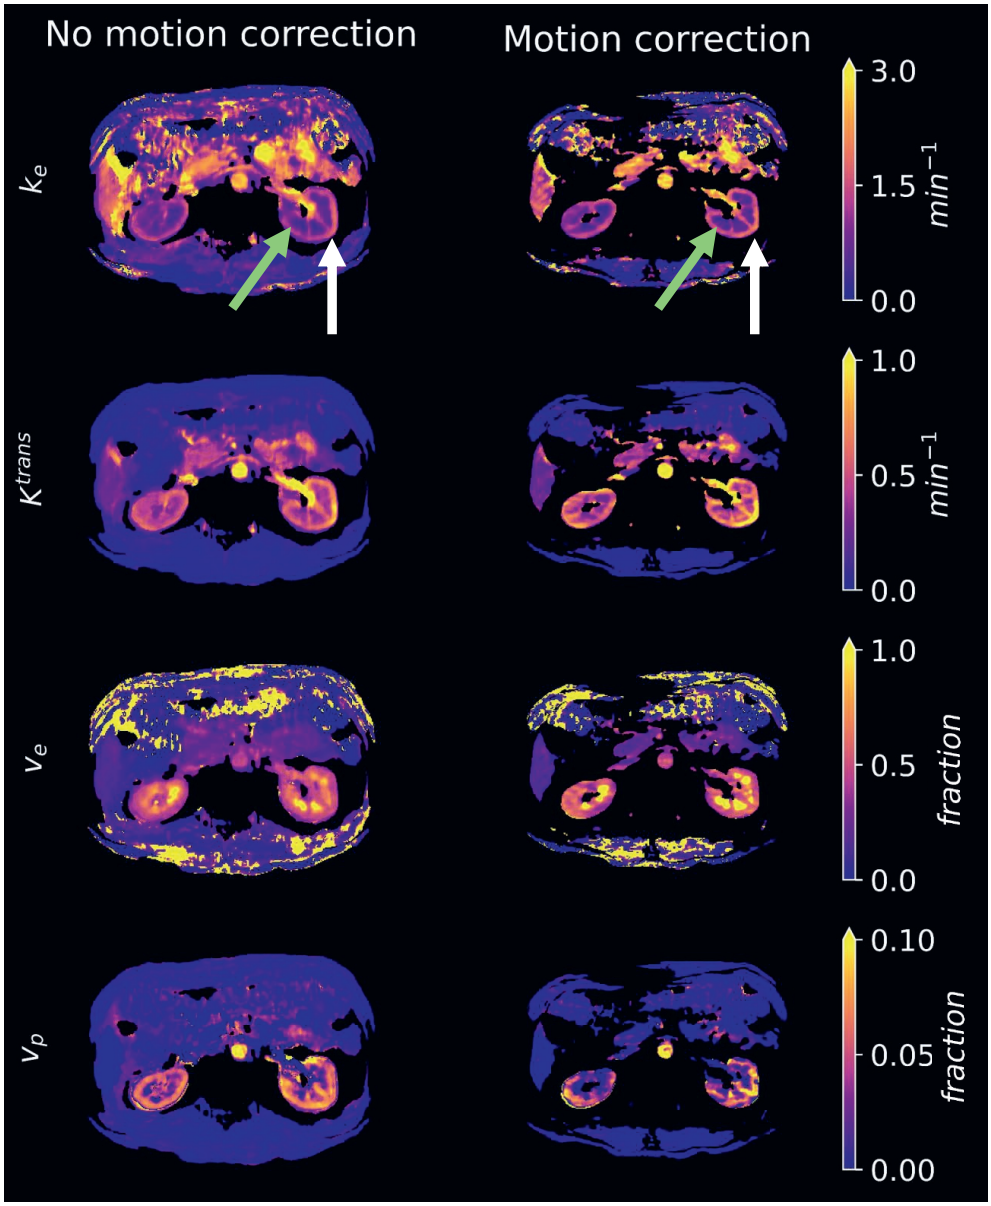

Figure 3.7: Quantitative pharmacokinetic maps for Patient 3, comparing reconstructions with and without motion correction. Parameter values are represented using the color scale shown next to each map. Green arrows on the k_e maps indicate the renal medulla, while white arrows indicate the renal cortex.

AIFs would have introduced additional variability in the reconstructed parametric maps, thereby complicating direct comparison between reconstruction methods.

When applying the extended Tofts model, we observed enhanced sharpness of the resulting parameters in renal medulla and cortex, while quantitative parameter values remained comparable across reconstruction methods. This indicates improved spatial definition without introducing systematic bias. The v_p maps in the kidneys exhibited higher noise levels in the motion-corrected reconstruction. Estimation of v_p is known to be sensitive to the AIF, as both parameters primarily influence the peak signal intensity within a voxel. Consequently, inaccuracies in the AIF can lead to instability in v_p estimation. Although the use of a patient-specific AIF may improve v_p estimation, in this study we deliberately applied the same AIF across reconstruction methods to ensure comparable conditions. This allowed us to isolate and evaluate the effect of reconstruction strategy on pharmacokinetic parameter estimation, independent of potential confounding effects arising from AIF modeling errors.

A strength of the proposed reconstruction framework is its seamless integration into the reconstruction process. Because motion correction is incorporated directly into the forward model, the approach can be readily embedded within existing reconstruction pipelines without requiring separate post-processing steps. This design allows the method to be flexibly combined with a variety of acquisition strategies and reconstruction frameworks. In particular, the approach is compatible with different non-Cartesian sampling schemes that oversample the center of k-space, enabling its application beyond the specific pseudo-spiral trajectory used in this study. Such flexibility facilitates adaptation to diverse acquisition protocols and scanner implementations, enhancing the general applicability of the proposed method.

The proposed reconstruction was demonstrated using a pseudo-spiral k-space trajectory. In practice, DCE-MRI does not follow a standardized acquisition protocol, and a wide range of sampling trajectories have been reported in the literature. An important alternative is radial stack-of-stars acquisition, which samples the center of k-space more frequently and offers greater flexibility in temporal resolution. Extending the proposed motion-corrected reconstruction framework to alternative sampling trajectories is therefore of considerable interest and will be the subject of future work.

3.4.1 Limitations

The method assumed a relatively regular respiratory pattern, as the motion vector fields were pre-estimated prior to the main reconstruction and were assumed to remain consistent throughout the entire scan. In cases of irregular or unstable respiration, characterized

by variations in breathing amplitude and/or frequency, the underlying motion patterns may change over time. In particular, it is known that inspiration is an unstable motion state [27], and treating it as a single state may be oversimplifying. Under such conditions, the current implementation may yield residual motion artifacts due to inaccuracies in the fixed motion model. Absolute binning may alleviate this issue, but would require the insertion of an imaging navigator which would interrupt steady state imaging and limit temporal resolution further. However, we believe that our method would partially deal with this motion, which is better than not deadling with it at all.

In addition, the total acquisition duration of approximately 3 minutes may be insufficient for robust fitting of the extended Tofts model, particularly for accurately capturing late contrast dynamics. In this study, the acquisition length was partly constrained by computational and memory limitations, as the proposed reconstruction is computationally demanding, with a runtime of approximately 8 hours per subject. Importantly, this is a technical limitation rather than a fundamental one, and could be alleviated with improved computational resources or further optimization of the reconstruction pipeline, enabling longer acquisitions in future work. Furthermore, the proposed method corrects only for respiratory motion, while other sources of motion, such as peristaltic motion and cardiac-induced motion, were not addressed and may still contribute to residual artifacts.

Finally, due to computational demands, several hyperparameters, including spatial and temporal regularization weights, image registration settings, the number of respiratory motion states, and temporal resolution were not exhaustively optimized in this study. The primary aim of this work was to introduce the proposed method and demonstrate its feasibility and potential benefits. However, systematic optimization of these parameters may further improve image quality and quantitative accuracy and will be necessary for robust application in clinical and research settings.

3.4.2 Conclusions

In this work, we presented a respiratory motion–corrected reconstruction framework for free-breathing abdominal DCE-MRI that integrates motion information directly into the reconstruction process. The proposed approach improved spatial sharpness and temporal signal consistency, leading to enhanced visualization of anatomical structures and more reliable pharmacokinetic parameter estimation. Importantly, these improvements were achieved without introducing bias in quantitative measures. The flexibility of the framework and its compatibility with different acquisition strategies make it a promising tool for both clinical and research applications.

Bibliography

[1] A. Tudorica, K. Y. Oh, S. Y. Chui, N. Roy, M. L. Troxell, A. Naik, K. A. Kemmer, Y. Chen, M. L. Holtorf, A. Afzal *et al.*, "Early prediction and evaluation of breast cancer response to neoadjuvant chemotherapy using quantitative dce-mri," *Translational oncology*, vol. 9, no. 1, pp. 8–17, 2016.

[2] H. Kim, D. E. Morgan, P. Schexnailder, R. M. Navari, G. R. Williams, J. B. Rose, Y. Li, and R. Paluri, "Accurate therapeutic response assessment of pancreatic ductal adenocarcinoma using quantitative dynamic contrast-enhanced magnetic resonance imaging with a point-of-care perfusion phantom: a pilot study," *Investigative radiology*, vol. 54, no. 1, pp. 16–22, 2019.

[3] J. O'connor, C. Rose, A. Jackson, Y. Watson, S. Cheung, F. Maders, B. Whitcher, C. Roberts, G. Buonaccorsi, G. Thompson *et al.*, "Dce-mri biomarkers of tumour heterogeneity predict crc liver metastasis shrinkage following bevacizumab and folfox-6," *British journal of cancer*, vol. 105, no. 1, pp. 139–145, 2011.

[4] A. Shukla-Dave, N. A. Obuchowski, T. L. Chenevert, S. Jambawalikar, L. H. Schwartz, D. Malyarenko, W. Huang, S. M. Noworolski, R. J. Young, M. S. Shiroishi *et al.*, "Quantitative imaging biomarkers alliance (qiba) recommendations for improved precision of dwi and dce-mri derived biomarkers in multicenter oncology trials," *Journal of Magnetic Resonance Imaging*, vol. 49, no. 7, pp. e101–e121, 2019.

[5] G. Wollny, P. Kellman, A. Santos, and M. J. Ledesma-Carbayo, "Automatic motion compensation of free breathing acquired myocardial perfusion data by using independent component analysis," *Medical image analysis*, vol. 16, no. 5, pp. 1015–1028, 2012.

[6] A. Melbourne, "Alignment of contrast enhanced medical images," Ph.D. dissertation, UCL (University College London), 2009.

[7] D. Atkinson, D. L. Hill, P. N. Stoyle, P. E. Summers, and S. F. Keevil, "Automatic correction of motion artifacts in magnetic resonance images using an entropy focus criterion," *IEEE Transactions on Medical imaging*, vol. 16, no. 6, pp. 903–910, 1997.

[8] G. Cruz, A. Hua, C. Munoz, T. F. Ismail, A. Chiribiri, R. M. Botnar, and C. Prieto, "Low-rank motion correction for accelerated free-breathing first-pass myocardial perfusion imaging," *Magnetic Resonance in Medicine*, vol. 90, no. 1, pp. 64–78, 2023.

[9] T. E. Olausson, M. L. Terpstra, N. R. Huttinga, C. Beijst, N. Blanken, D. Suchá, T. Correia, B. K. Velthuis, C. A. van den Berg, and A. Sbrizzi, "Free-running time-resolved first-pass myocardial perfusion using a multi-scale dynamics decomposition: Cmrmotus," *Magnetic Resonance Materials in Physics, Biology and Medicine*, pp. 1–14, 2025.

[10] A. Ghoul, J. Pan, A. Lingg, J. Kübler, P. Krumm, K. Hammernik, D. Rueckert, S. Gatidis, and T. Küstner, "Attention-aware non-rigid image registration for accelerated mr imaging," *IEEE transactions on medical imaging*, vol. 43, no. 8, pp. 3013–3026, 2024.

[11] M. Ippoliti, M. Lukas, W. Brenner, T. Schaeffter, M. R. Makowski, and C. Kolbitsch, "3d nonrigid motion correction for quantitative assessment of hepatic lesions in dce-mri," *Magnetic resonance in medicine*, vol. 82, no. 5, pp. 1753–1766, 2019.

[12] L. M. Gottwald, E. S. Peper, Q. Zhang, B. F. Coolen, G. J. Strijkers, A. J. Nederveen, and P. van Ooij, "Pseudo-spiral sampling and compressed sensing reconstruction provides flexibility of temporal resolution in accelerated aortic 4d flow mri: A comparison with k-t principal component analysis," *NMR in Biomedicine*, vol. 33, no. 4, p. e4255, 2020.

[13] J. Schoormans, G. J. Strijkers, A. C. Hansen, A. J. Nederveen, and B. F. Coolen, "Compressed sensing mri with variable density averaging (cs-vda) outperforms full sampling at low snr," *Physics in Medicine & Biology*, vol. 65, no. 4, p. 045004, 2020.

[14] D. L. Donoho, "Compressed sensing," *IEEE Transactions on information theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[15] M. Blumenthal and M. Uecker, "Motion corrections in bart." ISMRM workshop on Motion Correction in MR, 2024.

[16] T. Foo, M. A. Bernstein, A. M. Aisen, R. J. Hernandez, B. D. Collick, and T. Bernstein, "Improved ejection fraction and flow velocity estimates with use of view sharing and uniform repetition time excitation with fast cardiac techniques." *Radiology*, vol. 195, no. 2, pp. 471–478, 1995.

[17] M. Markl and J. Hennig, "Phase contrast mri with improved temporal resolution by view sharing: k-space related velocity mapping properties," *Magnetic resonance imaging*, vol. 19, no. 5, pp. 669–676, 2001.

[18] N. P. Wassenaar, O. J. Gurney-Champion, A.-S. van Schelt, T. Bruijnen, H. W. van Laarhoven, J. Stoker, A. J. Nederveen, J. H. Runge, and E. M. Schrauben, "Optimizing pseudo-spiral sampling for abdominal dce mri using a digital anthropomorphic phantom," *Magnetic Resonance in Medicine*, vol. 92, no. 5, pp. 2051–2064, 2024.

[19] J. L. Pech-Pacheco, G. Cristóbal, J. Chamorro-Martinez, and J. Fernández-Valdivia, "Diatom autofocusing in brightfield microscopy: a comparative study," in *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, vol. 3. IEEE, 2000, pp. 314–317.

[20] S. Pertuz, D. Puig, and M. A. Garcia, "Analysis of focus measure operators for shape-from-focus," *Pattern Recognition*, vol. 46, no. 5, pp. 1415–1432, 2013.

[21] G. J. Stanisz, E. E. Odrobina, J. Pun, M. Escaravage, S. J. Graham, M. J. Bronskill, and R. M. Henkelman, "T1, t2 relaxation and magnetization transfer in tissue at 3t," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 54, no. 3, pp. 507–512, 2005.

[22] P. S. Tofts, "Modeling tracer kinetics in dynamic gd-dtpa mr imaging," *Journal of magnetic resonance imaging*, vol. 7, no. 1, pp. 91–101, 1997.

[23] P. S. Tofts, G. Brix, D. L. Buckley, J. L. Evelhoch, E. Henderson, M. V. Knopp, H. B. Larsson, T.-Y. Lee, N. A. Mayr, G. J. Parker *et al.*, "Estimating kinetic parameters from dynamic contrast-enhanced t1-weighted mri of a diffusable tracer: standardized quantities and symbols," *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 10, no. 3, pp. 223–232, 1999.

[24] M. R. Orton, J. A. d'Arcy, S. Walker-Samuel, D. J. Hawkes, D. Atkinson, D. J. Collins, and M. O. Leach, "Computationally efficient vascular input function models for quantitative kinetic modelling using dce-mri," *Physics in Medicine & Biology*, vol. 53, no. 5, p. 1225, 2008.

[25] G. J. Parker, C. Roberts, A. Macdonald, G. A. Buonaccorsi, S. Cheung, D. L. Buckley, A. Jackson, Y. Watson, K. Davies, and G. C. Jayson, "Experimentally-derived functional form for a population-averaged high-temporal-resolution arterial input function for dynamic contrast-enhanced mri," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 56, no. 5, pp. 993–1000, 2006.

[26] H. Lu, C. Clingman, X. Golay, and P. C. Van Zijl, "Determining the longitudinal relaxation time (t1) of blood at 3.0 tesla," *Magnetic Resonance in Medicine: an Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 52, no. 3, pp. 679–682, 2004.

[27] R. George, S. S. Vedam, T. Chung, V. Ramakrishnan, and P. J. Keall, "The application of the sinusoidal model to lung cancer patient respiratory motion," *Medical physics*, vol. 32, no. 9, pp. 2850–2861, 2005.

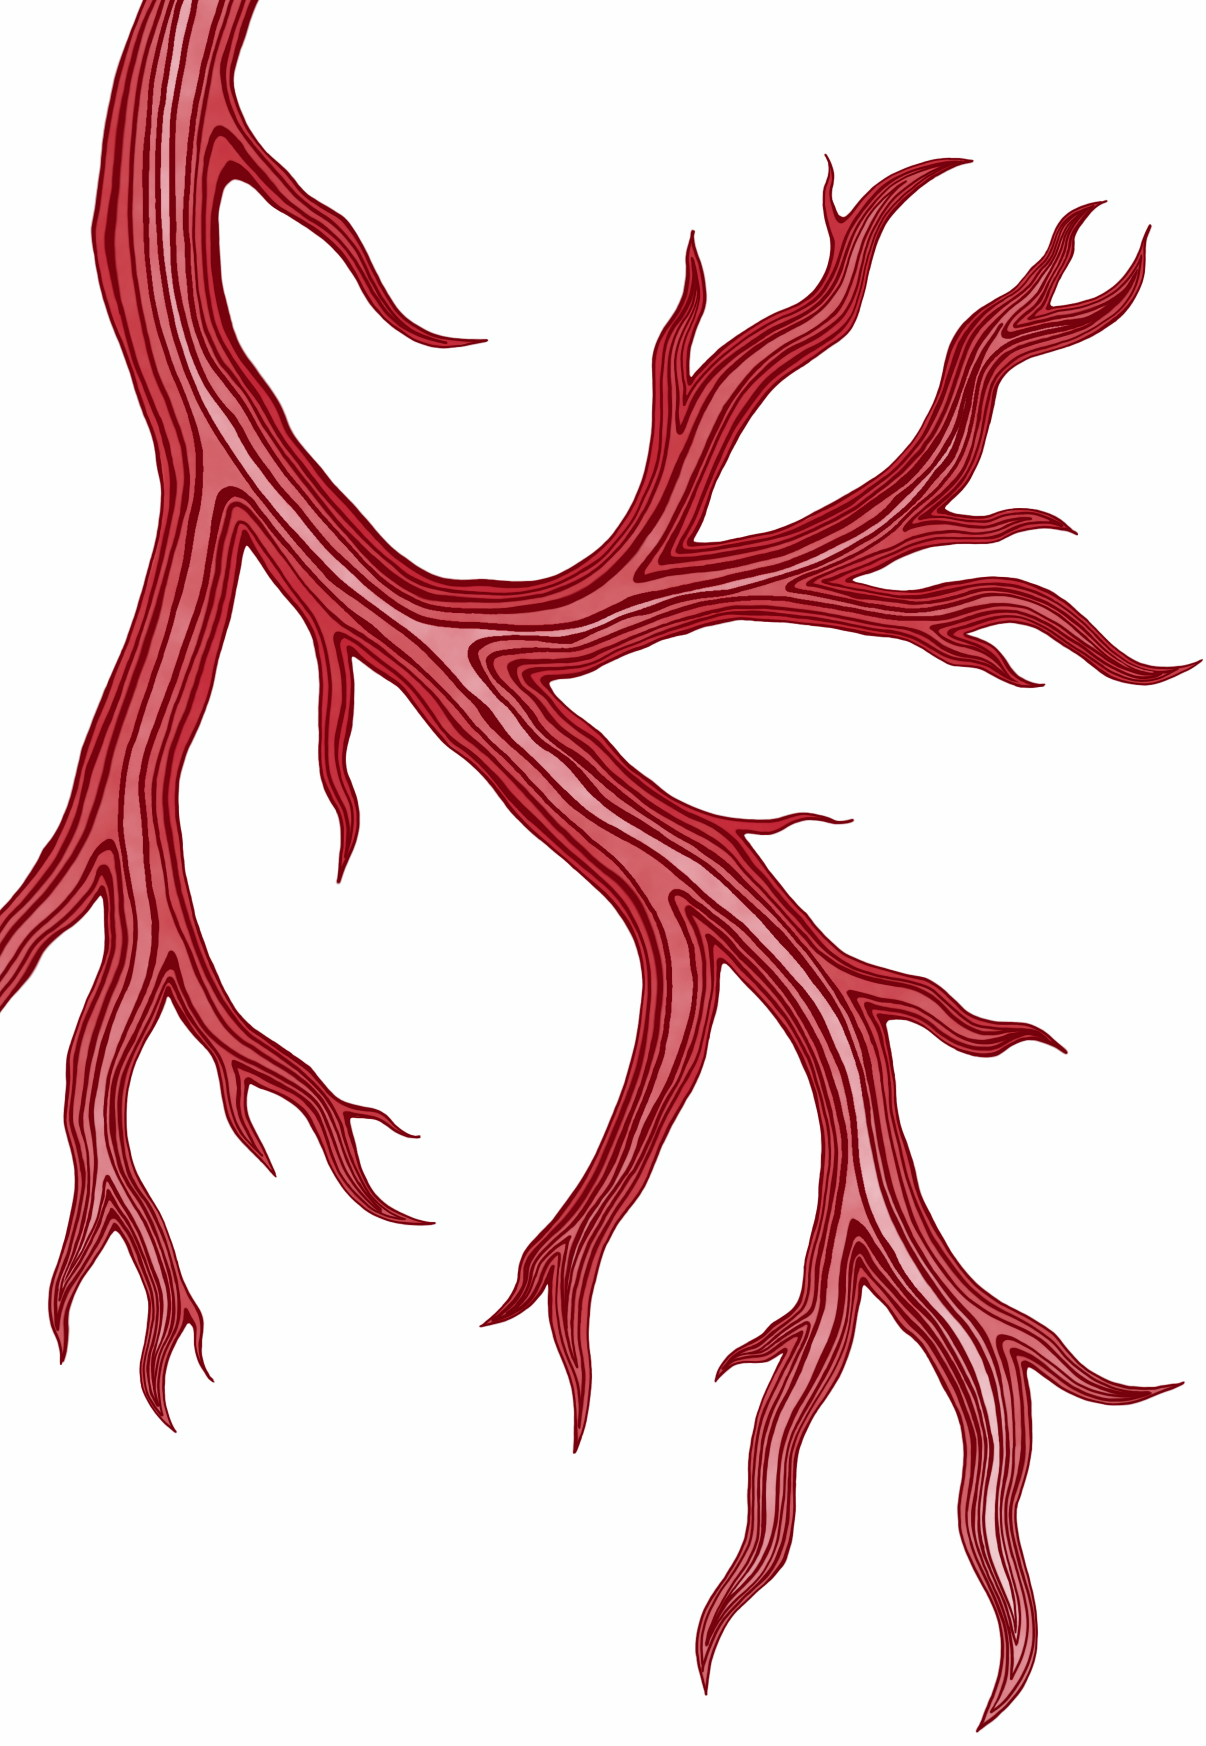

4

Model-based reconstruction for quantitative DCE-MRI parameters

Authors: Natalia V. Korobova, Susanne S. Rauh, Myrte Wennen, Eric M. Schrauben, Oliver Maier, Oliver J. Gurney-Champion

Published in: Magnetic Resonance Materials in Physics, Biology and Medicine (Online ahead of print, April 2026)

https://doi.org/10.1007/s10334-026-01350-x

Abstract

Object: In quantitative dynamic contrast-enhanced MRI (DCE-MRI), a fundamental trade-off exists between imaging speed, spatial resolution, and signal-to-noise ratio (SNR), driven by the amount of data acquired per dynamic frame. This work proposes a model-based reconstruction (MBR) framework for directly estimating pharmacokinetic parameters from raw k-space data, eliminating the need for intermediate image reconstruction and potentially mitigating this trade-off.

Materials and Methods: The extended Tofts model for pharmacokinetic modeling was integrated into an MBR framework - PyQMRI (code is shared). To validate the approach in a controlled setting, a simulated digital phantom of the abdominal region was used. Pharmacokinetic parameters were generated, and corresponding k-space data were calculated based on these values. Additionally, the feasibility of MBR was evaluated in vivo using liver DCE-MRI data from a healthy volunteer. The performance of MBR was compared to a conventional image-based fitting approach using a non-linear least squares (NLLS) algorithm.

Results: In simulations, MBR showed superior performance, producing more precise pharmacokinetic maps with accuracy comparable to or exceeding that of traditional image-based fitting. Notably, fine anatomical structures, such as blood vessels, were more clearly defined with MBR. This improvement was consistent across different temporal resolutions (1.5-16.5 s/frame). MBR did not show any sign of image degradation for shorter frame rates, and, in fact, performed best with the shortest tested frame rate (1.5 s/frame), highlighting the robustness of MBR image quality to higher frame-rates. In vivo, while the improvements offered by MBR were consistent with simulation results, they were less pronounced. Several aspects could have contributed to this discrepancy, including a difference between the simple extended Tofts model and the complex true in vivo pharmacokinetics, and data degradation due to motion combined with a complex landscape of the DCE loss.

Discussion: The proposed MBR framework offers a promising alternative to traditional DCE-MRI workflows by avoiding intermediate image reconstruction and relaxing the spatio-temporal trade-off. This approach may enable more accurate and robust estimation of pharmacokinetic parameters, particularly in scenarios where imaging constraints are severe.

4.1 Introduction

Angiogenesis - the formation of new blood vessels - plays a critical role in the growth and metastasis of malignant tumours, making it a key hallmark of cancer [1, 2]. In tumours, abnormal angiogenesis often fails to meet the high metabolic demands of rapidly proliferating cancer cells, leading to hypoxia [3]. Hypoxia is one of the main features of solid tumours and is associated with poor prognosis of cancer patients [4]. While hypoxia is lethal to many cells, some tumour cells can adapt to hypoxic conditions, developing resistance to chemo- and radiotherapy. Consequently, patients with hypoxic tumours often experience significantly lower overall and disease-free survival rates [5].

Dynamic contrast-enhanced (DCE) MRI has shown potential in assessing treatment response to (chemo-) radiotherapy across various cancers, including pancreatic [6], liver [7], oesophageal [8], rectal [9], cervical [10], breast [11], and brain tumours [12]. DCE-MRI acquires a series of T1-weighted images before and after the intravenous injection of a gadolinium-based contrast agent. The paramagnetic properties of the contrast agent shorten the T1 relaxation time in tissues and blood, increasing the signal intensity of the T1-weighted images. Pharmacokinetic and signal modelling are then applied to extract physiological parameters describing microvascular perfusion and vascular permeability from the DCE signal-time curves. Because these parameters determine how the DCE signal is formed, their estimation involves solving an inverse problem - typically approached through voxel-wise curve fitting.

However, the clinical adoption of quantitative DCE-MRI is hampered by several challenges that hinder its use in clinical routine. One issue is the inherent spatio-temporal trade-off in DCE-MRI: accurate pharmacokinetic modelling requires high temporal resolution to capture rapid contrast dynamics [13], while high quality (more sampling), high spatial resolution is desired for visualizing small structures such as tumours. Balancing these competing demands remains a significant challenge. As most clinical routines use a small amount of high-quality high-resolution post-contrast images (e.g., pre-contrast, arterial, equilibrium and wash-out phases), whereas close to none use DCE; this suggests that good image quality is desirable for clinical use. A second issue is that conventional fitting algorithms operate on a voxel-wise basis, making them more susceptible to noise and local minima. Finally, and partially as a result of the aforementioned issues, there are no large trials showing the definitive added benefit of DCE.

These challenges stem, in part, from fitting pharmacokinetic parameters in the image domain, which relies on acquiring sufficient k-space data to produce dynamic images. A large chunk of DCE-MRI literature is dedicated to reducing the amount of k-space needed

per dynamic image by, e.g. view-sharing [14, 15], compressed sensing [16] and low-rank reconstruction [17], or even deep learning [18]. An alternative approach would be estimating these parameters directly from the raw k-space data, omitting the need for images and potentially improving the quality of estimated maps across different temporal resolutions. Moreover, such a method could be used to re-analyze data from modern clinical contrast-enhanced MRI protocols that work with continuous sampling, such as radial and spiral acquisitions (GRASP, Radial VIBE, DISCO, StarVIBE, Wave-CAIPI), to acquire DCE parameter maps alongside the clinical images. This will greatly increase the ease of obtaining large DCE datasets in clinically relevant tissues and build up clinical evidence, without hindering routine clinical protocols.

Estimating quantitative MRI parameters directly from raw k-space data can be achieved with the model-based reconstruction (MBR) technique. MBR solves an inverse problem associated with the quantitative model with an iterative approach that compares measured k-space data to the results of the forward modelling of the estimated parameters. Initially introduced for relaxation mapping [19, 20], MBR has since been widely applied with promissing results to T1 and T2 mapping [21–23] and other MRI techniques, including arterial spin labelling [24], chemical exchange saturation transfer [25], diffusion MRI [26], diffusion tensor imaging (DTI) [27, 28], and combined intravoxel incoherent motion imaging and DTI [29]. Despite some studies exploring its use in quantitative DCE-MRI [30–33], MBR remains unestablished in this field.

While previous studies have shown promising results, some used simpler pharmacokinetic models, such as the Patlak model [32, 33], which neglects contrast reflux from the extravascular extracellular space (EES). The extended Tofts model, although more comprehensive, has been less frequently applied in MBR, likely due to its complexity. Unlike simpler models, it requires estimating three pharmacokinetic parameters along with arrival time and baseline signal, making it more challenging to fit. Additionally, some of these parameters exhibit strong correlations and exist on different scales, further complicating parameter estimation.

Studies that have incorporated the extended Tofts model in MBR [30, 31] faced key limitations. Felsted et al. did not use spatial regularisation [30], while Dikaios et al. [31] applied total variation (TV), which is known to introduce staircasing artifacts in parameter maps. In contrast, we employ total generalized variation (TGV), a higher-order regularisation technique that better preserves smooth transitions while maintaining edge sharpness [34]. Furthermore, while Dikaios et al. relied on stochastic gradient Langevin dynamics (SGLD) for optimisation - a sampling-based method that is computationally intensive and harder to tune - we used a deterministic primal-dual algorithm within a Gauss-

Newton framework. This offers faster, more stable convergence and is well suited for large-scale, GPU-accelerated reconstruction. Finally, to our knowledge, no previous work has evaluated the impact of temporal acceleration on MBR using the extended Tofts model. Our study fills this gap in understanding the feasibility and accuracy of MBR in accelerated imaging.

To address these challenges, we leverage PyQMRI [35, 36], an open-source MBR framework designed for quantitative MRI techniques (available at `https://gitlab.tugraz.at/ibi/mrirecon/software/PyQMRI`). PyQMRI incorporates TGV-based regularization, which introduces more specific prior information, helping to stabilize parameter estimation. Unlike traditional approaches, its TGV functional is adapted to multiparametric settings, allowing it to exploit shared structures across parameter maps. Moreover, PyQMRI optimizes volumetric data, making it well-suited for high acceleration rates by effectively utilizing spatial information across the entire image volume. This makes it a promising tool for improving MBR performance in accelerated DCE-MRI parameter estimation.

We hypothesize that MBR will allow us to overcome the spatio-temporal trade-off typically underlying DCE-MRI. To test this, we implemented the extended Tofts model into the PyQMRI MBR framework. We validated our approach using an adaptation of the MR-XCAT phantom designed to replicate the acquisition scheme of real in vivo DCE-MRI data. Furthermore, we compare the performance of MBR-based parameter estimation with conventional image-based fitting across different dynamic timings, assessing its accuracy, robustness, and ability to mitigate the trade-off. Finally, we demonstrated the feasibility of MBR for DCE parameter estimation in vivo in a healthy volunteer.

4.2 Materials and methods

4.2.1 Pharmacokinetic modelling

For pharmacokinetic modelling, we employed the extended Tofts model [37, 38], which is widely used for tumour characterisation [39]. It assumes that the contrast agent diffuses between two compartments: blood plasma and EES. The concentration of the contrast agent $C(t)$ is expressed through the pharmacokinetic parameters as follows:

$$C(t) = v_p C_p(t - t_0) + k_{ep} v_e \int_0^t C_p(\tau - t_0) e^{-k_{ep}(t-\tau)} d\tau, \tag{4.1}$$

where $k_{ep}\,[min^{-1}]$ is the mass transfer rate of reflux , $v_e\,[fraction]$ is the volume of EES, $v_p\,[fraction]$ is the volume of plasma, $t_0\,[min]$ is the contrast agent's transit time from the artery to the tissue, and $C_p\,[mM]$ is the concentration of the contrast agent in the blood

plasma, expressed as the arterial input function (AIF). The influx rate of contrast K^{trans} $[min^{-1}]$ is defined as the product of k_{ep} and v_e.

To reduce manual intervention in modelling, we assumed that the bolus arrived at the feeding artery at the start of the scan ($t = 0$). Under the assumption that the arterial arrival time is zero, t_0 directly reflects the actual delay of contrast arrival in the tissue relative to the scan start time.

We used the extended Tofts model implementation and fitting routines from Orton et al. [40] from the OSIPI GitHub repository [41] available at https://github.com/OSIPI/DCE-DSC-MRI_CodeCollection/tree/develop/src/original/OG_MO_AUMC_ICR_RMH_NL_UK.

This implementation includes an analytical representation of AIF C_p, which is parameterised using four variables θ_{AIF} that define its shape - capturing the initial bolus passage and subsequent wash-out phase. To apply this AIF model, we fitted these parameters to population-based AIF data [42]. Finally, the AIF was adjusted for an average hematocrit fraction of 42% in healthy adults [43].

Since the MBR operates in the signal domain, we converted contrast concentration measurements into MRI signal values. In DCE-MRI, the relaxation time $T_1(t)$ changes throughout the dynamic acquisition as a result of the interaction of blood and other tissues with the paramagnetic contrast agent:

$$T_1(t) = \frac{1}{rC(t) + \frac{1}{T_{10}}} \tag{4.2}$$

where r $[mM^{-1}s^{-1}]$ is the relaxivity of the contrast agent (set to 4.5 $mM^{-1}s^{-1}$), T_{10} $[s]$ is the precontrast baseline T_1 $[s]$ value, and $C(t)$ $[mmol/L]$ is the concentration of contrast agent expressed with the extended Tofts Eq.4.1.

The corresponding MRI signal was modeled using a fast gradient echo sequence:

$$\mathcal{S}(t) = M_0 \frac{\sin\alpha(1 - e^{-\frac{TR}{T_1(t)}})}{1 - (\cos\alpha)e^{-\frac{TR}{T_1(t)}}} \tag{4.3}$$

where α is the flip angle, TR is the repetition time, M_0 is the spin density, and $T_1(t)$ is the relaxation time. The parameters α and TR are acquisition-specific and typically taken from the scan protocol, as described in Section 4.2.3. To account for the fat-saturation pulse, we corrected TR according to:

$$TR = TR_{protocol} + \frac{time_{fat\ saturation\ pulse}}{number\ readouts\ per\ pulse} \tag{4.4}$$

The parameters M_0 and T_{10} can either be fitted alongside with pharmacokinetic parameters or pre-estimated from an additional T_1-weighted scan.

By placing Eq. 4.1 into Eq. 4.2, and the resulting equation into Eq. 4.3, we obtained the direct dependence of the signal $\mathcal{S}(\sqcup|\sqcap)$ on the estimated parameter set $[M_0, T_{10}, v_e, v_p, k_{ep}, t_0]$. This included the spin density, baseline T_{10} value, and four pharmacokinetic parameters. As our acquisition was not designed to estimate T_{10}, T_{10} was assumed known, and the fitted parameter set was $u = [M_0, v_e, v_p, k_{ep}, t_0]$.

4.2.2 Model-based reconstruction

In this work, we implemented the extended Tofts model within the open-source MBR toolbox PyQMRI [35, 36] available at `https://gitlab.tugraz.at/ibi/mrirecon/software/PyQMRI`. This algorithm was designed to solve a non-linear, ill-posed optimization problem, formulated as:

$$\begin{cases} \min_u ||A(u,t) - d(t)||_2^2 + \gamma R(u) \\ A(u,t) = \mathcal{P}\mathcal{F}\mathcal{C}\mathcal{S}(u,t) \end{cases}, \tag{4.5}$$

where u are the model's parameter maps, d corresponds to the measured raw k-space data, γ is the regularisation weight and $R(u)$ is the regularisation term, formulated as TGV, that incorporates prior knowledge about the estimated parameters u. Finally, $A(u,t)$ represents the forward operator, which combines the following components:

- Biophysical signal equation $\mathcal{S}$,
- Coil sensitivity maps $\mathcal{C}$,
- Fourier transformation $\mathcal{F}$,
- K-space sampling pattern $\mathcal{P}$.

The weighting factor γ controls the balance between data fidelity and prior information. In PyQMRI, it is dynamically adjusted during the iterative optimisation process to ensure stable convergence. Users can control three key parameters: the initial weighting factor γ_{init}, the minimum allowable value γ_{min}, and the decay rate γ_{dec}, which determins how quickly γ decreases with each iteration. The regularization term $R(u)$ is implemented as a second-order TGV [34,44] combined with a joint Frobenius norm [22,45,46], ensuring spatial smoothness and noise suppression. Reconstruction of off-grid sampling trajectories in PyQMRI is enabled through a non-uniform Fourier transform (NUFFT).

Although, in principle, d can represent each individual k-space readout, doing so is computationally demanding as a Fourier Transformation needs to be performed for each time-step within the iterative optimisation loop. In our implementation, the data was grouped

into segments consisting of multiple readout lines. As a result, even though the model was formulated in continuous time, the reconstruction was determined by these k-space segments. In this way, temporal resolution became an intrinsic property of the optimisation process that balances reconstruction time and required RAM memory with potential gain in accuracy.

Before the estimation of quantitative parameters, in PyQMRI, fast image reconstruction was calculated with the conjugate-gradient sensitivity encoding (CG-SENSE) [47] algorithm, and optionally coil sensitivities are established with a Newton-type method [48]. The objective function (Eq.4.5) was minimised using an iterative Gauss-Newton method, where each subproblem was solved with a Primal-Dual algorithm with a line search [49,50] to ensure convergence. A more detailed explanation of the PyQMRI algorithm can be found in the works by Maier et al [35, 36].

Here, we implemented the extended Tofts model within the PyQMRI framework (`https://gitlab.tugraz.at/ibi/mrirecon/software/PyQMRI/-/blob/master/pyqmri/models/DCE_fit_M0_numpy.py`). For this, we expressed the signal equation $\mathcal{S}$ of the forward operator $A(u,t)$ with combined equations 4.1,4.2, and 4.3. Physiological parameters u were represented as real numbers and constrained within physiologically relevant boundaries (Table 4.1) and M_0 was represented as complex numbers and unconstrained. Note that t_0 was bounded at a minimum of 0.17 minutes as the contrast agent was injected after this time. The initial guess for M_0 was the average of the precontrast image series reconstructed with the CG-SENSE algorithm using Eq.4.3. The initial guesses of the parameters are described below in the section 4.2.4 *Experiments*. The PyQMRI configuration files used in this work can be found in Appendix 4.A.

Table 4.1: Boundaries of fitted parameters

Parameter	Unit	Minimum value	Maximum value
v_p	[fraction]	$1e-8$	1.0
v_e	[fraction]	$1e-8$	1.0
k_{ep}	[min^{-1}]	$1e-8$	3.0
t_0	[min]	0.17	2.0

4.2.2.1 Image-based fitting

A conventional voxel-wise fit in the image space was performed to serve as a comparison for the MBR approach. The undersampled k-space data were first reconstructed into an image series, and then pharmacokinetic parameters were estimated using a non-linear least squares (NLLS) algorithm. The parameter M_0 was calculated from the average signal

measured before contrast injection using Eq.4.3. The fit boundaries were set identical to those in MBR (Table 4.1). Initial guesses for the fitted parameters were 0.6 $[min^{-1}]$ for k_{ep}, 1 $[fraction]$ for v_e, 0.2 $[min]$ for t_0, and 0.0025 $[fraction]$ for v_p.

4.2.3 Data

4.2.3.1 Digital phantom

To optimize and test DCE-MRI MBR in a controlled environment, we used the MRI adaptation of a digital phantom of extended cardiac-torso (MR-XCAT) anatomy [51, 52]. We further extended this phantom to incorporate pharmacokinetic dynamics and advanced sampling patterns (MR-XCAT with ERIC) [53]. The phantom is available at https://github.com/schrau24/XCAT-ERIC. To optimize computational efficiency, we selected a 2D axial slice of the phantom and simulated a DCE-MRI acquisition. Although DCE-MRI is typically acquired in 3D, we adapted the simulation by reducing the number of readouts to represent what would be expected from a 2D slice within a full 3D acquisition. The simulated dataset included various tissues and organs: muscle (1661 voxels), liver (1646 voxels), gallbladder (198 voxels), stomach wall (224 voxels) and contents (97 voxels), pancreas (458 voxels), kidney cortex (235 voxels) and medulla (68 voxels), adrenal gland (114), spleen (255 voxels), intestine (2214 voxels), cartilage (195 voxels), and blood (195 voxels).

Since the objective function of PyMQRI (Eq. 4.5) was defined in k-space, we simulated the acquisition of raw k-space data. First, we generated pharmacokinetic parameter maps for the phantom. For tissues in which pharmacokinetic parameters were known from literature, we used those values [54, 55]. For all other tissues, randomly sampled values from within a physiologically realistic range (Table 4.1) were used. Based on these parameters, we calculated the contrast agent concentration $C(t)$ as a function of time per voxel using Eq. 4.1. The distinct times at which simulated $C(t)$ needs to be calculated depend on the simulated acquisition. In our simulations, we implemented a stack-of-stars golden-angle acquisition with an in-slice sampling rate of 0.3 seconds (given 85 samples along z-axis and $TR = 3.5\,ms$, 85 samples per stack $\times\, 0.0035\,s$ per sample $= 0.3\,s$ per stack). Next, we applied Eqs. 4.2 and 4.3 to compute the corresponding signal curves $\mathcal{S}(t)$ in the image domain, ensuring that acquisition parameters matched those of a typical DCE-MRI liver scan (Table 4.2). This results in a 4D dynamic DCE image. Finally, per time-point, we applied a Fourier transformation $\mathcal{F}$ at the sampling locations, following a golden-angle radial stack-of-stars acquisition scheme. This was, each k-space line corresponds to a specific point along the DCE curve, as would be the case during an acquisition.

Then, we grouped the k-space lines into segments corresponding to different temporal resolutions that are intrinsically used in the optimisation process of PyQMRI. To test our hypothesis that MBR can overcome the conventional spatio-temporal trade-off, we reconstructed the data using four temporal resolutions: 1.5 s/frame (5 readouts, acceleration rate relative to Nyquist sampling $R = 35.2$), 2.4 s/frame (8 readouts, $R = 22.0$), 3.9 s/frame (13 readouts, $R = 13.5$), and 16.5 s/frame (55 readouts, $R = 3.2$). Here, the acceleration factor R relative to Nyquist sampling is defined as

$$R = \frac{\pi FOV}{2\Delta x N_s}, \tag{4.6}$$

where FOV is the extent of the image, Δx is the spatial resolution, and N_s is the number of acquired radial spokes.

To mimic real-world conditions, we added Gaussian-distributed noise to the sampled k-space data. Noise was adjusted to yield an image-domain SNR of 50 in the liver when averaged across the DCE time series. This SNR matched that of a fully sampled Cartesian reference dataset in our clinic.

Table 4.2: Acquisition parameters used in a simulated digital phantom and in vivo

Data	Digital phantom	In vivo
Contrast	Spoiled gradient echo	Spoiled gradient echo
Field strength	3.0 T	1.5 T
Flip angle	22°	12°
TE	1.4 ms	2.35 ms
TR	3.5 ms	4.3 ms
TFE shot length	-	28
Pre-pulse duration	-	37 ms
Relaxivity contrast agent	4.5 $mM^{-1}s^{-1}$	4.5 $mM^{-1}s^{-1}$
Acquisition matrix	1005, 1008, 1014, 1045 readouts, 224 sampling points per readout	28 partitions, 1780 readouts per partition, 512 sampling points per readout
Acquisition orientation	Axial	Axial
Acquired/reconstructed resolution	2 x 2 x 3 mm	1.5 x 1.5 x 3 mm
FOV	224 x 224 x 3 mm	380 x 380 x 216 mm
Reconstructed temporal resolution	1.5, 2.4, 3.9, 16.5 s/frame	2.0 s/frame
Scan length	5 min	5 min

4.2.3.2 In vivo data

To show the feasibility of MBR framework, we applied it to one liver DCE-MRI scan of a healthy volunteer.

One healthy volunteer underwent an abdominal DCE-MRI examination. The study was approved by the ethics committee of the Amsterdam UMC (NL72271.029.19), and the participant gave written informed consent before the start of the study. The MRI scan was performed on a 1.5 T system (MAGNETOM Sola, Siemens Healthcare, Erlangen, Germany), using a 18-channel phased-array anterior coil and a 32-channel phased-array posterior receive coil. After about 30 s, 2 mL/kg of 0.5 mmol/mL gadoterate meglumine (Dotarem, Guerbet, Villepinte, France) followed by a saline flush of the same volume were administered at a rate of 2 mL/s. The DCE protocol comprised a dynamic series consisting of a golden angle radial stack-of-stars spoiled gradient echo sequence with fat saturation. The imaging parameters are specified in Table 4.2. The participant was instructed to breathe freely during the entire DCE image acquisition. No motion correction was performed.

4.2.4 Experiments

4.2.4.1 Digital phantom

The characteristic effects of v_p and t_0 are most prominent in a small region of the curve, during the arrival of the contrast bolus. This means that these parameters only weakly contribute to the loss. Initial experiments illustrated that the fit had trouble converging to the correct solution when the initial guesses for v_p and t_0 were far from their optimal values. Hence, to prevent the fitting process from converging to these local minima, we implemented a three-step fitting strategy. First, to obtain an initial estimate of t_0, a voxel-wise algorithm was developed, which fits a modified hyperbolic tangent function to the first 30 seconds of the dynamic CG-SENSE image series after the injection of contrast agent. Then, using these initial t_0 estimates and an initial estimate $v_p = 0.03\,[fraction]$, a refined parameter t_0 and v_p were estimated by applying MBR to temporally truncated k-space data that included samples up to one minute after the signal peak. Finally, with the preliminary v_p and t_0 maps, we applied MBR to estimate all parameters (including v_p and t_0) from the entire dataset. In each step, v_e and k_{ep} were initialized with maps containing a constant value across the entire volume (with $v_e = 0.2\,[fraction]$ and $k_{ep} = 1.0\,[min^{-1}]$).

We repeated this experiment for each temporal resolution described in Section 4.2.3.2. In each run, we adjusted the weighting factor γ, as described in Appendix 4.A, with specific values selected using the procedure explained in Appendix 4.B. For the comparison, we performed image-based fitting of the simulated data, reconstructing image series with the same temporal resolutions. These image series were reconstructed with the GRASP method that combined Parallel Imaging and Compressed Sensing with regularization in the temporal domain (TV $\lambda = 0.01$).

4.2.4.2 In vivo data

Initial applications of MBR to in vivo data revealed that directly translating the phantom-based experiment did not yield comparable results. Due to different noise characteristics and potential mismatch of the extended Tofts model and true (more complex) tissue microstructural behavior, the optimisation problem has become more complex, with a higher chance of ending in local minima. Therefore, in vivo, we used an initial guess based on the pharmacokinetic map estimates derived from an NLLS fit performed on the CG-SENSE image series. Since the NLLS-generated maps were typically noisy, a histogram-based cutoff was applied to remove background noise. These refined initial guesses were then used to perform a one-step MBR fit on the full k-space data, with a temporal resolution of 2.0 s/frame (13 readouts, $R = 31$). As a comparison, image-based fitting was also performed on the in vivo data, reconstructing image series at the same temporal resolution. The dynamic series were reconstructed with the CG-SENSE algorithm implemented in PyQMRI at the temporal resolution of 2.0 s/frame (13 readouts, $R = 31$).

4.2.5 Evaluation

4.2.5.1 Digital phantom

We evaluated the performance of both methods in terms of accuracy, measured by the mean error (ME) between the simulated parameter maps and the estimated parameter maps, and precision, represented by the standard deviation (STD) of the error. ME and STD were calculated for each tissue and for each temporal resolution. For liver, pancreas, and spleen, we plotted the distribution of errors.

4.2.5.2 In vivo data

This data is meant as proof of principle. Since ground truth pharmacokinetic parameters are unavailable for in vivo data, we initially conducted a qualitative comparison between the reconstructed maps obtained using MBR and those generated with the image-based fitting algorithm. We further assessed the map quality of both methods quantitatively using two metrics, one for sharpness and one for noise. For sharpness, we selected a region of interest (ROI) in the liver that included both tissue and vessels, as vessels have good contrast, but are small and hard to depict. Within this ROI, we used the blur metric defined in [56], where higher blur values indicate reduced image sharpness and hence reduced image quality. For quantifying noise, we selected a homogeneous ROI in the liver (excluding visible vessels), and calculated STD of parameter values. A lower standard deviation in

this region was interpreted as reduced noise, and therefore, higher map precision.

4.3 Results

MBR produced sharper, less noisy and visually more pleasing parameter maps for all parameters across all different temporal resolutions (Figures 4.1 (v_p and K^{trans}) and 4.2 (v_e, k_{ep}, t_0)). In image-based fitting, the temporal resolution was defined during image reconstruction, after which pharmacokinetic parameters were fitted. In contrast, MBR incorporated temporal resolution intrinsically during the joint optimisation process of Eq.4.5.

At the higher temporal resolution (1.5 s/frame), MBR produced less noisy maps compared to image-based fitting, due to its ability to better utilise temporal information. The most notable improvement was seen in the estimation of v_p (yellow arrows on Figure 4.1), where image-based fitting exhibited radial undersampling artifacts (so-called 'streaking' artifacts, visible as pink and purple stripes), whereas MBR had no streaking artifacts as it mitigates the need for generating intermediate images. Furthermore, fine anatomical details, such as stomach wall and vessels in the k_{ep} maps, were more clearly defined with MBR (green arrows on Figure 4.2). It is important to note that the estimated k_{ep} values in the vessels were higher than the ground truth for both methods. This discrepancy arises from the ill-posed nature of the problem in this region, where the ground truth values for both v_e and k_{ep} are effectively zero.

MBR also demonstrated robustness across different temporal resolutions. Parameter maps at 2.4 s/frame and 3.9 s/frame closely resembled those at 1.5 s/frame, indicating that MBR preserved image quality and resolution even as temporal resolution decreased — something conventional methods struggled to maintain (white arrows on Figure 4.2). At lower temporal resolutions (16.5 s/frame), some MBR-derived parameter maps, particularly v_p, show a bias, but overall, even at these time frames, the MBR performance remained more accurate and sharper than the least squares fitting.

Overall, MBR produced more accurate and precise maps in most organ-parameter combinations (example organs shown in Figure 4.3). Specifically, MBR yielded more accurate estimates in 41/52, 43/52, 15/52, 41/52 tissue/parameter combinations (13 tissues and 4 parameters in total) at temporal resolutions of 1.5 s/frame, 2.4 s/frame, 3.9 s/frame, and 16.5 s/frame, respectively (Table 4.3). In terms of precision, MBR outperformed image-based fitting in 52/52, 52/52, 44/52, 30/52 tissue/parameter combinations at temporal resolutions of 1.5 s/frame, 2.4 s/frame, 3.9 s/frame, and 16.5 s/frame respectively.

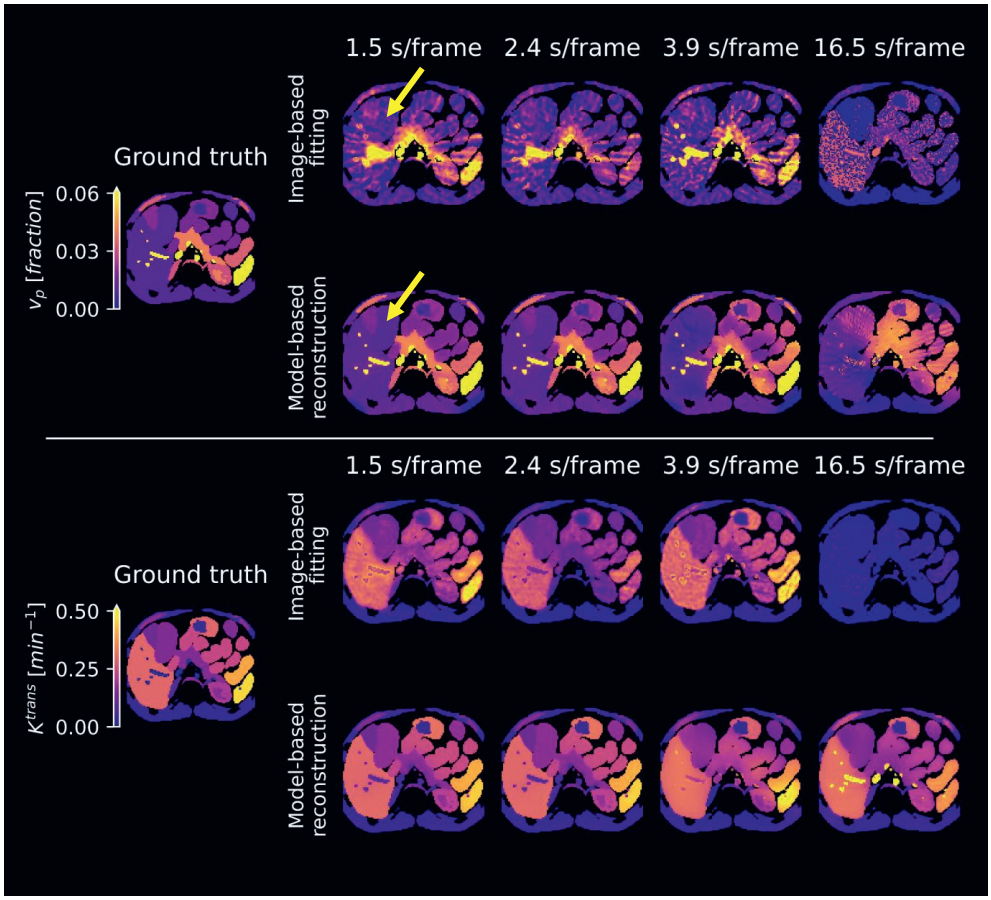

Figure 4.1: v_p and K^{trans} maps of XCAT phantom estimated using image-based fitting (rows 1 and 3) and MBR (rows 2 and 4) at different temporal resolutions. The images at the far left represent the ground truth references. Per method row, the left side represents the highest temporal resolution (highest acceleration rate), while the right side corresponds to the lowest temporal resolution (lowest acceleration rate). The yellow arrow highlights the difference between the two methods, showing that parameter v_p exhibits streaking artifacts with image-based fitting, whereas MBR yields artifact-free results.

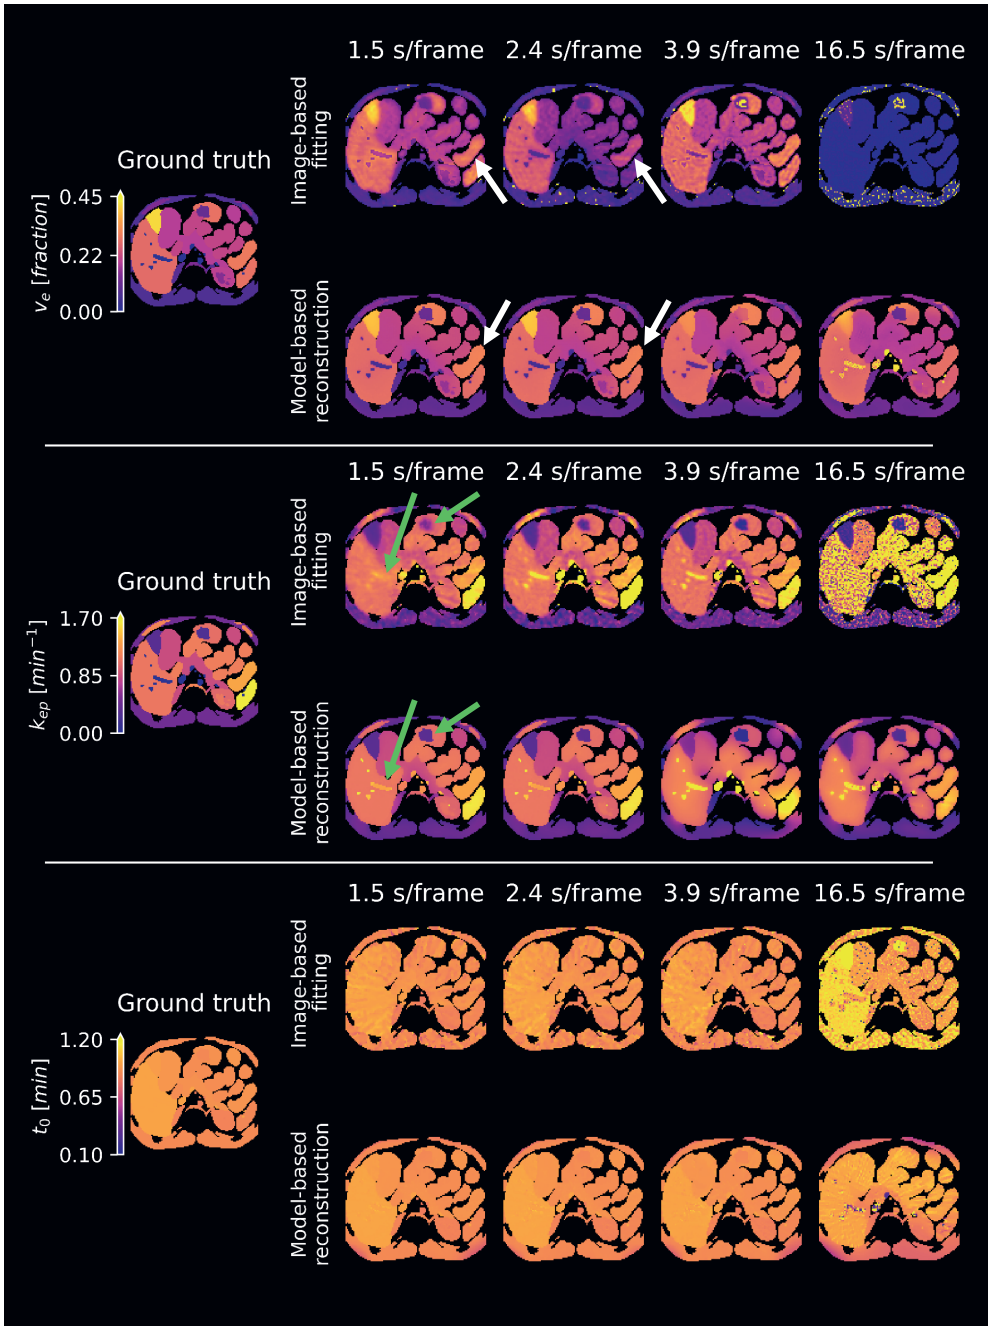

Figure 4.2: v_e, k_{ep}, and t_0 maps of XCAT phantom estimated using image-based fitting (rows 1, 3, and 5) and MBR (rows 2, 4, and 6) at different temporal resolutions. The images at the far left represent the ground truth references. Per method row, the left side represents the highest temporal resolution (highest acceleration rate), while the right side corresponds to the lowest temporal resolution (lowest acceleration rate). The white arrow indicates that MBR maintained image quality even at high acceleration rates, in contrast to image-based fitting. The green arrow demonstrates that MBR better preserved fine anatomical details, such as vessels.

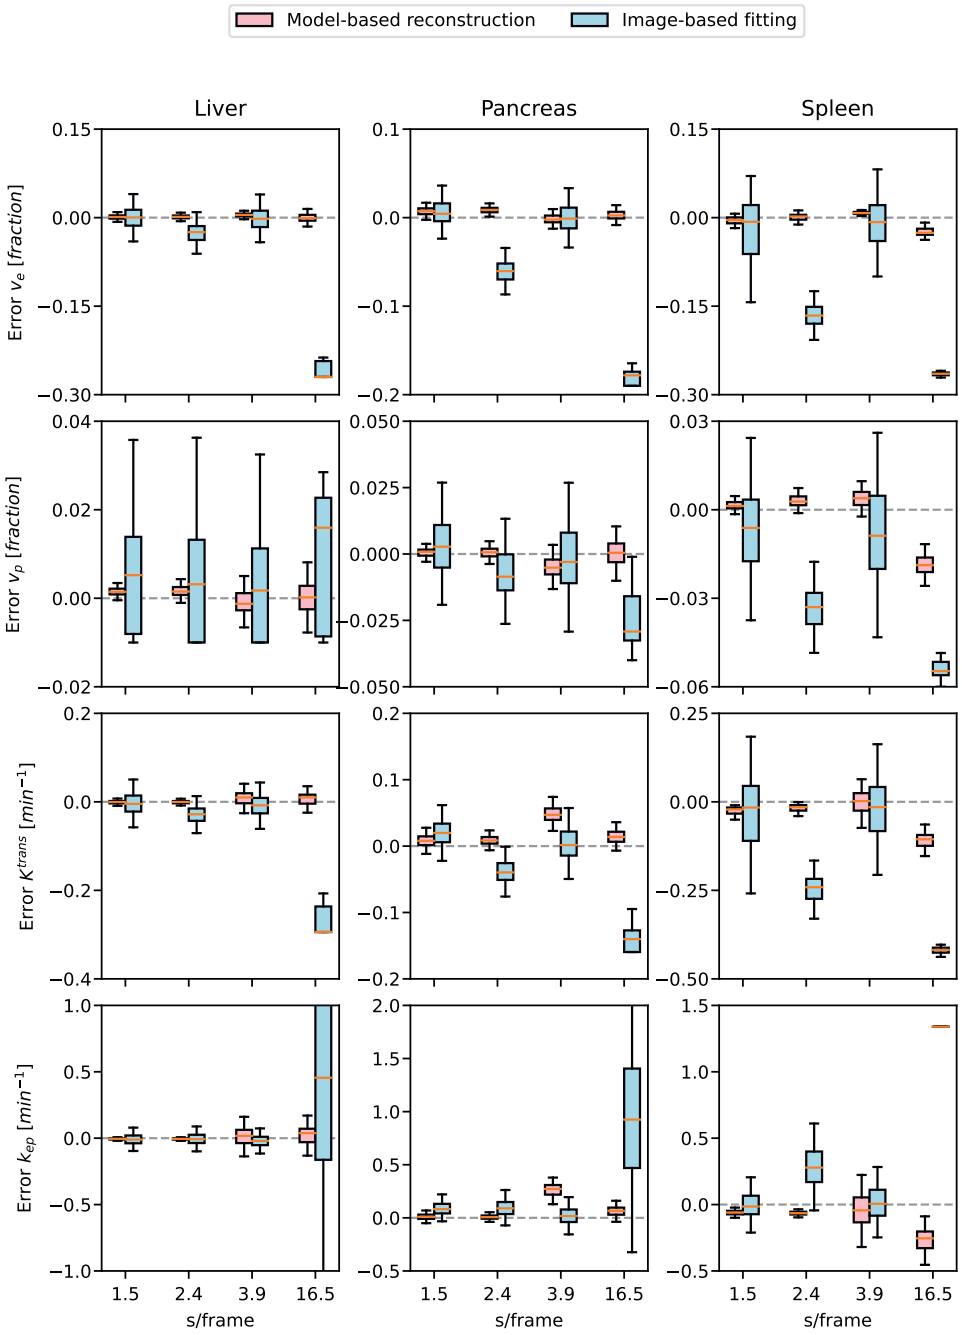

Figure 4.3: The distribution of errors in v_e, v_p, K^{trans}, and k_{ep} maps in the liver, pancreas, and spleen using MBR (pink) and image-based fitting (blue) at different temporal resolutions. The gray dashed line indicates zero error.

Table 4.3: Comparison of MBR and image-based fitting for all simulated tissues, temporal resolutions, and parameters. Simulations included 13 tissues in total.

Percentage of tissues where MBR outperformed NLLS	Parameter	Temporal resolution			
		1.5 s/frame	2.4 s/frame	3.9 s/frame	16.5 s/frame
Accuracy (ME)	v_p	85%	92%	46%	69%
	K^{trans}	85%	92%	38%	92%
	v_e	85%	85%	23%	92%
	k_{ep}	62%	62%	8%	62%
Precision (SD)	v_p	100%	100%	100%	69%
	K^{trans}	100%	100%	100%	23%
	v_e	100%	100%	92%	38%
	k_{ep}	100%	100%	46%	100%

Figure 4.4 presents axial and coronal views of pharmacokinetic maps of a healthy abdomen estimated using both image-based fitting method and MBR at a k-space temporal resolution of 2.0 s/frame (13 readouts, $R = 31$). Compared to the image-based fitting, MBR demonstrated a slight improvement in noise suppression and enhanced vessel contrast, particularly noticeable in the axial slice of the v_e and K^{trans} parameters (indicated by yellow arrows on Figure 4.4. However, similar to simulations, a slight discrepancy in the v_e parameter was observed in the liver with MBR compared to conventional fitting. Unlike the simulation results, we did not notice an improvement in the MBR v_p map.

ROI-based analysis (ROI is shown with a white rectangle on the axial view on Figure 4.4) confirmed that MBR generally produced sharper parameter maps. For example, the blur metric - where values range from 0 (least blur) to 1 (most blur) - was lower for MBR compared to image-based fitting for both v_e (0.238 vs. 0.318) and K^{trans} (0.206 vs. 0.283). However, for k_{ep} MBR showed an increase in blur (0.550 vs. 0.375), consistent with the visually more smoothed appearance of this parameter map in Figure 4.4. Analysis of a homogeneous ROI in liver tissue (ROI is shown with a light blue rectangle on the coronal view on Figure 4.4) demonstrated that MBR also yields more precise parameter maps with lower STD across all parameters: for v_e (0.041 vs. 0.052), K^{trans} (0.082 vs. 0.113), k_{ep} (0.086 vs. 0.139). We excluded v_p from the analysis due to its very low estimated values in the liver.

4.4 Discussion

In this work, we show that, by directly estimating pharmacokinetic parameters from k-space, MBR can overcome the spatio-temporal trade-off typically pestering DCE-MRI. In

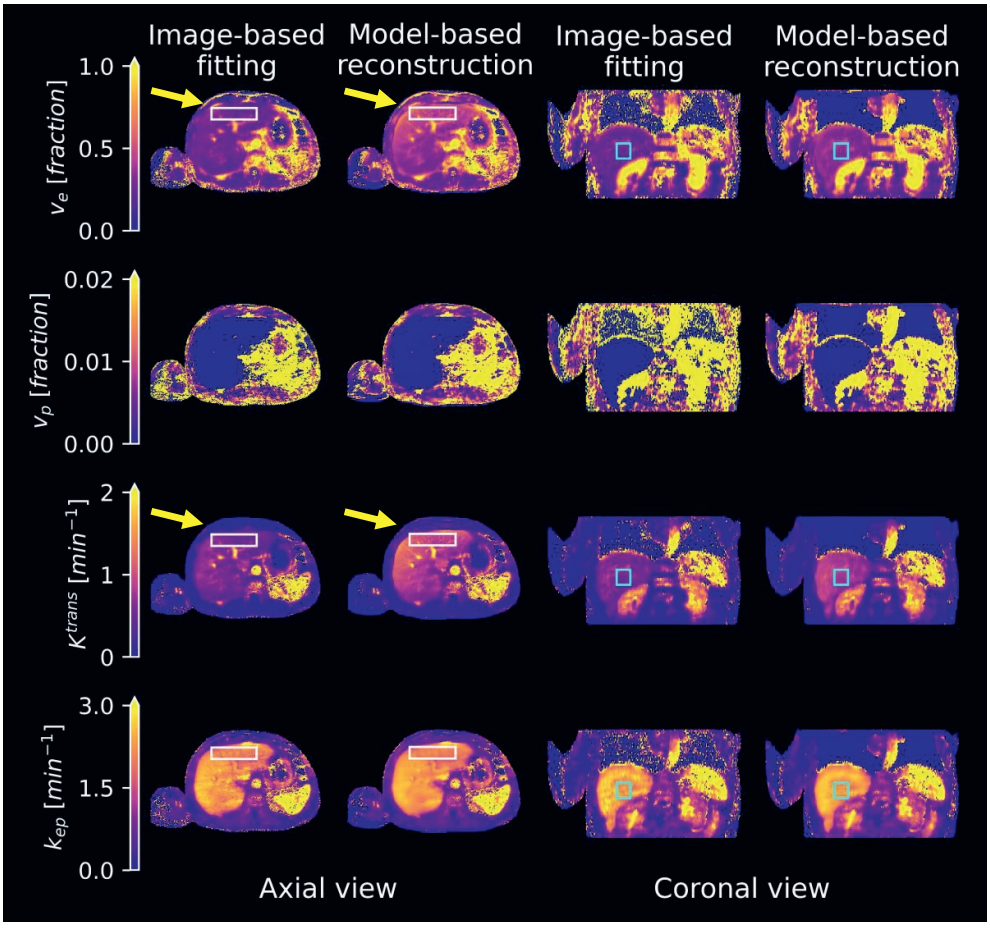

Figure 4.4: Pharmacokinetic parameter maps of a healthy abdomen estimated using image-based fitting method and MBR. The k-space data bins corresponded to the temporal resolution of 2.0 s/frame. In the axial view, the white rectangle marks a heterogeneous ROI used to assess map sharpness, while the light blue rectangle indicates a homogeneous ROI for evaluating parameter precision.

simulations, MBR allows for more accurate and precise images than conventional image-based fitting, and contrary to image-based methods, it does not suffer from streaking artifacts at high acceleration rates. In vivo, MBR resulted in sharper and and more precise parameter maps. To our knowledge, we are the first to implement MBR of extended Tofts DCE with a TGV regularisation (`https://gitlab.tugraz.at/ibi/mrirecon/software/PyQMRI/-/blob/master/pyqmri/models/DCE_fit_M0_numpy.py`) and the first to perform extensive testing and validation MBR for extended Tofts DCE in simulations. Our work shows how MBR is unhindered by temporal resolution and allows obtaining the high-quality perfusion maps clinicians desire to see. This can be the step that is needed to introduce DCE into a clinical workflow. However, we would like to highlight that despite the benefit of MBR for DCE, there are several challenges including a long reconstruction times and its sensitivity to hyperparameter tweaking.

Current clinical DCE-MRI protocols often include only a few temporal phases (typically pre-contrast, arterial, and washout), limiting their value for quantitative pharmacokinetic modeling. On many systems, DCE is acquired using radial trajectories and reconstructed at low temporal resolution to yield motion-robust, high-SNR images. While this suits qualitative interpretation, it suppresses the dynamic signal fluctuations needed for accurate modeling. If quantitative DCE-MRI is to enter clinical practice, these routinely acquired datasets represent an important starting point. However, higher temporal resolution is essential. Existing CS approaches improve image quality through strong temporal regularization but inadvertently dampen physiologically relevant dynamics. Our MBR approach directly addresses this limitation by reconstructing data at high temporal resolution without relying on temporal regularization, thereby preserving kinetic detail crucial for quantitative analysis. In doing so, it helps bridge a major technical gap hindering clinical translation of DCE-MRI.

Using a digital phantom, we showed that MBR generated more visually appealing pharmacokinetic maps, with accuracy and precission comparable to or exceeding that of traditional image-based fitting (Figures 4.1 and 4.2, and Table 4.3). Furthermore, we confirmed our hypothesis that MBR is more robust to variations in temporal resolution compared to image-based approaches: the performance of image-based method shows streaking artifacts (yellow arrow on Figure 4.1), edge enhancements (green arrow on Figure 4.2), and blurring at lower temporal resolutions (white arrow on Figure 4.2) whereas MBR shows consistent image quality even at 1.5 s/frame. In vivo, MBR also offered improvements by reducing noise and enhancing the visibility of fine details, such as vessels.

A key advantage of MBR lies in its ability to preserve k-space data fidelity more effectively than image-based methods, as it directly optimises quantitative maps from raw k-

space measurements. This direct reconstruction also helps overcome the spatio-temporal trade-off inherent in dynamic imaging, and is particularly important in cancer imaging, where both high temporal resolution (for pharmacokinetic modelling) and high spatial resolution (for precise tumour characterisation) are critical. Given the known limitations in repeatability of DCE-MRI [57], improving reconstruction techniques - such as MBR - represents an important step toward generating more reliable and informative pharmacokinetic maps. In contrast, image-based approaches rely on a two-step process - first reconstructing images and then estimating quantitative parameters - which can introduce cumulative errors in the final maps.

This direct relation to k-space, without the need for intermediate image reconstructions, also opens up full freedom in k-space sampling schemes. Potentially, there are more efficient sampling schemes possible that better capture the pharmacokinetics. However, current DCE-MRI protocols are typically optimised to balance spatial resolution, temporal resolution, and SNR of intermediate images. This trade-off results in moderate spatial and temporal resolution to maintain diagnostic quality. In this study, we aimed to challenge this balance with MBR. However, we were limited by the available in vivo datasets, which were acquired using conventional protocols and did not include data with different spatial-, temporal resolutions, and SNR. As a result, our ability to fully explore the potential of MBR in breaking the spatio-temporal trade-off was constrained. This limitation may also partially explain the relatively modest improvements observed in the in vivo results compared to simulations, where we had more control over sampling strategies and noise properties.

Beyond acquisition constraints, variations in MBR's performance between the digital phantom and in vivo data are likely influenced by additional real-world factors not present in simulations. For example, complexities such as imperfect fat suppression, B_1 inhomogeneities, and subject motion were not accounted for in the simulations. The extended Tofts model, while effective for controlled settings, simplifies pharmacokinetic behavior and may not fully capture the more intricate blood flow dynamics found in vivo. Additionally, the more complex morphology of internal organs in vivo may reduce the effectiveness of regularisation techniques. Despite these challenges, the digital phantom remains a powerful tool for validating new algorithms and establishing proof-of-concept. The strong performance of MBR in simulations confirms its theoretical advantages, while its consistent - albeit more modest - in vivo improvements demonstrate its practical feasibility and robustness under realistic conditions. These results suggest that with further refinement, MBR holds strong potential for clinical translation.

One challenge we encountered with MBR is the complex DCE loss surface characterized by multiple local minima. To address this, we provided initial guesses close to the ex-

pected optimum to reduce the risk of convergence to poor local minima. In practice, this meant using an NLLS-based initialization for the in vivo data and a three-step fitting strategy for the simulations. To highlight the complex loss function for a single voxel, Figure 4.5 shows a 2D projection of a 4D DCE loss function for a simulated single voxel with parameters $v_p = 0.05\,[fraction]$, $v_e = 0.1\,[fraction]$, $k_{ep} = 0.5\,[min^{-1}]$, $t_0 = 1.0\,[min]$. Simulated noise was added, sampled from a normal distribution with a standard deviation of 0.1. This example illustrates how, in a single voxel, the NLLS fit converges to different optimal solutions depending on the initial parameter guesses. Note that this projection fixes two parameters (v_e and k_{ep}) to the solution while only fitting the other two (v_p and t_0), reducing the actual degrees of freedom by half. MBR will fit five parameters multiplied by the number of voxels (e.g. 128*128*5=81,920 parameters). Unlike NLLS, MBR applies regularisation that smooths the loss surface, reducing the chance of being trapped in local minima. However, when the loss landscape is highly complex, even regularisation may struggle to guide the optimisation to the global minimum. In NLLS fitting, local minima occur independently at each voxel, often contributing to the characteristic noise observed in NLLS-based parameter maps. In MBR, on the other hand, convergence to a local minimum affects the entire image, and could cause a shift in intensity in a whole organ. This may be a reason for the offset in K^{trans} in the 3.9 s/frame simulations (Figure 4.1). Note that the NLLS is just an initialisation and MBR re-estimates all pharmacokinetic parameters jointly from k-space data, thereby incorporating additional data consistency constraints beyond what is available in image-based fitting.

The improved noise suppression observed in both simulations and in vivo data, along with the higher parameter precision in simulations, can partially be attributed to the TGV regularisation implemented in PyQMRI. TGV effectively balances noise reduction and edge preservation, leading to smoother and more stable pharmacokinetic maps. However, it is well-established that regularization techniques, including TGV, can introduce biases in the estimated parameters [58]. For example, in our study, we noted a bias in the v_e maps in vivo. While the regularization strength was optimized to balance noise suppression and detail preservation, the inherent trade-off between these objectives could have led to a slight underestimation or overestimation of v_e in specific regions.

4.4.1 Limitations

DCE-MRI is a dynamic acquisition technique and, consequently, it is susceptible to respiratory, bulk and possibly cardiac motion. In the digital phantom, motion was not simulated. However this does not affect the comparison between image-based fitting and MBR, as both methods were evaluated under the same conditions. While radial acquisition is

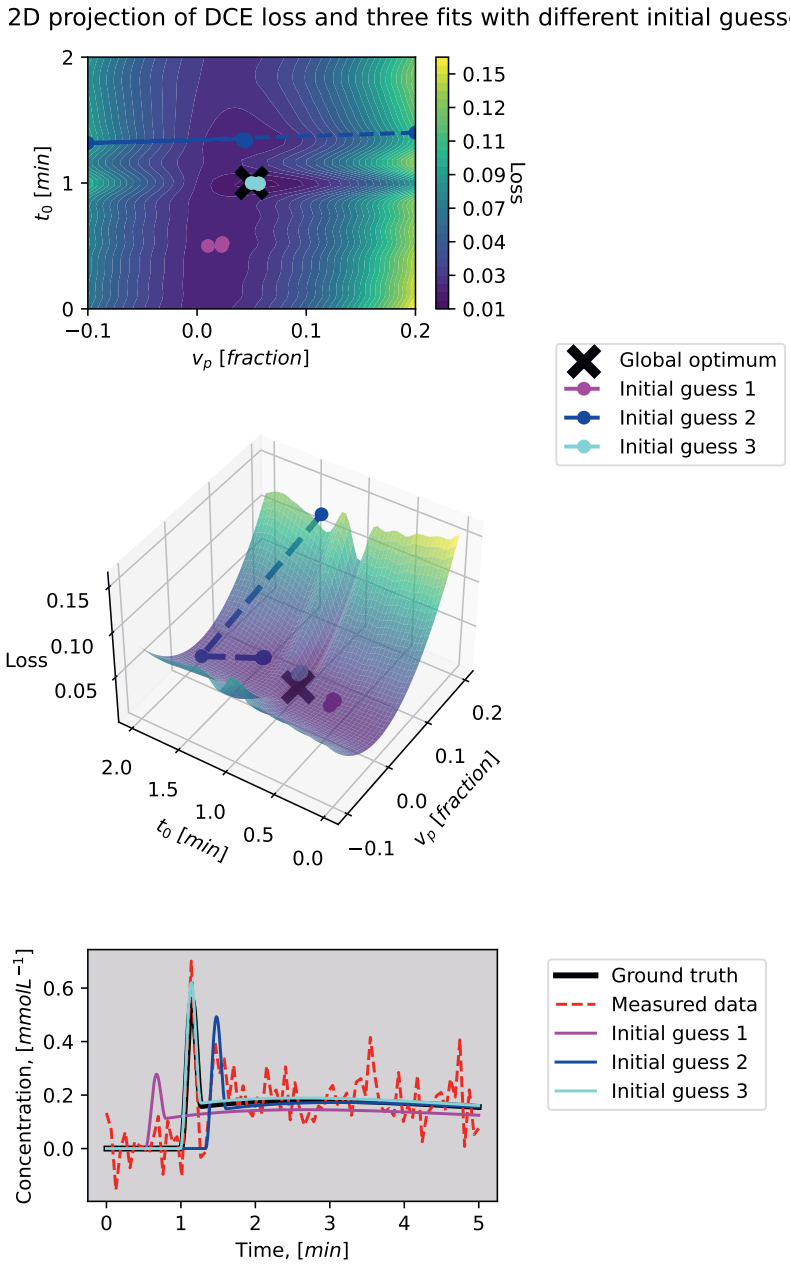

Figure 4.5: A 2D slice of the 4D DCE loss function for a simulated voxel with ground truth parameters: $v_p = 0.05\ [fraction]$, $v_e = 0.1\ [fraction]$, $k_{ep} = 0.5\ [min^{-1}]$, $t_0 = 1.0\ [min]$. The NLLS iteration paths and solutions obtained from different initial guesses are also displayed. As this is a 2D slice in 4D space, v_e and k_{ep} were fixed, while v_p and t_0 were fitted. The first panel shows a 2D view of the loss surface with NLLS solutions, the second panel presents the same in 3D for better visualization of the loss landscape, and the third panel compares the ground truth curve (based on simulated parameters), the measured curve used for fitting, and three resulting fits. This figure illustrates how initial guesses can lead to convergence to different local minima, highlighting the complexity of the loss function landscape.

known to be more robust to motion [59], residual artifacts can still appear, particularly in abdominal scans. No motion correction was applied, which may have impacted parameter map quality. In the future, incorporating a motion correction operator into the forward operator $A(u)$ could help mitigate motion-related artifacts. Implementing such an operator falls outside the scope of this article.

A drawback of MBR is the long computation time associated with the technique. For example, the run time of in vivo experiment was 4 days and 18 hours for a 5-minute 3D DCE-MRI acquisition with full liver coverage (scan details are provided in Table 4.2). In this study, we focused on accuracy rather than on speed. To ensure convergence in all cases, we used 10 Gauss-Newton iterations. To accelerate the reconstruction, one can reduce the number of Gauss-Newton iterations and/or the number of maximum Primal-Dual iterations. This could serve as a preliminary step towards faster reconstruction.

In this study, where we focused on the technical feasibility of using MBR for DCE, we chose the extended Tofts model in this work, as it remains one of the most widely used pharmacokinetic models in oncologic DCE-MRI, and has demonstrated predictive value for treatment response across different tumor sites [60–62]. Tumors are inherently heterogeneous, often containing regions with varying perfusion and vascular permeability, and the extended Tofts provides a pragmatic balance between physiological interpretability and numerical stability. More complex multi-compartment models can, in principle, better describe certain tissue characteristics [63], but they increase computational complexity and the risk of non-unique or unstable parameter estimation [64]. However, the in vivo data used in this study included only healthy abdominal organs, such as the liver and kidneys - organs for which the extended Tofts model is not ideally suited. For example, liver imaging typically requires a dual-inlet input function to accurately reflect its dual blood supply. As a result, the pharmacokinetic parameters estimated in the liver may be subject to bias. In addition, we used a population-based AIF, which assumes a certain injection rate, dose, and vascular dynamics, thereby neglecting individual variability. This simplification can introduce errors in AIF characterisation and subsequent parameter estimation. While patient-specific AIFs can mitigate these issues, their accurate measurement is technically demanding due to the need for high temporal resolution and sufficient SNR, as well as potential artefacts from inflow effects, partial-volume averaging, and flow pulsatility. As the model and AIF-related limitations affected both reconstruction approaches equally, it does not affect our primary goal, which was to show the technical feasibility of using MBR for DCE and compare it to existing state-of-the-art compressed sensing. Finally, it should be noted that the use of fat suppression introduces a non-steady-state signal behaviour, as described in [65], where an alternative signal equation accounting for

this effect has been proposed. However, this formulation is not yet widely adopted in the DCE-MRI community and, importantly, is not analytically invertible, making it less practical for applications where T_1 must be estimated from the measured signal. Therefore, we opted to use the conventional forward T_1 signal model that is commonly employed in DCE-MRI studies.

While the MBR framework was implemented using the extended Tofts model, it is inherently modular and can, in principle, support other pharmacokinetic models such as the standard Tofts, Patlak, or dual-input models. However, replacing the model is non-trivial and would require extensive adaptation of the Jacobian, re-optimisation of regularisation parameters, and rigorous validation to ensure convergence and avoid local minima. For these reasons, extending the framework to additional pharmacokinetic models falls outside the scope of this work, which focuses on feasibility and methodological comparison. Nevertheless, the framework is modular, and future work could explore integration of alternative models where clinically justified.

In this work, we demonstrated the in vivo feasibility of the MBR framework using data from a single healthy subject. While this serves as an important proof-of-concept, a more comprehensive evaluation will require testing across a broader cohort, including multiple healthy volunteers and patients with varying pathologies. Expanding the study population will allow us to assess the generalisability, robustness, and clinical relevance of the method under diverse physiological and pathological conditions.

Chapter appendix

4.A Configuration settings for PyQMRI

The configuration file used as input to PyQMRI framework. Depending on the temporal resolution or type of data, the regularization parameters gamma and gamma min were adjusted according to Table 4.A.1.

```
[TGV]
max_iters = 500
start_iters = 20
max_gn_it = 10
lambd = 1.0
gamma = 0.09
delta = 1.0
omega = 0.0
display_iterations = False
gamma_min = 1e-8
delta_max = 10000
omega_min = 0.0
rtol = 1e-06
atol = 1e-08
stag = 1e-10
delta_inc = 10.0
gamma_dec = 0.5
omega_dec = 0.5
beta = 1
```

4.B Selection of weighting factors

To determine suitable values, we empirically evaluated different combinations of γ and γ_{min}. In the first step of the MBR pipeline, when estimating the initial v_p map, we computed the structural similarity metric (SSIM) for v_p between reconstructions and ground

Table 4.A.1: Adapted parameters gamma and gamma_min for different temporal resolutions or type of data

Data	Temporal resolution	gamma	gamma_min
Phantom	1.5 s/frame	0.09	1e-08
Phantom	2.4 s/frame	0.07	1e-08
Phantom	3.9 s/frame	0.06	1e-03
Phantom	16.5 s/frame	0.02	1e-02
In vivo	2.0 s/frame	0.07	0.01

truth. We then selected γ-γ_{min} combinations that achieved both a high SSIM and a visually acceptable results. In the second step, which involves full pharmacokinetic fitting, we conducted a more comprehensive evaluation by computing multiple metrics for all four model parameters: mean absolute error, standard deviation, SSIM, peak SNR, and blur; and ranked the regularization parameter combinations accordingly. As the ranking sets equal weight to each attribute and parameter, it will select good algorithms, but its best is somewhat arbitrary; therefore the top-performing configurations were visually inspected and the combination with the best overall quantitative performance and visual quality was chosen for the final reconstruction.

Bibliography

[1] P. Carmeliet and R. K. Jain, "Angiogenesis in cancer and other diseases," *nature*, vol. 407, no. 6801, pp. 249–257, 2000.

[2] W. W. Li, "Tumor angiogenesis: molecular pathology, therapeutic targeting, and imaging," *Academic radiology*, vol. 7, no. 10, pp. 800–811, 2000.

[3] B. Muz, P. De La Puente, F. Azab, and A. Kareem Azab, "The role of hypoxia in cancer progression, angiogenesis, metastasis, and resistance to therapy," *Hypoxia*, pp. 83–92, 2015.

[4] G. L. Semenza, "Defining the role of hypoxia-inducible factor 1 in cancer biology and therapeutics," *Oncogene*, vol. 29, no. 5, pp. 625–634, 2010.

[5] M. Hockel and P. Vaupel, "Tumor hypoxia: definitions and current clinical, biologic, and molecular aspects," *Journal of the National Cancer Institute*, vol. 93, no. 4, pp. 266–276, 2001.

[6] R. K. Do, M. Reyngold, R. Paudyal, J. H. Oh, A. S. Konar, E. LoCastro, K. A. Goodman, and A. Shukla-Dave, "Diffusion-weighted and dynamic contrast-enhanced mri derived imaging metrics for stereotactic body radiotherapy of pancreatic ductal adenocarcinoma: preliminary findings," *Tomography*, vol. 6, no. 2, p. 261, 2020.

[7] S. J. Hectors, S. Lewis, P. Kennedy, O. Bane, D. Said, M. Segall, M. Schwartz, E. Kim, and B. Taouli, "Assessment of hepatocellular carcinoma response to 90y radioembolization using dynamic contrast material–enhanced mri and intravoxel incoherent motion diffusion-weighted imaging," *Radiology: Imaging Cancer*, vol. 2, no. 4, p. e190094, 2020.

[8] S. E. Heethuis, L. Goense, P. S. van Rossum, A. S. Borggreve, S. Mook, F. E. Voncken, A. Bartels-Rutten, B. M. Aleman, R. van Hillegersberg, J. P. Ruurda *et al.*, "Dw-mri and dce-mri are of complementary value in predicting pathologic response to neoadjuvant chemoradiotherapy for esophageal cancer," *Acta Oncologica*, vol. 57, no. 9, pp. 1201–1208, 2018.

[9] A. M. Hötker, L. Tarlinton, Y. Mazaheri, K. M. Woo, M. Gönen, L. B. Saltz, K. A. Goodman, J. Garcia-Aguilar, and M. J. Gollub, "Multiparametric mri in the assessment of

response of rectal cancer to neoadjuvant chemoradiotherapy: A comparison of morphological, volumetric and functional mri parameters," *European radiology*, vol. 26, pp. 4303–4312, 2016.

[10] W. Yang, J. W. Qiang, H. P. Tian, B. Chen, A. J. Wang, and J. G. Zhao, "Multi-parametric mri in cervical cancer: early prediction of response to concurrent chemoradiotherapy in combination with clinical prognostic factors," *European Radiology*, vol. 28, pp. 437–445, 2018.

[11] C. Wang, J. K. Horton, F.-F. Yin, and Z. Chang, "Assessment of treatment response with diffusion-weighted mri and dynamic contrast-enhanced mri in patients with early-stage breast cancer treated with single-dose preoperative radiotherapy: initial results," *Technology in cancer research & treatment*, vol. 15, no. 5, pp. 651–660, 2016.

[12] R. L. Jensen, M. L. Mumert, D. L. Gillespie, A. Y. Kinney, M. C. Schabel, and K. L. Salzman, "Preoperative dynamic contrast-enhanced mri correlates with molecular markers of hypoxia and vascularity in specific areas of intratumoral microenvironment and is predictive of patient outcome," *Neuro-oncology*, vol. 16, no. 2, pp. 280–291, 2014.

[13] M. Rata, D. J. Collins, J. Darcy, C. Messiou, N. Tunariu, N. Desouza, H. Young, M. O. Leach, and M. R. Orton, "Assessment of repeatability and treatment response in early phase clinical trials using dce-mri: comparison of parametric analysis using mr-and ct-derived arterial input functions," *European radiology*, vol. 26, pp. 1991–1998, 2016.

[14] J. J. Van Vaals, M. E. Brummer, W. Thomas Dixon, H. H. Tuithof, H. Engels, R. C. Nelson, B. M. Gerety, J. L. Chezmar, and J. A. Den Boer, ""keyhole" method for accelerating imaging of contrast agent uptake," *Journal of Magnetic Resonance Imaging*, vol. 3, no. 4, p. 671–675, Jul 1993.

[15] L. A. Tudorica, K. Y. Oh, N. Roy, M. D. Kettler, Y. Chen, S. L. Hemmingson, A. Afzal, J. W. Grinstead, G. Laub, X. Li *et al.*, "A feasible high spatiotemporal resolution breast dce-mri protocol for clinical settings," *Magnetic resonance imaging*, vol. 30, no. 9, pp. 1257–1267, 2012.

[16] J. Chen, C. Huang, K. Shanbhogue, D. Xia, M. Bruno, Y. Huang, K. T. Block, H. Chandarana, and L. Feng, "Dce-mri of the liver with sub-second temporal resolution using grasp-pro with navi-stack-of-stars sampling," *NMR in Biomedicine*, vol. 37, no. 12, p. e5262, 2024.

[17] O. Arif, H. Afzal, H. Abbas, M. F. Amjad, J. Wan, and R. Nawaz, "Accelerated dynamic mri using kernel-based low rank constraint," *Journal of medical systems*, vol. 43, pp. 1–11, 2019.

[18] A. Koçanaoğulları, C. Ariyurek, O. Afacan, and S. Kurugol, "Learning the regularization in dce-mr image reconstruction for functional imaging of kidneys," *IEEE Access*, vol. 10, pp. 4102–4111, 2021.

[19] K. T. Block, M. Uecker, and J. Frahm, "Model-based iterative reconstruction for radial fast spin-echo mri," *IEEE transactions on medical imaging*, vol. 28, no. 11, pp. 1759–1769, 2009.

[20] C. Graff, Z. Li, A. Bilgin, M. Altbach, A. Gmitro, and E. Clarkson, "Iterative t2 estimation from highly undersampled radial fast spin-echo data," in *Proc. Intl. Soc. Mag. Reson. Med*, vol. 14, 2006, p. 925.

[21] T. Hilbert, T. J. Sumpf, E. Weiland, J. Frahm, J.-P. Thiran, R. Meuli, T. Kober, and G. Krueger, "Accelerated t2 mapping combining parallel mri and model-based reconstruction: Grappatini," *Journal of Magnetic Resonance Imaging*, vol. 48, no. 2, pp. 359–368, 2018.

[22] O. Maier, J. Schoormans, M. Schloegl, G. J. Strijkers, A. Lesch, T. Benkert, T. Block, B. F. Coolen, K. Bredies, and R. Stollberger, "Rapid t1 quantification from high resolution 3d data with model-based reconstruction," *Magnetic resonance in medicine*, vol. 81, no. 3, pp. 2072–2089, 2019.

[23] V. Roeloffs, X. Wang, T. J. Sumpf, M. Untenberger, D. Voit, and J. Frahm, "Model-based reconstruction for t1 mapping using single-shot inversion-recovery radial flash," *International Journal of Imaging Systems and Technology*, vol. 26, no. 4, pp. 254–263, 2016.

[24] O. Maier, S. M. Spann, D. Pinter, T. Gattringer, N. Hinteregger, G. G. Thallinger, C. Enzinger, J. Pfeuffer, K. Bredies, and R. Stollberger, "Non-linear fitting with joint spatial regularization in arterial spin labeling," *Medical Image Analysis*, vol. 71, p. 102067, 2021.

[25] M. Huemer, C. Stilianu, O. Maier, M. S. Fabian, M. Schmidt, A. Doerfler, K. Bredies, M. Zaiss, and R. Stollberger, "Improved quantification in cest-mri by joint spatial total generalized variation," *Magnetic Resonance in Medicine*, 2024.

[26] T. Hüfken, J. M. Arbogast, A.-K. Bracher, M. Beer, H. Neubauer, and V. Rasche, "Accelerated model-based quantitative diffusion mri: a feasibility study for musculoskeletal application," *Zeitschrift für Medizinische Physik*, vol. 32, no. 2, pp. 240–247, 2022.

[27] Z. Dong, E. Dai, F. Wang, Z. Zhang, X. Ma, C. Yuan, and H. Guo, "Model-based reconstruction for simultaneous multislice and parallel imaging accelerated multishot diffusion tensor imaging," *Medical physics*, vol. 45, no. 7, pp. 3196–3204, 2018.

[28] F. Knoll, J. G. Raya, R. O. Halloran, S. Baete, E. Sigmund, R. Bammer, T. Block, R. Otazo, and D. K. Sodickson, "A model-based reconstruction for undersampled radial spin-echo dti with variational penalties on the diffusion tensor," *NMR in Biomedicine*, vol. 28, no. 3, pp. 353–366, 2015.

[29] S. S. Rauh, O. Maier, O. J. Gurney-Champion, M. T. Hooijmans, R. Stollberger, A. J. Nederveen, and G. J. Strijkers, "Model-based reconstructions for intravoxel incoherent motion and diffusion tensor imaging parameter map estimations," *NMR in Biomedicine*, vol. 36, no. 8, p. e4927, 2023.

[30] B. K. Felsted, R. T. Whitaker, M. Schabel, and E. V. DiBella, "Model-based reconstruction for undersampled dynamic contrast enhanced mri," in *Medical Imaging 2009: Biomedical Applications in Molecular, Structural, and Functional Imaging*, vol. 7262. SPIE, 2009, pp. 841–850.

[31] N. Dikaios, "Stochastic gradient langevin dynamics for joint parameterization of tracer kinetic models, input functions, and t1 relaxation-times from undersampled k-space dce-mri," *Medical Image Analysis*, vol. 62, p. 101690, 2020.

[32] Y. Guo, S. G. Lingala, Y. Zhu, R. M. Lebel, and K. S. Nayak, "Direct estimation of tracer-kinetic parameter maps from highly undersampled brain dynamic contrast enhanced mri," *Magnetic resonance in medicine*, vol. 78, no. 4, pp. 1566–1578, 2017.

[33] A. Rastogi and P. K. Yalavarthy, "Comparison of iterative parametric and indirect deep learning-based reconstruction methods in highly undersampled dce-mr imaging of the breast," *Medical Physics*, vol. 47, no. 10, pp. 4838–4861, 2020.

[34] K. Bredies, K. Kunisch, and T. Pock, "Total generalized variation," *SIAM Journal on Imaging Sciences*, vol. 3, no. 3, pp. 492–526, 2010.

[35] O. Maier, S. M. Spann, M. Bödenler, and R. Stollberger, "Pyqmri: an accelerated python based quantitative mri toolbox," *Journal of Open Source Software*, vol. 5, no. 56, p. 2727, 2020.

[36] O. Maier, M. Schloegl, K. Bredies, and R. Stollberger, "3d model-based parameter quantification on resource constrained hardware using double-buffering," *Proc. 27th Annu. Meet. ISMRM, Montréal, QC, Canada*, vol. 4839, 2019.

[37] P. S. Tofts, G. Brix, D. L. Buckley, J. L. Evelhoch, E. Henderson, M. V. Knopp, H. B. Larsson, T.-Y. Lee, N. A. Mayr, G. J. Parker *et al.*, "Estimating kinetic parameters from dynamic contrast-enhanced t1-weighted mri of a diffusable tracer: standardized quantities and symbols," *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 10, no. 3, pp. 223–232, 1999.

[38] P. S. Tofts, "Modeling tracer kinetics in dynamic gd-dtpa mr imaging," *Journal of magnetic resonance imaging*, vol. 7, no. 1, pp. 91–101, 1997.

[39] X. Yang and M. V. Knopp, "Quantifying tumor vascular heterogeneity with dynamic contrast-enhanced magnetic resonance imaging: a review," *BioMed Research International*, vol. 2011, no. 1, p. 732848, 2011.

[40] M. R. Orton, J. A. d'Arcy, S. Walker-Samuel, D. J. Hawkes, D. Atkinson, D. J. Collins, and M. O. Leach, "Computationally efficient vascular input function models for quantitative kinetic modelling using dce-mri," *Physics in Medicine & Biology*, vol. 53, no. 5, p. 1225, 2008.

[41] P. J. van Houdt, S. Ragunathan, M. Berks, Z. Ahmed, L. E. Kershaw, O. J. Gurney-Champion, S. Tadimalla, J. Arvidsson, Y. Sun, J. Kallehauge *et al.*, "Contrast-agent-based perfusion mri code repository and testing framework: Ismrm open science initiative for perfusion imaging (osipi)," *Magnetic resonance in medicine*, vol. 91, no. 5, pp. 1774–1786, 2024.

[42] G. J. Parker, C. Roberts, A. Macdonald, G. A. Buonaccorsi, S. Cheung, D. L. Buckley, A. Jackson, Y. Watson, K. Davies, and G. C. Jayson, "Experimentally-derived functional form for a population-averaged high-temporal-resolution arterial input function for dynamic contrast-enhanced mri," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 56, no. 5, pp. 993–1000, 2006.

[43] H. Lu, C. Clingman, X. Golay, and P. C. Van Zijl, "Determining the longitudinal relaxation time (t1) of blood at 3.0 tesla," *Magnetic Resonance in Medicine: an Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 52, no. 3, pp. 679–682, 2004.

[44] F. Knoll, K. Bredies, T. Pock, and R. Stollberger, "Second order total generalized variation (tgv) for mri," *Magnetic resonance in medicine*, vol. 65, no. 2, pp. 480–491, 2011.

[45] K. Bredies, "Recovering piecewise smooth multichannel images by minimization of convex functionals with total generalized variation penalty," in *Efficient Algorithms for Global Optimization Methods in Computer Vision: International Dagstuhl Seminar, Dagstuhl Castle, Germany, November 20-25, 2011, Revised Selected Papers*. Springer, 2014, pp. 44–77.

[46] F. Knoll, M. Holler, T. Koesters, R. Otazo, K. Bredies, and D. K. Sodickson, "Joint mr-pet reconstruction using a multi-channel image regularizer," *IEEE transactions on medical imaging*, vol. 36, no. 1, pp. 1–16, 2016.

[47] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "Sense: sensitivity encoding for fast mri," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–962, 1999.

[48] M. Uecker, T. Hohage, K. T. Block, and J. Frahm, "Image reconstruction by regularized nonlinear inversion—joint estimation of coil sensitivities and image content," *Magnetic Resonance in Medicine*, vol. 60, no. 3, pp. 674–682, 2008.

[49] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *Journal of mathematical imaging and vision*, vol. 40, pp. 120–145, 2011.

[50] Y. Malitsky and T. Pock, "A first-order primal-dual algorithm with linesearch," *SIAM Journal on Optimization*, vol. 28, no. 1, pp. 411–432, 2018.

[51] W. P. Segars, G. Sturgeon, S. Mendonca, J. Grimes, and B. M. Tsui, "4d xcat phantom for multimodality imaging research," *Medical physics*, vol. 37, no. 9, pp. 4902–4915, 2010.

[52] L. Wissmann, C. Santelli, W. P. Segars, and S. Kozerke, "Mrxcat: Realistic numerical phantoms for cardiovascular magnetic resonance," *Journal of Cardiovascular Magnetic Resonance*, vol. 16, pp. 1–11, 2014.

[53] E. Schrauben, "The eric phantom-extra-dimensional respiration and inflow of contrast: an open-source dce digital simulation tool."

[54] N.-N. Sun, C. Liu, X.-L. Ge, and J. Wang, "Dynamic contrast-enhanced mri for advanced esophageal cancer response assessment after concurrent chemoradiotherapy," *Diagnostic and Interventional Radiology*, vol. 24, no. 4, p. 195, 2018.

[55] M. D. Holland, A. Morales, S. Simmons, B. Smith, S. R. Misko, X. Jiang, D. A. Hormuth, C. Christenson, R. P. Koomullil, D. E. Morgan *et al.*, "Disposable point-of-care portable perfusion phantom for quantitative dce-mri," *Medical physics*, vol. 49, no. 1, pp. 271–281, 2022.

[56] F. Crete, T. Dolmiere, P. Ladret, and M. Nicolas, "The blur effect: perception and estimation with a new no-reference perceptual blur metric," in *Human vision and electronic imaging XII*, vol. 6492. SPIE, 2007, pp. 196–206.

[57] R. Klaassen, O. J. Gurney-Champion, J. W. Wilmink, M. G. Besselink, M. R. Engelbrecht, J. Stoker, A. J. Nederveen, and H. W. van Laarhoven, "Repeatability and correlations of dynamic contrast enhanced and t2* mri in patients with advanced pancreatic ductal adenocarcinoma," *Magnetic resonance imaging*, vol. 50, pp. 1–9, 2018.

[58] J. A. Fessler, "Model-based image reconstruction for mri," *IEEE signal processing magazine*, vol. 27, no. 4, pp. 81–89, 2010.

[59] L. Feng, "Golden-angle radial mri: basics, advances, and applications," *Journal of Magnetic Resonance Imaging*, vol. 56, no. 1, pp. 45–62, 2022.

[60] X. Li, L. R. Arlinghaus, G. D. Ayers, A. B. Chakravarthy, R. G. Abramson, V. G. Abramson, N. Atuegwu, J. Farley, I. A. Mayer, M. C. Kelley *et al.*, "Dce-mri analysis methods for predicting the response of breast cancer to neoadjuvant chemotherapy: Pilot study findings," *Magnetic resonance in medicine*, vol. 71, no. 4, pp. 1592–1602, 2014.

[61] B. Liu, Z. Sun, W.-L. Ma, J. Ren, G.-W. Zhang, M.-Q. Wei, W.-H. Hou, B.-X. Hou, L.-C. Wei, Y. Huan *et al.*, "Dce-mri quantitative parameters as predictors of treatment response in patients with locally advanced cervical squamous cell carcinoma underwent ccrt," *Frontiers in Oncology*, vol. 10, p. 585738, 2020.

[62] M. Vivoda Tomšič, P. Korošec, V. Kovač, S. Bisdas, and K. Šurlan Popovič, "Dynamic contrast-enhanced mri in malignant pleural mesothelioma: prediction of outcome based on dce-mri measurements in patients undergoing cytotoxic chemotherapy," *BMC cancer*, vol. 22, no. 1, p. 191, 2022.

[63] S. P. Sourbron and D. L. Buckley, "On the scope and interpretation of the tofts models for dce-mri," *Magnetic resonance in medicine*, vol. 66, no. 3, pp. 735–745, 2011.

[64] ——, "Classic models for dynamic contrast-enhanced mri," *NMR in Biomedicine*, vol. 26, no. 8, pp. 1004–1027, 2013.

[65] M. Wennen, W. Stehling, J. T. Marcus, J. P. Kuijer, C. Lavini, L. M. Heunks, G. J. Strijkers, B. F. Coolen, A. J. Nederveen, and O. J. Gurney-Champion, "A signal model for fat-suppressed t1-mapping and dynamic contrast-enhanced mri with interrupted spoiled gradient-echo readout," *NMR in Biomedicine*, vol. 38, no. 1, p. e5289, 2025.

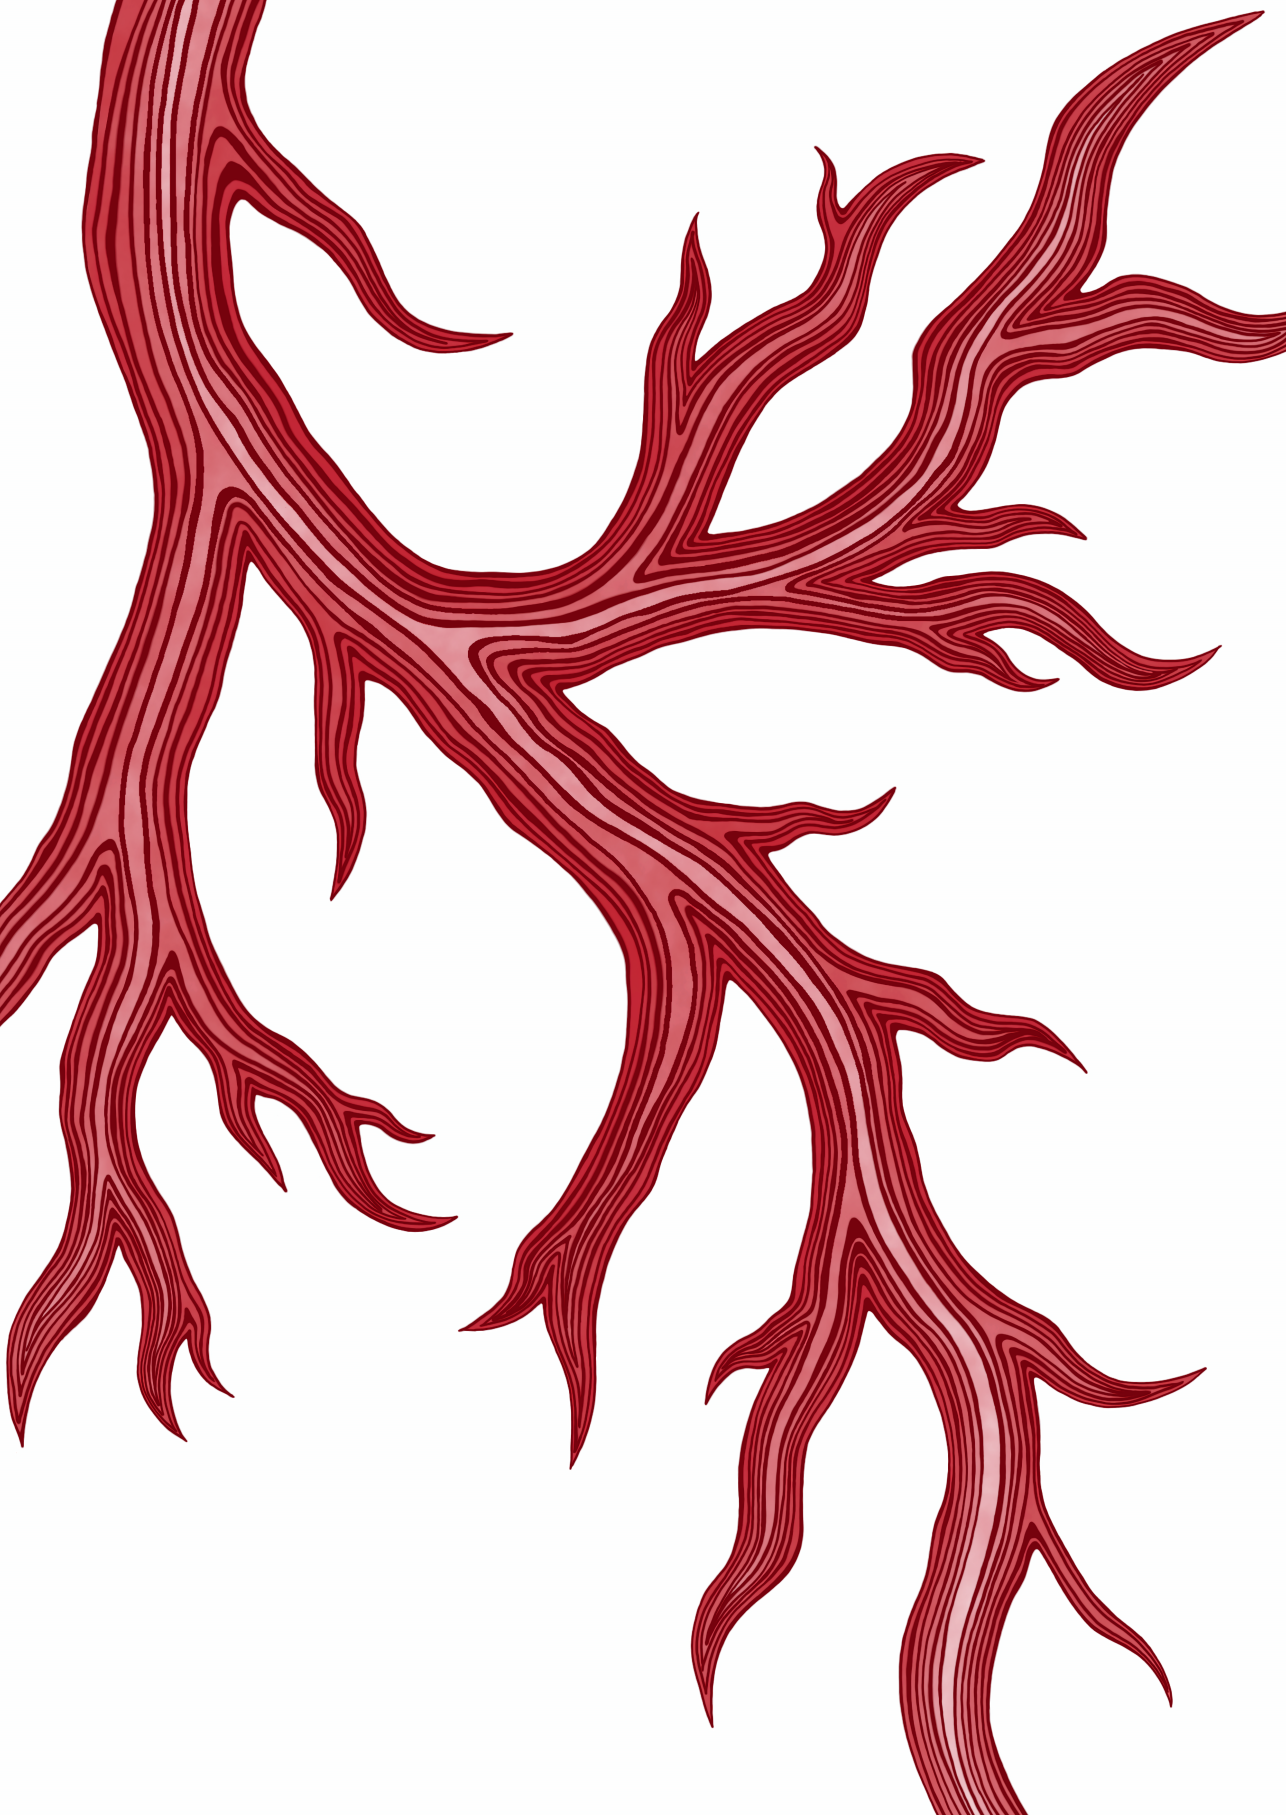

5

Uncertainty estimates in pharmacokinetic modelling of DCE-MRI

Authors: Jonas M. Van Elburg*, Natalia V. Korobova*, Mohammad M. Islam, Marian A. Troelstra, Oliver J. Gurney-Champion (*equal contribution)

Published in: Medical Image Analysis (Volume 109, December 2025)

https://doi.org/10.1016/j.media.2025.103881

Abstract

Dynamic contrast-enhanced (DCE) MRI is a powerful technique for detecting and characterising various diseases by quantifying tissue perfusion. However, accurate perfusion quantification remains challenging due to noisy data and the complexity of pharmacokinetic modelling. Conventional non-linear least squares (NLLS) fitting often yields noisy parameter maps. Although deep-learning algorithms generate smoother, more visually appealing maps, these may lure clinicians into a false sense of security when the maps are incorrect. Hence, reliable uncertainty estimation is crucial for assessing model performance and ensuring clinical confidence.

Therefore, we implemented an ensemble of mean-variance estimation (MVE) neural networks to quantify perfusion parameters alongside aleatoric (data-driven) and epistemic (model-driven) uncertainties in DCE-MRI. We compared MVE with NLLS and a physics-informed neural network (PINN), both of which used conventional covariance matrix-based uncertainty estimation.

Simulations demonstrated that MVE achieved the highest accuracy in perfusion and uncertainty estimates. MVE's aleatoric uncertainty strongly correlated with true errors, whereas NLLS and PINN tended to overestimate uncertainty. Epistemic uncertainty was significantly higher for the data deviating from what was encountered in training (out-of-distribution) in both MVE and PINN ensembles. In vivo, MVE produced smoother and more reliable uncertainty maps than NLLS and PINN, which exhibited outliers and overestimation. Within a liver region of interest, MVE's uncertainty estimates matched the standard deviation of the data more closely than NLLS and PINN, making it the most accurate method.

In conclusion, an MVE enhances quantitative DCE-MRI by providing robust uncertainty estimates alongside perfusion parameters. This approach improves the reliability of AI-driven MRI analysis, supporting clinical translation.

5.1 Introduction

Dynamic contrast-enhanced (DCE) MRI is widely used in clinical research for the detection, characterisation, and monitoring of various diseases, including oncological, renal, neurovascular and neurodegenerative diseases [1–4]. DCE-MRI involves intravenous injection of a gadolinium-based tracer, or contrast agent, whilst acquiring a time series of T1-weighted images to study the contrast dynamics. As the contrast agent perfuses through the tissues, it causes a local decrease in the T1 relaxation time of the water molecules, increasing the signal intensity as function of contrast concentration. By converting the signal as a function of time into contrast concentration curves over time and fitting a pharmacokinetic model to this curve, tissue microvascular perfusion and vascular permeability can be quantified [5].

Estimating pharmacokinetic parameters from the concentration curve presents a complex inverse problem that can result in inaccurate estimates due to several contributions. First, due to the need for fast acquisitions to capture the rapid dynamics of the contrast agent, DCE-MRI acquisitions are typically characterized by low signal-to-noise ratios (SNR). Second, pharmacokinetic models, such as the extended Tofts-Kety model, can be hard to fit, due to a strong cross-correlation between pharmacokinetic parameters, for example, fractional volume v_e and reflux rate k_e. Finally, with noisy data and complex models, choosing an appropriate fitting algorithm becomes paramount. Conventional algorithms, such as non-linear least squares (NLLS) fit, are sensitive to noise and local minima, causing poor repeatability of parameter estimates [6]. As NLLS is iterative, it is slow and fitting all voxels in an image can take minutes to hours, depending on implementation.

Deep learning neural networks have emerged as an alternative to NLLS fitting for estimating pharmacokinetic parameters [7–12]. Unlike NLLS, these networks can learn prior information about typical pharmacokinetic parameters during training, making them more robust to outliers. Additionally, they do not rely on iterative optimization during inference, allowing faster computation and reducing the risk of ending in local minima. As a result, they are less affected by noise and have demonstrated superior precision [7, 12], accuracy [7, 8, 12], and fitting speed [7, 8, 11, 12].

However, it turns out that neural networks can produce visually appealing parameter maps in regions where they are wrong/uncertain and even miss important pathologies [13]. This happens especially when applied to noisy or out-of-distribution (OOD) data [14] (OOD refers data distributions that deviate from those encountered during training). This is problematic because it masks uncertainty, making it difficult for radiologists or clini-

cians to distinguish between reliable and unreliable predictions. As a result, there is a risk of overtrusting the output, potentially leading to misdiagnosis, delayed treatment, or false reassurance. In clinical practice, where decisions are made based on the apparent quality of parameter maps, such "false clarity" can be more dangerous than obvious failure.

To ensure confidence in model predictions and prevent clinicians from being misled, it is essential to quantify the uncertainties associated with predicted values. Providing uncertainty estimates helps identify which regions can be reliably assessed, providing a measure of certainty in diagnosis. Uncertainty can be categorized into aleatoric and epistemic uncertainty. Aleatoric uncertainty arises from inherent noise in the data, while epistemic uncertainty arises from a lack of knowledge or model limitations [15]. Distinguishing between these two types allows us to determine whether uncertainty is an intrinsic property of the data or a result of model shortcomings, e.g. when encountering OOD data.

Conventional NLLS methods often incorporate post hoc aleatoric uncertainty quantification, based on covariance matrix estimation [16]. More recently, deep learning capable of estimating aleatoric and / or epistemic uncertainty have been developed for various quantitative MRI applications, including liver fat quantification [17], chemical exchange transfer rate mapping in the brain [18], diffusion-weighted imaging in the brain [19] and in the prenatal placenta MRI [20], and DCE-MRI in the brain [8] and in the liver [9]. Most of these methods [8, 9, 19, 20] adopt a Bayesian approach where aleatoric and / or epistemic uncertainties are expressed as an approximate posterior distribution over model parameters. While conceptually appealing, Bayesian inference requires explicit specification of prior distributions and approximation strategies to handle the otherwise intractable true posterior. By contrast, deep ensembles offer a simpler and more practical alternative for estimating epistemic uncertainty. They approximate posterior variability by exploiting randomness in training, primarily from weight initialization, to generate diverse models whose predictions capture epistemic uncertainty without requiring explicit priors.

In this study, we establish a different framework for quantifying both aleatoric and epistemic uncertainties in quantitative DCE-MRI. Specifically, aleatoric uncertainty was modelled with a mean-variance estimation (MVE) neural network, which avoids explicit prior assumptions and learns uncertainty directly from data. Epistemic uncertainties were captured with deep ensembles, which allow detection of OOD data. These methods were compared against a post hoc uncertainty quantification approach applied to pharmacokinetic parameters calculated by conventional NLLS fitting and DCE-NET, a deep-learning neural network specifically developed for DCE-MRI data. Additionally, we evaluated our framework against a Bayesian inference method, μGUIDE [19], originally developed for diffusion-weighted MRI and adapted here for DCE-MRI. To guide our study, we define three hypothe-

ses:

H1 The MVE network provides more accurate and reliable aleatoric uncertainty estimates compared to a post hoc approach applied to NLLS and DCE-NET predictions, while maintaining parameter accuracy.

H2 Deep ensembles capture epistemic uncertainty, enabling the identification of OOD data in DCE-MRI.

H3 Application of MVE networks and deep ensembles to in-vivo DCE-MRI data is feasible and informative about aleatoric and epistemic uncertainties.

5.2 Theory

This chapter provides the theoretical background of deep learning for quantitative DCE-MRI, focusing on the extended Tofts-Kety model [21, 22], a widely used pharmacokinetic model for assessing tumour microvascularisation [23].

5.2.1 Extended Tofts-Kety model

The extended Toft-Kety model assumes the diffusion of the injected contrast agent between two compartments: blood plasma and extracellular-extravascular space (EES). The concentration of contrast agent in tissue C_t at a time t is modelled using pharmacokinetic parameters $\boldsymbol{u} = \{k_e, v_e, v_p\}$:

$$C_t(t, \boldsymbol{u}, C_p(t)) = v_p C_p(t) + v_e k_e \int_0^t C_p(\lambda) e^{-k_e(t-\lambda)} d\lambda, \tag{5.1}$$

where v_p $[fraction]$ is the volume of the plasma, v_e $[fraction]$ is the volume of the EES, k_e $[min^{-1}]$ is the contrast agent reflux rate, and $C_p(t)$ $[mM]$ is the arterial input function (AIF), representing the contrast concentration in blood plasma.

An extended implementation of the model, as proposed in [24], includes an additional parameter dt, resulting in the parameter set $\boldsymbol{u} = \{k_e, v_e, v_p, dt\}$. The parameter dt models the delay in bolus arrival between the feeding artery and the tissue. The modified equation becomes:

$$C_t(t, \boldsymbol{u}, C_p(t)) = v_p C_p(t) + v_e k_e \int_0^t C_p(\lambda) e^{-k_e(t-(\lambda+dt))} d\lambda, \tag{5.2}$$

.

5.2.2 Uncertainty in non-linear least squares fit

NLLS fit is commonly used to estimate pharmacokinetic parameters $\hat{\boldsymbol{u}}$ by minimizing the mean squared error between the observed concentration curve $\tilde{C}_t$ and the fitted concentration curve $\hat{C}_t := C_t(\hat{\boldsymbol{u}})$ [25–29].

To quantify uncertainty, NLLS methods often rely on post hoc covariance matrix estimation [16]. This involves computing the Hessian matrix, which contains second-order partial derivatives of the objective function with respect to the fitted parameters. At the optimal parameter values, the Hessian describes how tightly the model is constrained around the fit, directly relating to parameter uncertainty. The inverse of the Hessian matrix, when appropriately scaled, forms the covariance matrix, from which uncertainty is typically expressed as the variance of fitted parameters error (its diagonal elements). While mathematically accurate, in practice, this approach can yield inaccurate uncertainty estimates when the Hessian inverse is poorly conditioned. Additionally, when the solution converged to a local minimum, the Hessian (and its inverse) will be misleading.

5.2.3 Deep learning for parameter and aleatoric uncertainty estimation

Neural networks have been explored as an alternative to NLLS for estimating DCE-MRI parameters [7–12]. Compared to NLLS, neural networks offer greater robustness to outliers and faster inference times. These models are typically trained using either a supervised or physics-informed loss function.

Supervised loss operates in the parameter space and requires ground truth reference values for training:

$$\mathcal{L}_{\text{supervised}} = (\boldsymbol{u} - \hat{\boldsymbol{u}})^2, \tag{5.3}$$

where $\boldsymbol{u}$ represents ground-truth parameters and $\hat{\boldsymbol{u}}$ denotes predicted parameters. In contrast, physics-informed loss operates in the concentration curve space and does not require ground truth values:

$$\mathcal{L}_{\text{physics-informed}} = (\tilde{C}_t - \hat{C}_t)^2, \tag{5.4}$$

where $\tilde{C}_t$ is the measured concentration curve, and $\hat{C}_t$ is estimated concentration curve derived from predicted parameters $\hat{\boldsymbol{u}}$.

Supervised training is a more direct training form, but may not guarantee good performance in vivo due to the data shift between simulated and in vivo data (OOD). Physics-informed training allows overcoming this limitation by enabling training in vivo, but could

run in to similar errors as NLLS as they optimize the same objective function.

To quantify aleatoric uncertainty, neural networks can use mean-variance estimation (MVE) [30, 31]. Instead of predicting a single parameter value, such an uncertainty-aware network predicts a Gaussian distribution characterised by a mean $\hat{\boldsymbol{u}}_{\boldsymbol{\mu_i}}$ and variance $\hat{\boldsymbol{u}}_{\boldsymbol{\sigma_i^2}}$. Such a network can be trained using the Gaussian negative log-likelihood loss (GNLL):

$$\mathcal{L}_{\text{GNLL}} = \frac{1}{2}\left(\frac{(\boldsymbol{u} - \hat{\boldsymbol{u}_\mu})^2}{\hat{\boldsymbol{u}}_{\sigma^2}} + \log \hat{\boldsymbol{u}}_{\sigma^2}\right). \tag{5.5}$$

5.2.4 Epistemic uncertainty quantification

While MVE effectively estimates aleatoric uncertainty, it does not account for epistemic uncertainty. In DCE-MRI, epistemic uncertainty arises from dataset imbalances, over-/under-fitting, or OOD data.

Deep ensembles [32] address epistemic uncertainty by analysing the spread of predictions from N independently trained models. In case aleatoric uncertainty is already estimated by the model, each model provides predictions $(\hat{\boldsymbol{u}}_{\mu_i}, \hat{\boldsymbol{u}}_{\sigma_i^2})$, which are treated as samples from an approximate posterior distribution $p(\hat{\boldsymbol{u}}|\tilde{C}_t, \mathcal{D})$, where $\tilde{C}_t$ is the input concentration curve and $\mathcal{D}$ is the training dataset. The ensemble mean $\hat{\boldsymbol{u}}^*_\mu$ and variance $\hat{\boldsymbol{u}}^*_{\sigma^2}$ are computed as:

$$\hat{\boldsymbol{u}}^*_\mu = \frac{1}{N}\sum_i^N \hat{\boldsymbol{u}}_{\mu_i}, \tag{5.6}$$

$$\hat{\boldsymbol{u}}^*_{\sigma^2} = \underbrace{\frac{1}{N}\sum_{i=1}^N \hat{\boldsymbol{u}}_{\sigma_i^2}}_{\text{Aleatoric}} + \underbrace{\frac{1}{N}\sum_{i=1}^N \hat{\boldsymbol{u}}^2_{\mu_i} - \left(\hat{\boldsymbol{u}}^*_\mu\right)^2}_{\text{Epistemic}}. \tag{5.7}$$

This formulation separates aleatoric and epistemic uncertainty, enabling independent evaluation of each component. Deep ensembles can be applied to any model that estimates uncertainty as a Gaussian variance, making them a versatile approach for robust uncertainty quantification.

5.3 Methods

All code is shared on `https://github.com/JonasElburgUVA/uncertainty-in-quantitative-dce-mri-public`. All analyses were performed using Python (v3.12.5). All neural networks were trained to map DCE-MRI contrast concentration curves C_t

to extended Tofts-Kety pharmacokinetic parameters $\boldsymbol{u}$. We used a population-based AIF for all concentration curves [33] and the extended Tofts-Kety model implementation from Open Science Initiative for Perfusion Imaging (OSIPI) [34], specifically from OG_MO_AUMC_ICR_RMH_NL_UK directory [24]. This implementation expresses the AIF in an analytical form (5.B), allowing for an exact analytical solution to the convolution in the Tofts-Kety equation. Additionally, this implementation of extended Tofts-Kety model introduces an extra parameter dt that models the delay in the bolus arrival between the feeding artery and the tissue.

5.3.1 Data

5.3.1.1 Synthetic datasets

We simulated 10,000,000 contrast agent concentration curves using the extended Tofts-Kety model (Equation 5.1) by simulating realistic pharmacokinetic parameters first. To obtain a realistic parameter distribution, we trained and applied a physics-informed DCE-NET on the training portion of an in vivo abdominal data set (Section 5.3.1.2.1) and used its predictions on the same subset to fit a multivariate Gaussian distribution (Appendix 5.E.1). Pharmacokinetic parameter were then sampled from this distribution, with any values falling outside physiologically relevant ranges (Table 5.1) excluded. Parameters v_e and v_p were restricted between 0 and 1, as they represent fractional volumes. The upper bound for k_e was set to reflect values observed in the in vivo training data, while the overall parameter ranges were chosen to capture values reported in the literature for typical abdominal tissues [35]. Importantly, the validation and test subsets of the in vivo data were not used to construct simulations and were reserved exclusively for evaluation.

Gaussian noise was added to each concentration curve, with a standard deviation randomly selected from the range $[0.01 - 0.2\,mM]$ with increments of 0.01. The synthetic data was divided into training (70%), validation (15%) and test (15%) sets.

Parameter	Range
$k_e, [min^{-1}]$	[0,3]
$v_e, [fraction]$	[0,1]
$v_p, [fraction]$	[0,1]
$dt, [min]$	[0,2]

Table 5.1: Physiologically relevant ranges of the pharmacokinetic parameters.

To simulate a data shift between training and test datasets, we excluded a part of the training set, creating out-of-distribution (OOD) test values (Appendix 5.E.1). For k_e, dt, and

v_p, all data points exceeding two standard deviations above the mean were removed from the training data. For v_e the cut-off was set to one standard deviation above the mean due to its broader distribution. Test values that remained represented in the training set are referred to as in-distribution (ID). The validation set followed the same distribution as the test set.

5.3.1.2 In vivo data

5.3.1.2.1 Abdominal DCE-MRI in healthy volunteers

Eleven healthy volunteers underwent a DCE-MRI scan of the liver. The study was approved by the ethics committee of Amsterdam UMC location AMC (Amsterdam, The Netherlands), and all volunteers provided written informed consent (NL54274.018.15) before the study. A detailed description of the acquisition protocol is provided in 5.C.

Before performing experiments, the in vivo data were preprocessed in several steps. First, to exclude background and weakly perfused voxels from the dataset, we created a foreground mask containing all voxels with an average temporal signal above the third quartile of all non-zero voxels. A morphological opening operator with a 3x3 structuring element was then applied to remove isolated noisy voxels. Next, the signals of the masked voxels were converted to contrast concentration using steps described in 5.A. We assumed a constant baseline T_{10} value of 812 ms, corresponding to the liver at 3T [36]. Finally, the processed scans were split into a training (N_{scans}=6, $N_{\text{voxels}} = 4,101,132$), validation ($N_{\text{scans}}$=3, $N_{\text{voxels}} = 2,056,236$) and test ($N_{\text{scans}}$=2, $N_{\text{voxels}} = 1,373,859$) data sets.

5.3.1.2.2 Brain DCE-MRI in tumor patients

We used a publicly available dataset from the OSIPI DCE Challenge as an OOD in vivo dataset. This dataset included DCE-MRI scans of eight patients with glioblastoma acquired on a 1.5 T MRI system, with lesion segmentations provided for each scan. In addition to the DCE scans, a variable flip-angle GRE sequence was acquired for T1 mapping, which was performed using NLLS fitting. As with the abdominal dataset, DCE signal intensities were converted to concentration curves following the steps described in 5.A. A detailed description of acquisition protocols is provided in [37].

5.3.2 Pharmacokinetic parameter and uncertainty estimation methods

5.3.2.1 Neural networks

We used three architectures of neural networks: a physics-informed neural network (PINN), a supervised MVE model, and a Bayesian inference method μGUIDE [19] (Figure 5.1). PINN and MVE networks were built upon the DCE-NET architecture [7], available on GitHub - oliverchampion/DCENET.

DCE-NET predicts four target parameters: three extended Tofts-Kety pharmacokinetic parameters $\hat{\boldsymbol{u}}$ and a timing offset parameter dt. We used the Gated Recurrent Unit (GRU)-based version of DCE-NET, which employs four stacked bidirectional GRU layers. For each predicted parameter, the model applies an attention mechanism over the GRU outputs to compute an attention-weighted representation. Each of these representations is then passed through two fully connected layers with ELU activation to produce the final prediction. To ensure predictions remained within physiologically plausible ranges, a scaled sigmoid function was applied to the estimates, soft-bounding them within $\pm 25\%$ of the total parameter range defined in Table 5.1 (Equation 5.8):

$$S(x, \boldsymbol{b}) = b_{\text{lower}} + \frac{(b_{\text{upper}} - b_{\text{lower}})\,(5 - e^{-x})}{4\,(1 + e^{-x})}, \tag{5.8}$$

where $\boldsymbol{b} = [b_{\text{lower}}, b_{\text{upper}}]$ represents the lower and upper boundaries of parameters u as given in Table 5.1.

The training objective of the PINN was a physics-informed loss (Equation 5.4). As there is no MVE equivalent for the physics-informed loss, the PINN employed a post hoc uncertainty quantification method similar to that used by `scipy.curve_fit` to evaluate uncertainty (red block on Figure 5.1). However, instead of using numerical approximations to compute the Jacobian, PyTorch's automatic differentiation was utilized, providing precise derivatives while keeping all other aspects of the uncertainty estimation procedure unchanged. To address potential issues in taking the inverse of an ill-conditioned Hessian matrix, we replaced it with the pseudo-inverse approximation.

An uncertainty-aware MVE network was built upon DCE-NET by extending the output layer from four to eight nodes enabling the prediction of log-variances alongside point estimates (green block on Figure 5.1). The variance was retrieved by taking the exponential of the added outputs. The MVE model was trained with the GNLL learning objective (Equation 5.5).

The Bayesian inference network μGUIDE was implemented with its default hyperparameters, using a three-layer multi-layer perceptron combined with a neural posterior estimator to estimate non-parametric posterior distributions [19] (blue block on Figure 5.1).

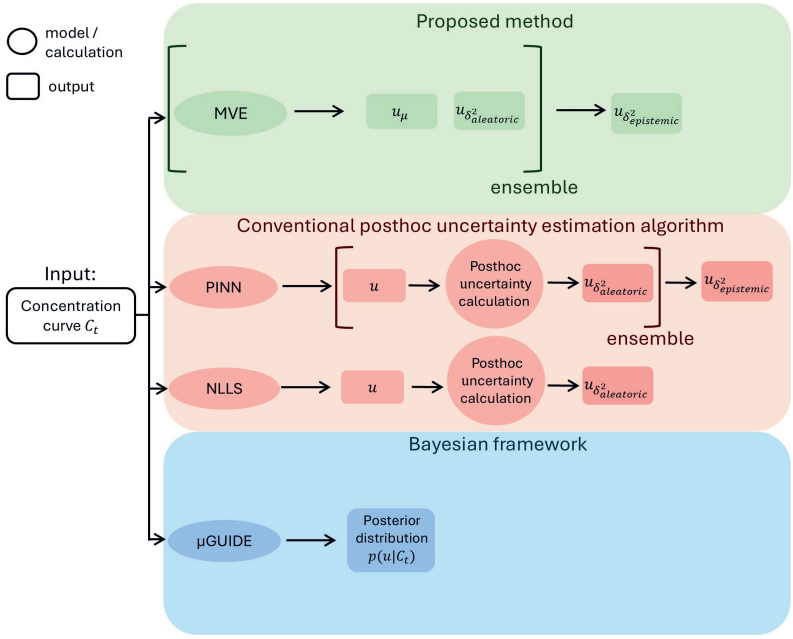

Figure 5.1: Graphical summary of implemented methods. Circles represent models or algorithms, while rectangles with rounded corners denote inputs and outputs. Green block: proposed framework, where an MVE network estimates pharmacokinetic parameters along with aleatoric uncertainty. A deep ensemble of MVEs additionally provides epistemic uncertainty. Red block: conventional approaches using uncertainty estimation based on the covariance matrix. Here, NLLS and PINN estimate parameters, with post hoc uncertainty quantification providing aleatoric uncertainty. PINNs were further extended to deep ensembles to capture epistemic uncertainty. Blue block: an alternative Bayesian inference method (μGUIDE), which estimates posterior distribution of parameters.

In adapting μGUIDE from diffusion MRI to DCE-MRI, we omitted the diffusion-specific signal normalization step, since converting the MR signal to contrast-agent concentration inherently compensates for baseline effects.

Training hyper-parameters of all networks are provided in Table 5.2.

5.3.2.2 Non-linear least squares fit

For comparison, a trust-region reflective NLLS algorithm with post hoc uncertainty quantification (`scipy.curve_fit`) was evaluated on the simulated test dataset. The bounds for the pharmacokinetic parameters were set to the physiologically relevant ranges (Table 5.1). The initial guess was set to the mean of the simulated distribution described in Section 5.3.1.1 ($k_e = 0.87\ min^{-1}, v_e = 0.31, v_p = 0.02, dt = 0.79\ min$). In line with the post-hoc method applied to the PINN model, `scipy` uses the pseudo-inverse of the Hessian to prevent computational issues (red block on Figure 5.1).

Hyper-parameter	DCE-NET (PINN and MVE)	μGUIDE
Max number of epochs	200	500
Number of burn-in epochs for MVE models	10	-
Early stopping patience	20	10
Optimizer	Adam	Adam
Learning rate	1e-4	1e-3
Learning rate scheduler patience	10	-
Learning rate scheduler factor	0.1	-
Samples drawn from posterior	-	50000

Table 5.2: Model-training hyper-parameters. Burn-in epochs refer to the number of epochs trained with a deterministic objective (Equations 5.3,5.4) before switching to the GNLL objective (Equations 5.5, 5.15) in MVE models. The learning rate scheduler multiplied the learning rate by 0.1 after 10 epochs of non-improvement in the validation loss.

5.3.3 Experiments

This section describes the experiments conducted to address the study's hypotheses (H1-H3) by evaluating methods on synthetic and in vivo data.

5.3.3.1 Experiment A: In distribution accuracy and aleatoric uncertainty

To test our first hypothesis (H1), we evaluated the uncertainty estimates of an MVE neural network and compared them to post hoc approaches applied to NLLS and PINN and to Bayesian inference method μGUIDE. We also assessed the parameter accuracy of all four methods. This was done in a controlled environment using the synthetic dataset.

Parameter accuracy was evaluated using the mean absolute error (MAE) between predicted $\hat{\boldsymbol{u}}_{\boldsymbol{\mu}}$ values and ground-truth $\boldsymbol{u}$ values, and analysed across the varying noise levels. For μGUIDE, the estimated parameter value was defined as maximum a posteriori (MAP). Additionally, we calculated the MAE between the estimated $\hat{C}_t$ and the observed $\tilde{C}_t$ concentration curves for all four methods, which we refer to as the reconstruction error.

Numerical instabilities in the inversion of the Hessian matrix resulted in extremely high predicted uncertainties for some voxels. To mitigate this, we defined outliers as predictions with uncertainty values exceeding the physiologically reasonable boundaries specified in Table 5.1 and were removed from the analysis. These outliers were excluded from further analysis, and the percentage of excluded outliers was reported.

For aleatoric uncertainties estimated by MVE and by the conventional posthoc covariance-based method, the predicted variance $\hat{\boldsymbol{u}}_{\boldsymbol{\sigma^2}}$ was converted into predicted absolute error by expressing it as the expected absolute deviation of a parameter from its

mean value:

$$\mathbb{E}[|X - \hat{\boldsymbol{u}}_{\boldsymbol{\mu}}|] = \sqrt{\frac{2\,\hat{\boldsymbol{u}}_{\boldsymbol{\sigma^2}}}{\pi}}. \tag{5.9}$$

For μGUIDE, which produces complex distributions, we adopted an alternative strategy. For each input, 50.000 samples were drawn from the posterior and the predicted absolute error was computed as the mean absolute difference between the MAP estimate and the posterior samples.

Finally, to evaluate the calibration of uncertainty estimates, we examined the correspondence between predicted absolute error and the observed MAE, and visualized their correlation.

After removing outliers, the correlation between mean absolute error and mean predicted absolute error in the simulated test set was analysed by plotting MAE as a function of mean predicted absolute error sorted in ten equal-frequency bins.

5.3.3.2 Experiment B: Epistemic uncertainty

To test our second hypothesis (H2), we investigated whether deep ensembles could capture epistemic uncertainty to identify OOD data. For this, we constructed two ensembles: one containing five PINN models and another containing five MVE models. Both were trained on synthetic training data with limited representation of parameter values and evaluated on OOD and ID test values (Appendix 5.E.1).

Because epistemic uncertainty predicted by ensembles is often miscalibrated, we applied a variance temperature scaling approach, following [38, 39]. This ensures that the predicted variance more closely reflects the observed error distribution, improving comparability between models.

For comparison, we also trained μGUIDE on the same dataset. While in the original publication, authors did not explicitly categorize the type of uncertainty estimated, μGUIDE produces posterior distributions over parameters. In our OOD setting, we interpret the spread of these posteriors as epistemic uncertainty, since broad posteriors reflect lack of knowledge or non-identifiability rather than irreducible data noise. To quantify this, we measured posterior dispersion using the interquartile range (IQR), normalized by the prior range for each parameter. This normalization ensured that uncertainty values remained comparable across parameters with different scales. We expected OOD samples, not represented in training, to yield broad posteriors resembling the uninformative uniform prior and thus large normalized IQR values. Conversely, ID samples should produce more concentrated posteriors and smaller IQRs.

The difference in epistemic uncertainty predictions between the OOD and ID values was

assessed using a one-sided Wilcoxon sum-rank test (significance level $\alpha = 0.05$). Plots of true MAE and predicted epistemic uncertainty as a function of parameter value were generated.

5.3.3.3 Experiment C: Feasibility in vivo

Finally, we tested our third hypothesis (H3). We explored the feasibility of applying the networks to in vivo DCE-MRI data and addressed the informativeness of the predicted aleatoric and epistemic uncertainties. The PINN models were retrained on the abdominal in vivo data 5.3.1.2.1. Due to the lack of ground-truth pharmacokinetic parameters in vivo, the MVE models from Experiments A and B, which were trained on simulated data, were evaluated on the in vivo test data without retraining. μGUIDE was not applied to in vivo data because it requires supervised training with ground-truth parameters and, as shown in Section 5.4.2, it exhibited overconfident uncertainty estimates on OOD data, making its use in this context inappropriate.

5.3.3.3.1 Aleatoric uncertainty

Aleatoric uncertainties were estimated on the abdominal in vivo dataset 5.3.1.2.1. The simulated dataset used to train the MVE models shared physiological similarities with this in vivo abdominal data, as the latter was used to derive the pharmacokinetic parameter distributions for simulation.

Uncertainty maps, expressed as predicted absolute error, were visually evaluated for NLLS, PINN, and MVE models. Similarly to Experiment A, data points with predicted uncertainty that exceeded physiologically relevant bounds were marked as outliers.

As we do not have ground truth in vivo, we cannot directly compare uncertainty with errors. Instead, we evaluated aleatoric uncertainties within a small region of interest (ROI) in the liver. Assuming that the tissue within the ROI is homogenous, any spread in parameter estimates was attributed primarily to aleatoric uncertainty. The standard deviation (SD) of parameter predictions within ROI was thus used as a reference for aleatoric uncertainty estimates. The predicted uncertainty $\hat{\boldsymbol{u}}_{\boldsymbol{\sigma}^2}$ was converted into standard deviations (SD), and compared to the mean SD of parameter predictions within the ROI.

5.3.3.3.2 Epistemic uncertainty

To assess epistemic uncertainty in an OOD setting, we evaluated a publicly available brain tumor dataset (Section 5.3.1.2.2) using a deep ensemble of MVE models. Predicted uncertainties in the liver (from the abdominal scans) were compared with those in the tumor

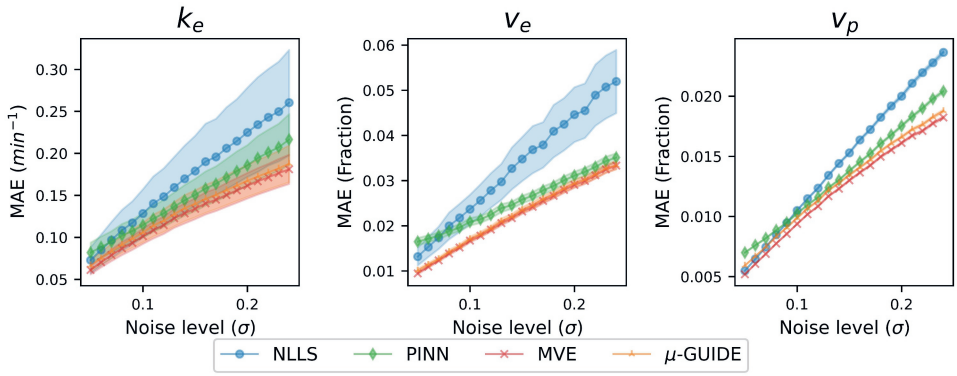

Figure 5.2: Accuracy (MAE) as a function of simulated noise level on each parameter (left: k_e, middle: v_e, right: v_p) over 20 equal-width bins. The error variance per model and bin is visualised around the average values of each bin. NLLS was the most affected by high noise levels, while MVE obtained the best accuracy for all noise levels and all parameters.

regions (from the brain scans). Tumor segmentations were provided in the dataset, while the full liver volumes were segmented semi-manually using nnInteractive [40], an interactive 3D segmentation tool based on a pretrained NN. Statistical differences between liver and tumor uncertainty distributions were assessed using a non-paired t-test for parameters k_e and v_e (normally distributed, Shapiro–Wilk $p > 0.05$), and a Mann–Whitney U test for v_p (non-normal distribution, Shapiro–Wilk $p < 0.05$), with significance defined as $p < 0.05$.

5.4 Results

5.4.1 Experiment A: In distribution accuracy and aleatoric uncertainty

MVE showed the highest parameter accuracy (lowest MAE), followed by μGUIDE, PINN and then NLLS. The performance gap between MVE/μGUIDE and NLLS/PINN was widening at higher noise levels (Figure 5.2) in k_e and v_p. In contrast, NLLS achieved the lowest reconstruction error across all noise levels, followed by PINN , μGUIDE, and then MVE (Table 5.3). These results highlight MVE as the most reliable method for accurate parameter estimation, while the lower parameter accuracy combined with lower reconstruction error of NLLS and PINN suggests a greater degree of overfitting to noise. The lower parameter accuracy of μGUIDE may be partly explained by its conversion of the full posterior distribution into a MAP point estimate, which may not optimally reflect the underlying distribution.

	NLLS	PINN	MVE	μGUIDE
Simulations (N=1,500,000)	**0.1126 (0.05)**	0.1138 (0.05)	0.1131 (0.05)	0.1135 (0.11)
In vivo (N=1,373,859)	**0.1903 (0.60)**	0.1926 (0.57)	0.2224 (0.62)	

Table 5.3: Mean absolute reconstruction error $[mM]$ of the models in the simulated and in vivo test-set. μGUIDE was not evaluated in vivo data because it exhibited overconfident uncertainty estimates on OOD data, making its use in this context inappropriate. Data are represented as mean (standard deviation). All differences were significant according to a Student's t-test with Bonferroni correction ($p<.016$)

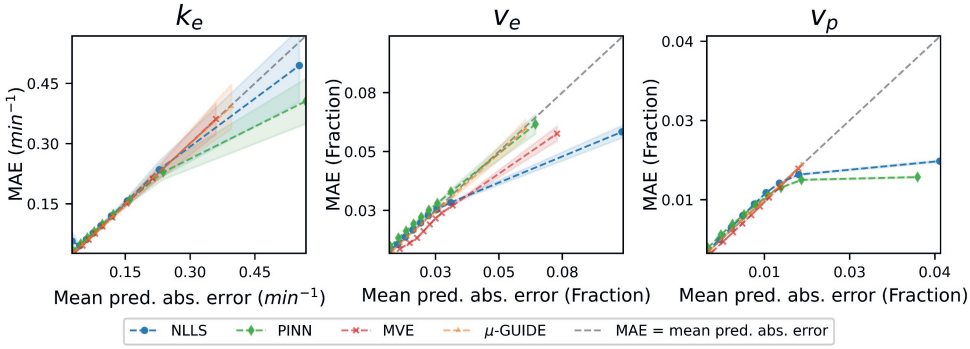

Figure 5.3: The relation between MAE and the mean predicted absolute error derived from uncertainty predictions (Eq. 5.9). Data points represent averages within ten equal-frequency bins, each containing 10% of the data, sorted by the predicted absolute error. The gray diagonal line shows perfect agreement between observed and predicted errors.

The MVE and μGUIDE did not produce any outliers in their predictions. NLLS with post hoc uncertainty quantification produced the highest percentage of outliers in uncertainty estimations: 2.17% (k_e), 7.76% (v_e), and 6.27% (v_p). Post hoc uncertainty quantification applied to a PINN reduced outliers to 1.85% (k_e), 0.05% (v_e), and 0.05% (v_p).

Both MVE and μGUIDE models showed a consistent one-to-one relationship between predicted and observed errors (Figure 5.3). A minor overestimation of uncertainty by MVE was observed for v_e, but overall, the MVE model demonstrated accurate error modelling across all parameters. NLLS and PINN models tended to overestimate uncertainties for larger errors. This effect was particularly pronounced in the k_e and v_p parameters.

Without outlier rejection, NLLS and PINN do a considerably worse job, as shown in Appendix 5.E.2.

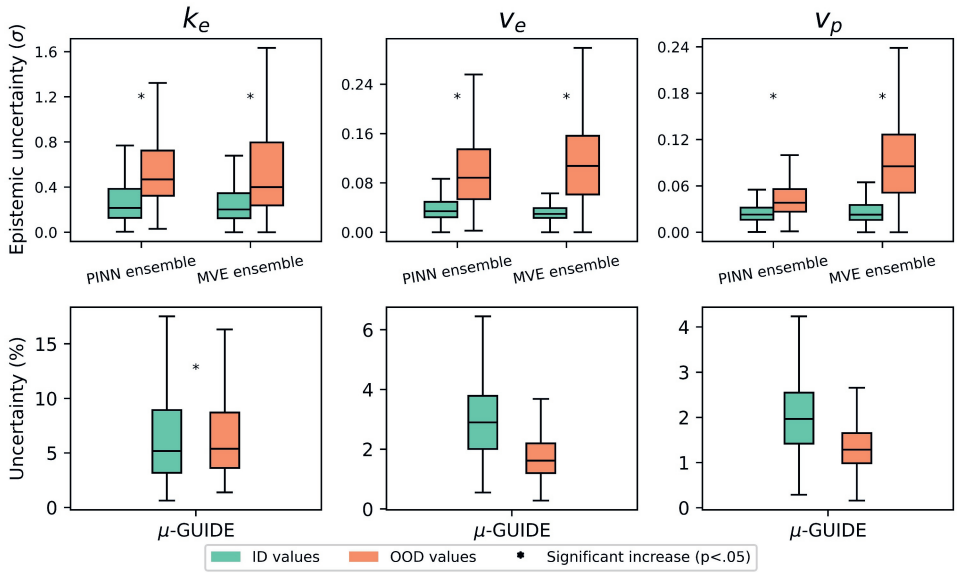

Figure 5.4: Distribution of epistemic uncertainty predicted by ensembles and μGUIDE for simulated ID and OOD test values. Epistemic estimates from ensambles were calculated using Equation 5.7. Epistemic uncertainty from μGUIDE was calculated as the posterior IQR normalized by the prior range. An asterisk (*) indicates significantly higher predicted uncertainty in OOD compared to ID data (one-sided Wilcoxon signed-rank test, p$<$.05). Boxplots show the IQR with a line at the median; whiskers extend to the most extreme data points within 1.5 × the IQR.

5.4.2 Experiment B - Epistemic uncertainty

Ensembles of MVE and PINN models produced significantly higher epistemic uncertainty in all parameters for OOD in comparison to ID data (Figure 5.4). When using PINN, the median uncertainty estimate in OOD data was 4.7, 6.8 and 2.8 times larger than the median estimate on ID data for k_e, v_e, and v_p respectively. Using MVE, the median OOD estimates were 3.9, 13.1 and 14.3 times higher. Bayesian inference method, μGUIDE, did not manage to distinguish between ID and OOD datasets across all parameters (Figure 5.4).

For both ensembles, we observed higher errors for OOD data compared to ID data, accompanied by a corresponding increase in epistemic uncertainty (Figure 5.5). This effect was particularly pronounced for the parameters v_e and v_p, where both the true MAE and predicted uncertainty rose sharply beyond the OOD cutoff. For k_e, the increase was more gradual across both methods. Although the curves of MAE and mean epistemic uncertainty did not perfectly overlap, they generally followed the same trends and fell within comparable ranges. In contrast, μGUIDE failed to detect OOD parameters reliably, often assigning lower uncertainty values to OOD samples than to ID ones. Moreover, since μGUIDE

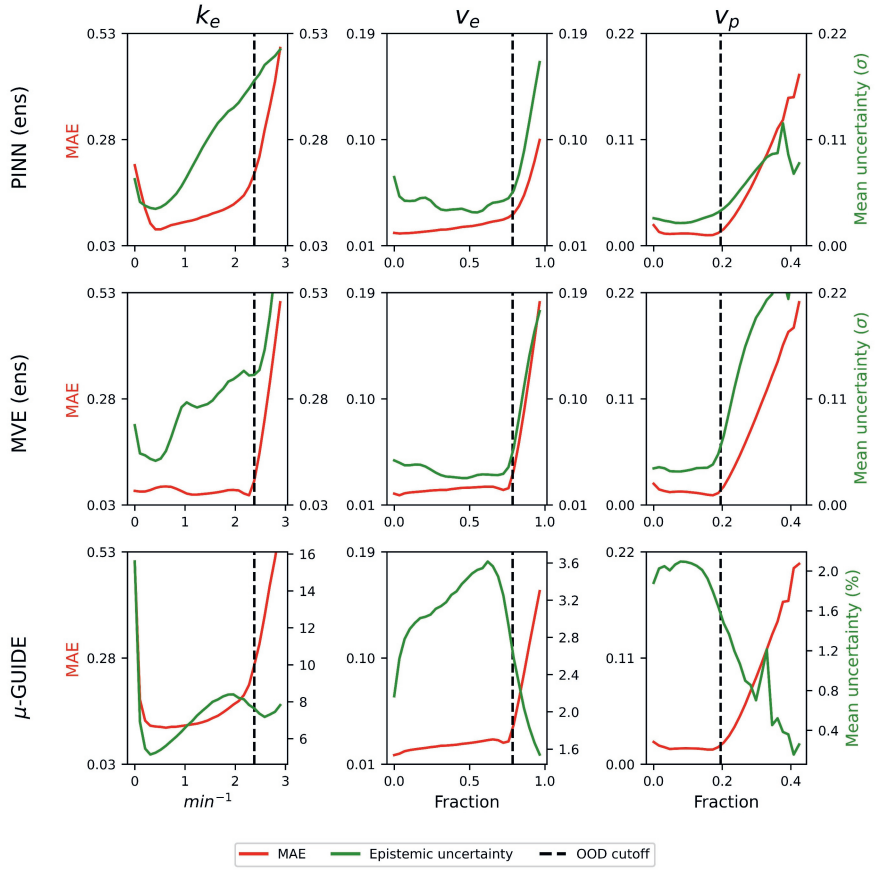

Figure 5.5: Average error (left y-axis) and average predicted epistemic uncertainty (right y-axis) across pharmacokinetic parameter values. Results are shown for ensembles of PINNs (top row), ensembles of MVEs (middle row), and μGUIDE (bottom row). Both ensembles report epistemic uncertainty as standard deviation, whereas μGUIDE reports it as a percentage. The black dashed line indicates the cutoff separating ID from OOD values. Error and uncertainty values were aggregated into 30 equally spaced parameter bins.

expresses uncertainty as a percentage rather than in absolute units, it limits the direct comparison between observed errors and predicted uncertainties.

5.4.3 Experiment C - Feasibility in vivo

Figure 5.6 shows parameter maps of the evaluated base models. NLLS, while providing a low mean reconstruction error (table 5.3), resulted in noisier k_e and v_e maps than deep models. The supervised MVE produced smoother parameter maps, but showed a more dimmed parameter map (avoiding extreme values), which is especially visible in the v_p parameter map. Presumably, the in vivo data noise is different from simulations (e.g. motion,

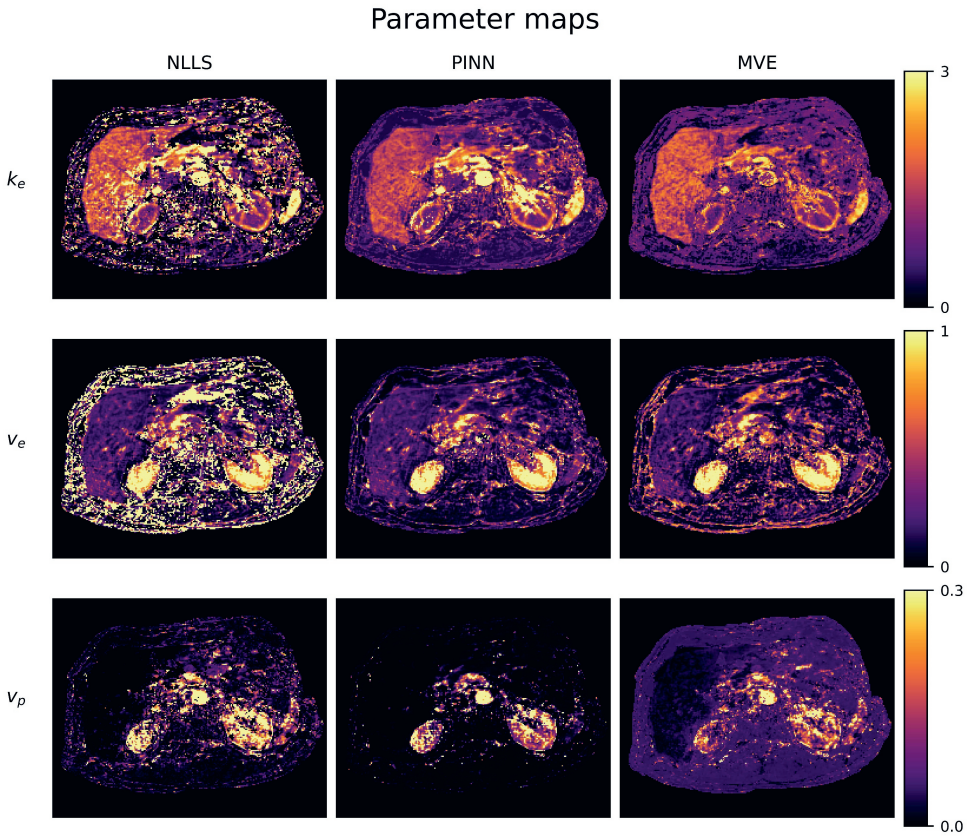

Figure 5.6: Spatial parameter maps of pharmacokinetic estimates derived from the assessed models. NLLS resulted in more erratic maps in contrast to neural networks. Supervised neural networks introduced a noticeable bias, which can be seen in its tendency to avoid extreme values.

artefacts), and this dimming is a bias towards the mean of the simulated training data. Furthermore, the NLLS and PINN models may be more affected by noise in the input data due to their physics-informed objective.

5.4.3.1 Aleatoric uncertainty

MVE-generated uncertainty maps were smoother than those from the post hoc methods, especially visible in the organs (e.g. liver and kidneys) in v_e and v_p (Figure 5.7), and remained within a range similar to the values found using simulations. As all networks work on a per-voxel base, the network likely recommends more consistent uncertainty estimates.

Post hoc aleatoric uncertainty estimates from NLLS showed numerous outliers

(28.81%), especially for v_e and v_p in the liver. Post hoc uncertainty estimates from PINN contained fewer outliers (3.37%), but k_e and v_p uncertainty estimates fell within the range where the post-hoc method tended to overestimate uncertainty in simulations (Figure 5.3), reducing their reliability. The v_e uncertainty map was noisy and inconsistent. Aleatoric uncertainty maps without marked outliers are shown in Appendix 5.E.3.

Predicted uncertainty maps of PINN and MVE differed notably for v_e. In noisy curves, PINN predicted higher uncertainty than MVE, likely reflecting its difficulty fitting noisy data, whereas MVE leveraged prior knowledge for more confident estimates (Appendix 5.E.4c). For less noisy curves, both models produced similar v_e fits and uncertainty estimates (Appendix 5.E.4d). In some cases, higher MVE uncertainty reflected failed fits, likely because the supervised network relied on priors from simulated data that could not capture atypical high curve peaks observed in vivo. While these failures were correctly flagged with high uncertainty, they also reveal a partial overlap between aleatoric and epistemic uncertainty.

ROI analysis revealed that MVE provided the most accurate uncertainty estimates, with a slight tendency toward underestimation, especially in v_e (Figure 5.8). As liver tissue is never completely homogeneous, some additional anatomically driven true variance in the reference value may be expected. MVE was the only model that produced consistent, smooth uncertainty maps and accurately estimated uncertainty in k_e and v_p.

5.4.3.2 Epistemic uncertainty

Overall, the ensembles of MVE models predicted significantly higher epistemic uncertainty in the OOD setting than in the ID (5.9). In particular, both the mean and median predicted uncertainties were significantly lower for v_e and k_e in the liver compared to the tumor, and their distributions exhibited a narrower spread. For v_p, predicted uncertainties were higher in the liver than in the tumor, although this difference was not statistically significant.

The mean reconstruction error of the MVE was larger than those obtained by NLLS and PINN, trained on in vivo data (Table 5.3). This discrepancy suggests that data shift (simulation training data vs in vivo test data) may have impacted parameter estimates. This effect may be overlooked by the aleatoric uncertainty produced by a single MVE network, showing the necessity of epistemic uncertainty estimates.

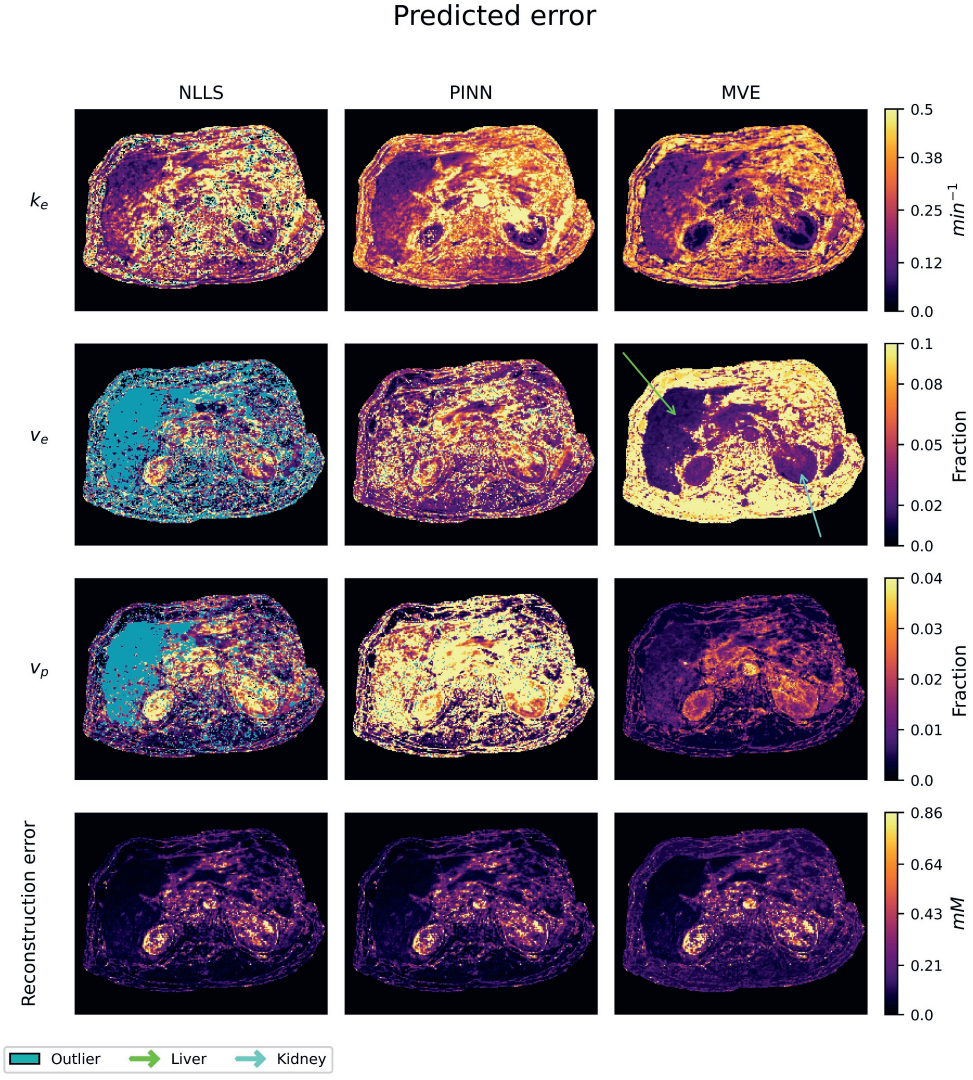

Figure 5.7: Predicted aleatoric uncertainty maps expressed as errors generated for pharmacokinetic parameters by single models (no ensembles). Outliers were marked in blue. To prevent pseudo-inverse instability-related overestimates from dominating the contrast between voxels, we set the maximum visualised value to the values where overestimation started dominating in simulations (Figure 5.3). NLLS displayed a large number of outliers in the v_e and v_p parameters. The distinct uncertainty patterns between parameters highlight differences in error attribution. The figure without marked outliers can be seen in Appendix 5.E.3.

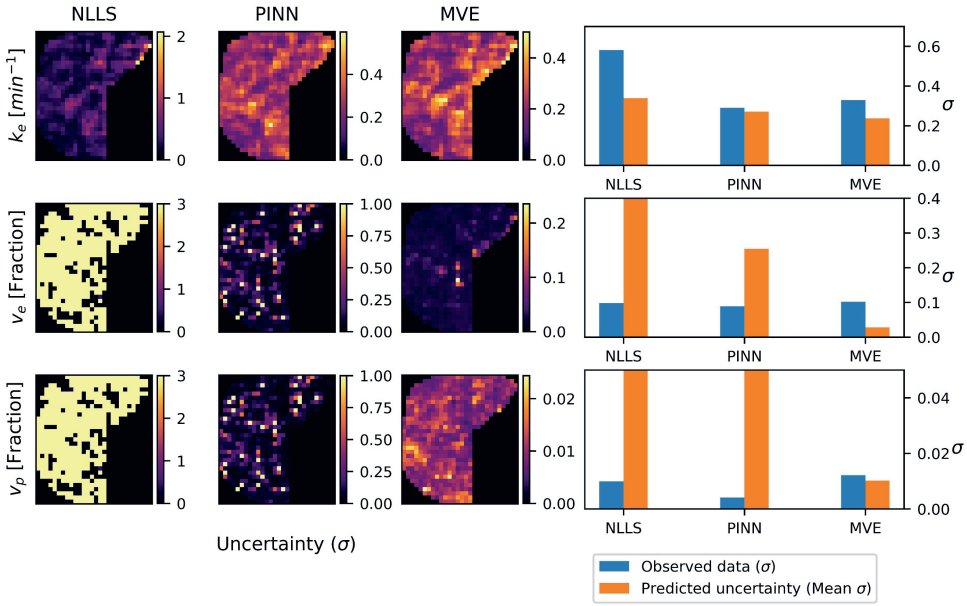

Figure 5.8: Aleatoric uncertainty maps of the liver ROI. On the right side, the standard deviation over parameter predictions in the ROI is compared to the predicted standard deviation based on the aleatoric uncertainty estimate. Note that colorbar ranges differ between parameters and methods due to large variations in values. A box-plot version is shown in Appendix 5.E.5.

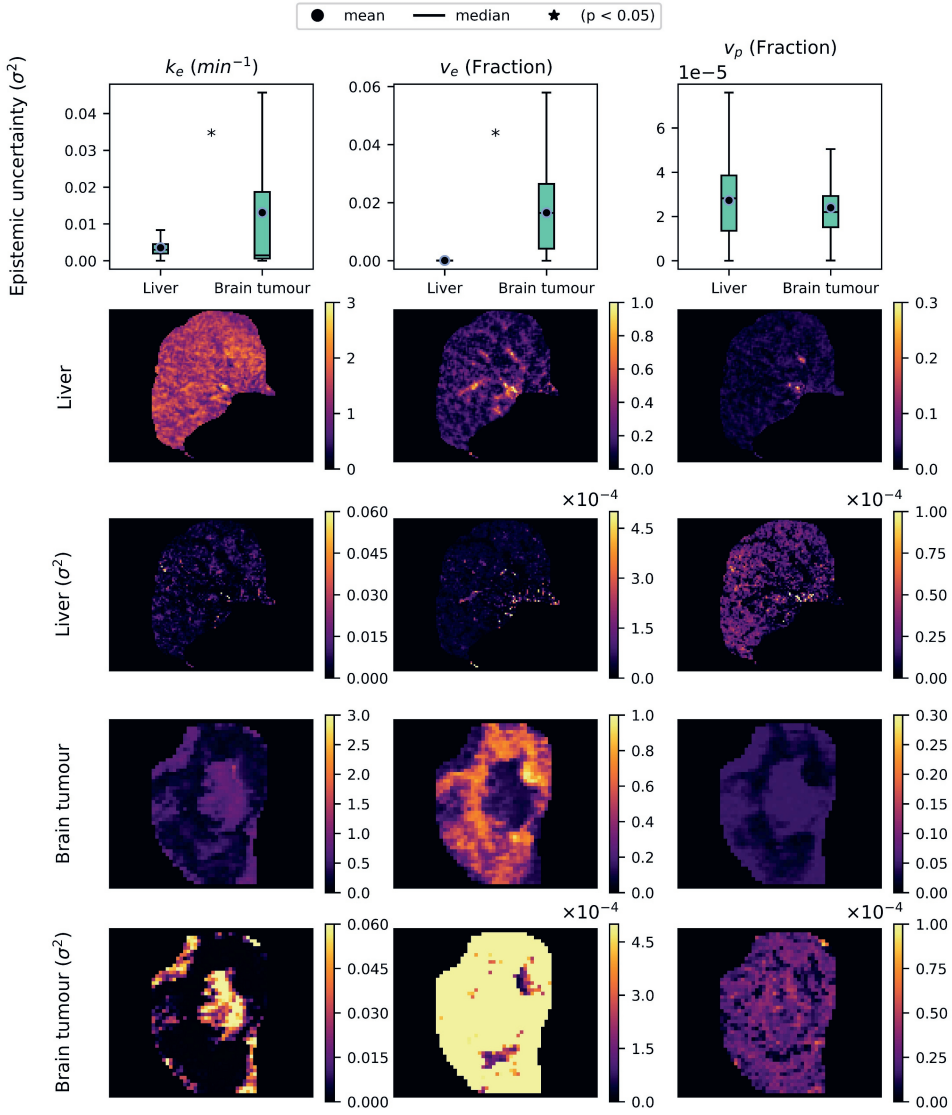

Figure 5.9: First row: Distributions of epistemic uncertainty estimated with an ensemble of MVE models in the liver (ID) and brain tumor (OOD). Second and fourth rows: pharmacokinetic parameter maps in the liver and brain tumor. Third and fifth rows: corresponding epistemic uncertainty maps for the predicted parameters.

5.5 Discussion

We implemented deep learning for aleatoric and epistemic uncertainty quantification in extended Tofts-Kety pharmacokinetic modelling using ensembles of uncertainty-aware MVE neural networks. We demonstrate that MVE networks provide more accurate and reliable aleatoric uncertainty estimates than post-hoc uncertainty estimates from PINNs and NLLS. By separating aleatoric from epistemic uncertainty, we can identify whether the uncertainties are inherent to the data, or whether the model needs further tuning, for example by showing it more data. Our results suggest that the MVE ensemble best predicts both aleatoric and epistemic uncertainty, however, translation to in vivo data is more challenging given its need for supervised training. These tools can prevent clinicians from being overconfident in smooth AI-generated parameter maps in the presence of uncertainty.

Our simulations demonstrated that incorporating uncertainty estimation did not compromise parameter accuracy - MVE produced the most accurate and precise parameter estimates compared to NLLS and PINN across various noise levels. In a separate comparison, MVE also outperformed a standard supervised neural network in estimating the v_e and v_p parameters, suggesting that incorporating uncertainty not only improves interpretability but may also enhance robustness and accuracy of the predictions (results not shown).

The aleatoric uncertainties predicted by MVE and a Bayesian inference method μGUIDE scaled linearly with noise level, whereas post hoc methods applied to NLLS and PINN tended to overestimate uncertainty at higher noise levels. These outliers likely occurred due to an ill-conditioned Hessian matrix, resulting in a numerically unstable calculation of its pseudo-inverse. However, even when the most extreme outliers were rejected (Figure 5.3), post hoc uncertainty estimates from NLLS and PINN do not perform as well as MVE. Hence, uncertainty-aware NNs such as MVE and μGUIDE seem the best at accurately estimating aleatoric uncertainty.

Further simulations using ID and OOD datasets confirmed that deep ensembles can effectively capture epistemic uncertainty. Both PINN- and MVE-based ensembles exhibited statistically significant higher uncertainty predictions in OOD datasets compared to ID datasets, with MVE ensembles having the strongest correlation to actual MAE. By contrast, μGUIDE often provided lower uncertainty values for OOD than ID, which limited its utility for reliable OOD detection in our experiments.

Taken together, our results suggest that while μGUIDE provides a principled Bayesian alternative, ensembles of MVEs offer a more accurate and practically flexible approach to separating aleatoric and epistemic uncertainty in quantitative DCE-MRI as they do not require explicit posterior modelling. Although our analysis of uncertainty relied on MAE,

which implicitly assumes Gaussian error distributions, this simplification did not reduce the accuracy of parameter or uncertainty estimates in practice, even compared to μGUIDE, which can model more complex posterior distributions. Moreover, the simpler architecture of MVE enables faster inference than Bayesian methods: for a single concentration curve with 80 samples, MVE was 4.6 times faster than μGUIDE (20.8 ms vs. 95.5 ms). While the inference time increases when epistemic uncertainty is estimated with ensembles, this overhead is justified, since μGUIDE did not provide accurate epistemic uncertainty in our experiments.

Our results in vivo were consistent with findings in simulations. Post hoc aleatoric uncertainty maps contained multiple outliers, while MVE-generated uncertainty maps were smooth and more aligned with the standard deviation of noise within a homogenous liver ROI. Note that liver tissue is not expected to be perfectly homogeneous, given its vasculature. The variation within the ROI hence describes a combination from true variation in tissue properties, and noise in the parameter maps. Hence, as our networks only estimate the uncertainty, and not the anatomical variation, a small underestimation may be expected.

In addition to aleatoric uncertainty, we also evaluated epistemic uncertainty in an OOD setting by comparing liver scans (ID) with brain tumor scans (OOD). Ensembles of MVE models predicted significantly higher epistemic uncertainty in the tumor regions for v_e and k_e, confirming that the method is sensitive to domain shifts. The uncertainty increase in v_p was smaller and not statistically significant, which may be explained by several experimental factors. In particular, although abdominal scans were considered ID, the MVE models were trained on simulations, which do not perfectly replicate in vivo conditions. For example, noise may differ in vivo; there is cardiac and respiratory motion that are not simulated; and the physiological model in the tissue may not match the simulated model. As discussed above, the liver has a complex vascular architecture, which means the simulated contrast arrival is not representative, causing additional uncertainty in parameter estimates. This is especially the case for v_p, as the liver is provided by two inputs. Nonetheless, the clear increase in epistemic uncertainty in tumor regions supports the role of deep ensembles as a practical tool for detecting domain shifts.

Our findings suggest that MVE ensembles offer a promising strategy for uncertainty-aware DCE-MRI analysis. This approach not only improves the accuracy of pharmacokinetic parameter estimation but also provides meaningful uncertainty quantification, distinguishing between aleatoric and epistemic sources. Across all relevant evaluation points, such as prediction accuracy, robustness to noise, and ability to flag OOD inputs, the MVE ensembles method outperformed conventional methods. Based on this, we rec-

ommend the MVE ensemble as a robust and interpretable tool for advancing quantitative DCE-MRI.

The definition of "outlier" could be argued differently. In our definition, we have been very conservative, as not to identify real voxels as outliers, and have predominantly coupled it to physiologically feasible values (v_e, v_p). This approach ensures that reported outlier statistics reflect true extreme deviations rather than normal biological variability. Regardless of the exact definition, from Figure 5.7, it is clear even if less conservative thresholds were applied, NLLS exhibits the most outliers, followed by PINN, and then MVE.

Deep ensembles have a major drawback - their computational cost, which increases linearly with the ensemble size. A more computationally efficient alternative is the Monte Carlo dropout [41] method. Monte Carlo dropout estimates epistemic uncertainty by using dropout during inference and performing multiple forward passes with different dropout configurations. Despite its efficiency, Verdoja et al [42] demonstrated that uncertainty estimates derived from Monte Carlo dropout depend on the predicted parameter values, contradicting the fundamental nature of uncertainty. We observed similar behaviour (data not shown) when applying Monte Carlo dropout to our MVE neural network, suggesting that this method introduces biases in uncertainty estimation. Hence, we chose ensembles, despite their computational cost.

The proposed MVE neural network is supervised and requires ground-truth data for training. While ground-truth data can be generated through simulations, this data may not be representative for in vivo data. For example, real-world data may adhere to more complex pharmacokinetics and real-world noise may be more complex (e.g. due to image artefacts, Rician behaviour of noise in signal domain, motion, etc). This essentially renders in vivo data OOD. This is reflected in the fact that the reconstruction error in vivo is substantially larger than in simulations. So, likely, part of the epistemic uncertainty estimated in vivo rises due to a domain gap. As PINNs can be trained on real data, this would not be the case for PINNs. Hence, a promising solution could be a combination of physics-informed loss with MVE, so the uncertainty-aware model can be trained on in vivo data. However, our attempts to implement this approach resulted in unstable training due to challenges in uncertainty backpropagation (5.D, see models/mve_pinn.py on https://github.com/JonasElburgUVA/uncertainty-in-quantitative-dce-mri-public).

In this work, we classified uncertainties into two categories: aleatoric and epistemic [43]. While this classification is widely adopted, it remains somewhat ambiguous [44]. The distinction between these two types of uncertainty is often context-dependent and may vary across studies. To address this, future research should explore which evaluation frameworks could facilitate more meaningful comparison of uncertainty

quantification models in quantitative MRI.

The nature of uncertainties is partially explained by its classification as aleatoric or epistemic, but a deeper understanding of its origins can be achieved through explainable AI. Integrating these methods into uncertainty quantification could provide valuable insights into why specific predictions are made and which factors contribute most to uncertainty. For instance, feature attribution techniques could help distinguish whether uncertainty stems primarily from biological variability or poor data quality. If uncertainty is found to arise mainly due to low-quality imaging or acquisition artefacts, scanning protocols could be adjusted to improve data reliability, ultimately enhancing the accuracy and efficiency of further analyses. Moreover, explainable AI could help identify which acquisition points are most informative to minimize the uncertainty, allowing for optimised acquisition protocols.

In this work, we focused specifically on isolating the uncertainty inherent in the concentration curves themselves, as this provides the clearest framework for evaluating which methodology best captures it. In practice, however, many other sources contribute to uncertainty, including day-to-day physiological variability, errors in T1-estimation and signal-to-concentration conversion, and variability in data interpretation and contouring. Future networks could be extended to jointly model these additional sources of uncertainty. For instance, one could envision architectures that directly estimate DCE parameters from concatenated VFA and DCE signal data, leveraging the underlying physics in a manner similar to the uncertainty derivatives applied here. Even end-to-end multi-task networks incorporating automated contouring could be explored, though such extensions fall outside the scope of this study.

Building on this, another promising direction is the development of acquisition-independent uncertainty estimation networks. In DCE-MRI, acquisition protocols can vary significantly across organs, clinical applications, and research studies, with differences in temporal resolution and sampling patterns. The DCE-NET used in this work already supports variable input lengths through its gated recurrent unit (GRU) layers. Accommodating datasets with different sampling rates would require explicitly encoding temporal information into the network, either as an additional input feature or injected directly into the hidden state. Note that AIF parameters were already injected into the hidden state by Ottens et al [7]. These conventional GRUs do require regular sampling intervals. Alternatively, neural controlled differential equation (NCDE) networks, which are specifically designed for irregularly sampled and incomplete time-series data of variable length, could provide a more general solution [45]. The implementation and optimization of sequence-independent DCE-NET falls outside the scope of this paper.

Finally, future work could extend beyond parameter uncertainty to include model uncertainty, particularly in the choice of the signal or pharmacokinetic model. This would provide a framework for assessing whether the selected model adequately captures the underlying physiology for a given organ or pathology. Bayesian model comparison, for example, could quantify confidence in competing models or even highlight cases where none of the existing models accurately represent the observed dynamics [46, 47]. Detecting such mismatches would not only guide model selection in a data-driven way but also identify opportunities to refine or develop new models that better capture the physiology of interest.

Communicating uncertainty of parameters to clinicians is not routine, and we have several suggestions how this could be done. First, one could add the uncertainty maps for the clinician to interpret, but this would give clinicians even more challenging-to-interpret images to assess. On the other hand, clinicians are already used to interpreting noise in quantitative MRI parameter maps. Hence, one could use the uncertainty maps to add local noise to the deep learnt parameter maps of which the standard deviation depends on the local uncertainty. This would allow clinicians to flicker between low-noise deep learnt parameter maps, and noisy parameter maps to assess local precision (see Appendix E.17). A third approach would be to use the maps as a quality assurance tool by the technicians, both assessing the quality of the data (aleatoric uncertainty) and the model (epistemic uncertainty).

5.5.1 Limitations

One limitation is that the DCE-MRI scans of the liver were acquired from healthy volunteers in this study. While we implemented the extended Tofts model - commonly used for tumour characterisation - we did not evaluate our approach on pathological data. Since tumours are known for their heterogeneous microenvironment, it remains an open question whether the MVE network can accurately distinguish between atypical parameter estimates arising from noise and those reflecting true biological variations. While tumour imaging could potentially benefit from uncertainty-aware models that flag unreliable parameter estimates, this application is outside the scope of the current work and would require dedicated datasets and validation strategies.

Another limitation of this study is the use of a single input function when evaluating extended Tofts parameters in the liver. While our network design could accommodate a dual-input function—including both the arterial input function (AIF) and the portal venous input function—our analysis was based solely on the AIF. Using a dual-input function could potentially improve the accuracy of liver perfusion estimates. Additionally, it could lower

the epistemic uncertainty as a result of improved fit. We chose a single input function because population-based input functions are only available for AIF and have demonstrated the highest repeatability [48]. In contrast, the patient-specific portal venous input function in our dataset exhibited high variability, with values fluctuating depending on the ROI used for its calculation.

Another important consideration is that the deep-learning models used in this study are voxel-based, and thus treat voxels as independent observations. In practice, in vivo data exhibit a hierarchical structure, with voxels nested within subjects. Explicitly modeling this structure could improve pharmacokinetic parameter estimation and lead to smaller uncertainty estimates. Future extensions could explore hierarchical deep learning architectures that better capture voxel correlations and subject-level effects.

5.5.2 Conclusion

In conclusion, we developed a neural network capable of predicting pharmacokinetic parameters based on the extended Tofts model while simultaneously estimating aleatoric and epistemic uncertainties. This allows us to separate data and model confidence. By providing both parameter estimates and their associated uncertainties, this framework enhances the reliability of pharmacokinetic modelling and can help clinicians make more informed decisions.

Chapter appendix

5.A Signal model

Most pharmacokinetic models describe the temporal evolution of contrast agent concentration within tissue. Converting the measured signal into concentration values requires several processing steps.

First, baseline T_{10} maps must be estimated from a separate scan, for example, using a variable flip angles acquisition or inversion recovery. Then a signal model $S(t)$ is applied to derive the effective $T_1(t)$ values at each time point of DCE-MRI acquisition. In this work, for our in vivo data, we used the spoiled gradient echo signal equation, assuming a short echo time [49]:

$$S(t) = M_0 \sin\alpha \frac{1 - e^{-\frac{TR}{T_1(t)}}}{1 - e^{-\frac{TR}{T_1}} \cos\alpha} \tag{5.10}$$

where TR is the repetition time, α is the flip angle, and M_0 is the net magnetisation. By inverting Equation 5.10, we obtain $T_1(t)$ at each time point. The corresponding contrast agent concentration $C(t)$ is then calculated using:

$$C(t) = \frac{1}{r}\left(\frac{1}{T(t)} - \frac{1}{T_{10}}\right) \tag{5.11}$$

where r is the relaxivity of the contrast agent. Once $C(t)$ is determined using Equations 5.10 and 5.11, pharmacokinetic modelling can be applied to estimate physiologically relevant parameters.

5.B Arterial input function

An AIF model used in this article and earlier proposed in [24] is expressed as:

$$C_p(t, \theta_{AIF}) = \begin{cases} a_B(1 - cos(\mu_B t)) + a_B a_G f(t, \mu_G), & \text{for } 0 \leq t \leq t_B \\ a_B a_G f(t_B, \mu_G) e^{-\mu_G(t - t_B)}, & \text{for } t > t_B \end{cases} \tag{5.12}$$

where $\theta_{AIF} = \{a_B, a_G, \mu_B, \mu_G\}, t_B = 2\pi\mu_B^{-1}$, and

$$f(t,\alpha) = \frac{1}{\alpha}(1 - e^{-\alpha t}) - \frac{1}{\alpha^2 + \mu_B^2}(\alpha \cos \mu_B t + \mu_B \sin(\mu_B t) - \alpha e^{-\alpha t})$$

.

Parameters θ_{AIF} define the shape of AIF.

5.C DCE-MRI acquisition protocol

MRI scans were acquired on a 3T system (Ingenia, Philips, Best, The Netherlands) with a 70 cm bore diameter, using a 16-channel phased-array anterior coil and a 16-channel phased-array posterior coil. After approximately 30 s, 0.2 mmol/kg of 0.5 mmol/mL gadoterate meglumine (Dotarem, Guerbet, Villepinte, France) was administered at a rate of 5 mL/s followed by a 15 mL saline flush at the same rate. The DCE protocol comprised of a dynamic series of a 3D golden angle radial stack of stars fast field echo (FFE) spoiled gradient echo sequence with SPIR fat suppression [50]. Scan parameters can be found in Table 5.C.1. The dynamic series were reconstructed in MATLAB (R2021b, MathWorks) with a GRASP algorithm [51] with regularisation in the temporal domain (λ = 0.01) at a temporal resolution of 3.8 s/frame.

Contrast	Spoiled gradient echo
Sampling trajectory	Golden angle radial stack-of-stars
Field strength	3.0 T
Flip angle	22°
TE	1.39 ms
TR	3.48 ms
Relaxivity contrast agent	4.3 L/mmol-s
Acquisition matrix	68 reconstructed slices, 1045 readouts per slice, 480 sampling points per readout
Acquisition orientation	Axial
Acquired and reconstructed spatial resolution	1.5 x 1.5 x 3.5 mm
Matrix size	240 x 240
Number of slices	68
Total scan duration	5.33 min
Reconstructed temporal resolution	4 s/frame

Table 5.C.1: Acquisition parameters of DCE-MRI protocol of the abdomen in healthy subjects.

5.D Physics-Informed Mean-Variance Estimation

Our research indicates a physics-informed mean-variance estimation (PI-MVE) model could provide a superior uncertainty-aware model in vivo. In this work, some attempts have been made at realising this. Attempts at including uncertainty propagation within the model training were hindered by unstable training dynamics, and ultimately did not lead to the expected results.

5.D.1 Theory

The MVE loss normally assumes supervised training. We further attempted to extend the MVE framework to the physics-informed paradigm. This was achieved by propagating both the mean estimates and the uncertainty estimates through the extended Tofts-Kety model to the loss-domain. As in the deterministic PINN objective, the mean estimates were propagated from parameter space to time-concentration space by calculating the reconstructed mean curve $\hat{C}_{t\mu} \coloneqq C_t(\hat{\boldsymbol{u}}_\mu)$. The variance in parameters was approximately propagated through the extended Tofts-Kety model using the linear error propagation formula [52]:

$$\hat{C}_{t\sigma^2} \approx \boldsymbol{J}^{\odot 2}\hat{\boldsymbol{u}}_{\sigma^2}, \tag{5.13}$$

where $\boldsymbol{J}^{\odot 2}$ is the element-wise square of the Jacobian matrix with entries $\frac{\partial \hat{C}_{t\mu}}{\partial \hat{\boldsymbol{u}}_\mu}$ representing partial derivatives. This enables the formulation of the GNLL in data space, creating a physics-informed MVE objective:

$$\mathcal{L}_{\text{PI-GNLL}} = \frac{1}{2}\left(\frac{(\hat{C}_{t\mu} - \tilde{C}_t)^2}{\hat{C}_{t\sigma^2}} + \log \hat{C}_{t\sigma^2}\right). \tag{5.14}$$

The linear error propagation formula assumes small errors in parameters, and may not fully capture non-linearities in the extended Tofts-Kety model. In preliminary experiments, we found that the complex relationship between uncertainty in the data and uncertainty in parameter predictions in combination with the approximation can cause optimisation issues, including convergence to local minima where the parameter estimates are very poor. As a means to prevent this, a regularisation term was added to the loss, which was based on the difference between the PI-MVE model's predicted mean estimates $\hat{\boldsymbol{u}}_{\mu\text{PI-MVE}}$ and the predictions made by a deterministic physics-informed model $\hat{\boldsymbol{u}}_{\text{PINN}}$. This regularisation helped to prevent predictions from violating the assumptions of linear error propagation, thereby stabilising the optimisation process. The final objective

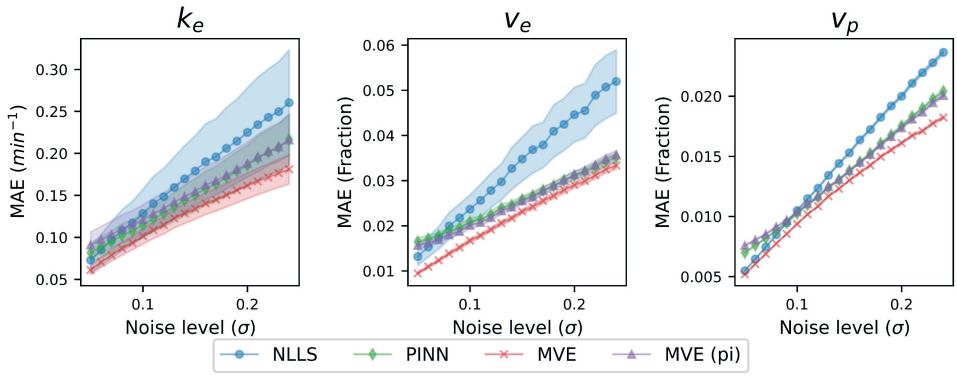

Figure 5.D.1: Accuracy results including the physics-informed MVE model

is thus formulated as follows:

$$\mathcal{L}_{\text{PI-GNLL-R}} = \frac{1}{2}\left(\frac{(\hat{C}_{t\mu} - \tilde{C}_t)^2}{\hat{C}_{t\sigma^2}} + \log \hat{C}_{t\sigma^2}\right) + \lambda_{\text{MVE}} e^{|\hat{\boldsymbol{u}}_{\mu\text{PI-MVE}} - \hat{\boldsymbol{u}}_{\mu\text{PINN}}|}, \tag{5.15}$$

where λ is a tunable regularisation factor. We used a relatively high $\lambda = 5$.

5.D.2 Results

The regularised PI-MVE neural network trained on simulations obtained similar accuracy as a regular PINN model, as was expected due to the high regularisation of mean estimates (Figure 5.D.1).

While some correlation between predicted and observed error was found, the PI-MVE model produced 100% outliers in the k_e and v_p parameters (Figure 5.D.2). This indicates significant scaling issues in the uncertainty estimates. It was noted that the loss at testing time was very high, which is consistent with the inflated uncertainty estimates. The observations are likely a result of optimisation issues, severely impacting the interpretability of the uncertainty estimates. Further research may attempt to use a different architecture such as a transformer to better minimise the loss objective, or a multivariate Gaussian negative log likelihood-based loss to better model the relationship between pharmacokinetic parameters and the reconstructed curve.

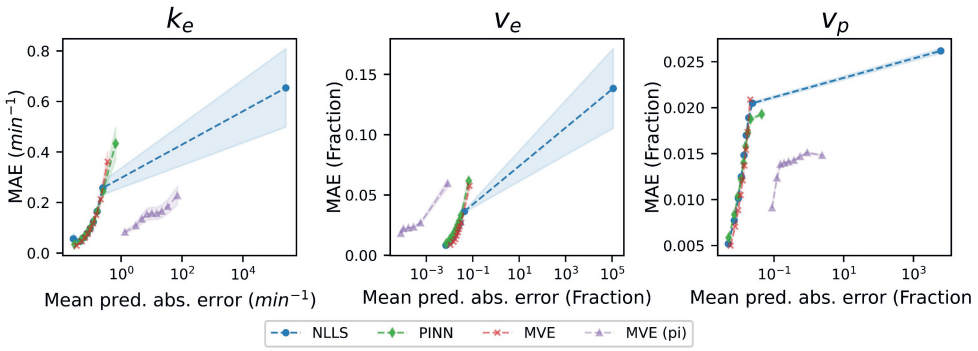

Figure 5.D.2: Observed by predicted absolute error including outliers and the physics-informed mean-variance estimation model.

5.E Additional visualisations

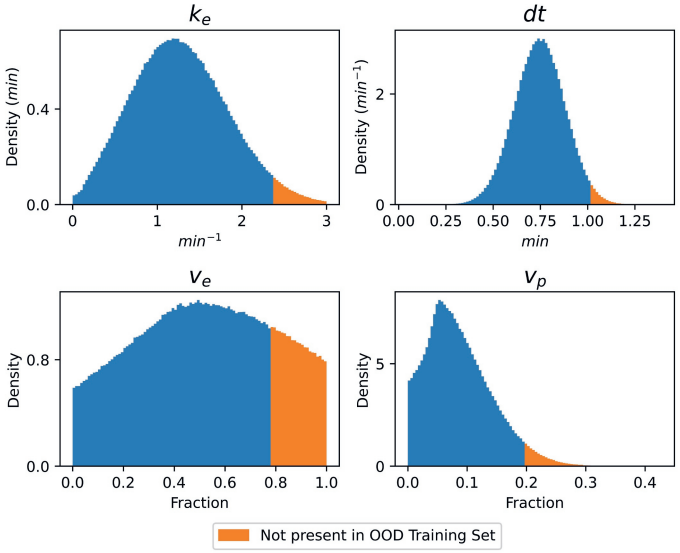

Figure 5.E.1: Normalized distribution of the simulated dataset, generated by fitting a multivariate Gaussian distribution to pharmacokinetic parameters estimated by DCE-NET from in vivo scans. Data not included in the training set during the OOD experiment marked in orange.

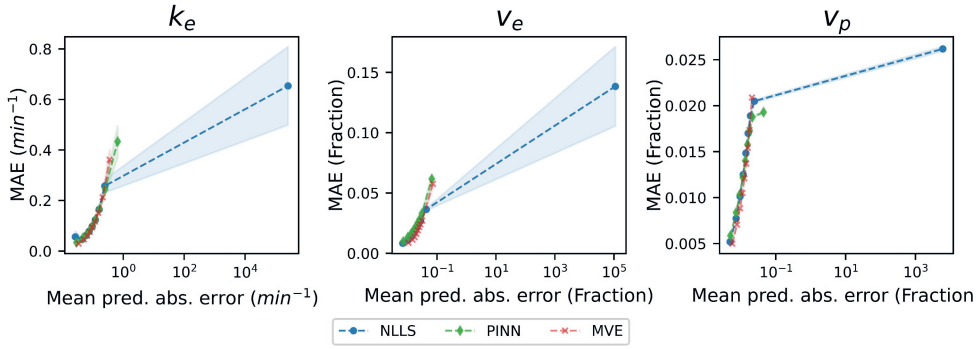

Figure 5.E.2: Average error by average predicted error including outliers. Note that the x-axis is in log-scale. The outliers had a very large effect on the predicted absolute error in the final bin of NLLS, and to a lesser extent, PINN estimates.

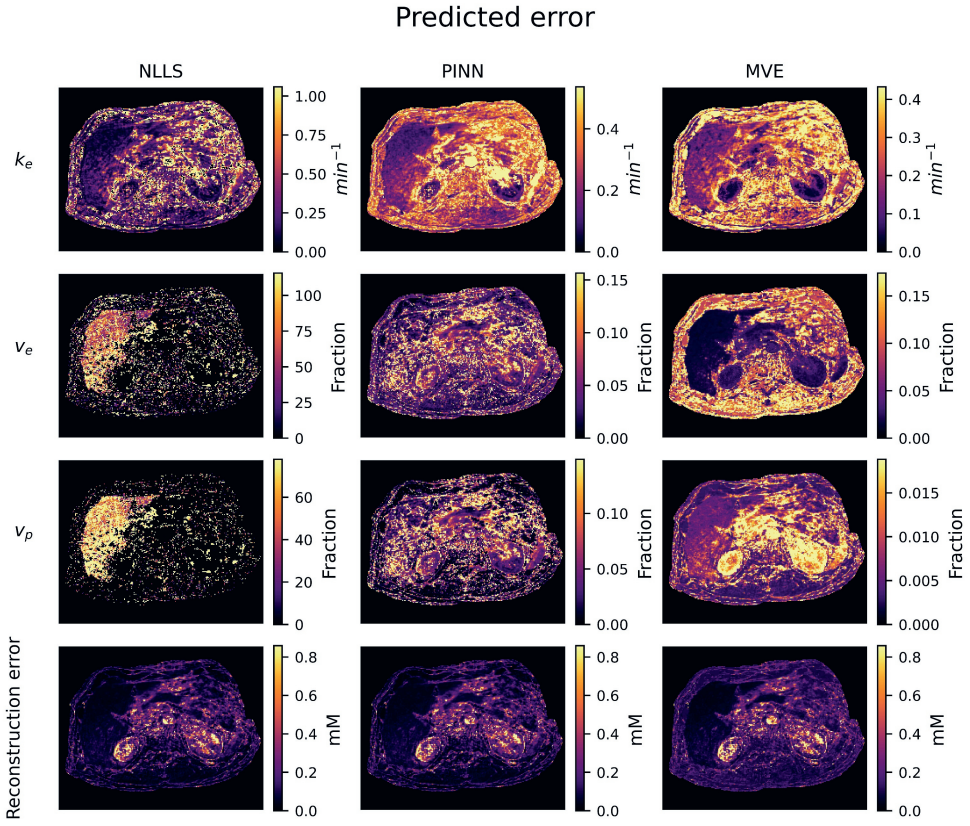

Figure 5.E.3: Predicted error map including outliers, with each visualisation having its separate scale.

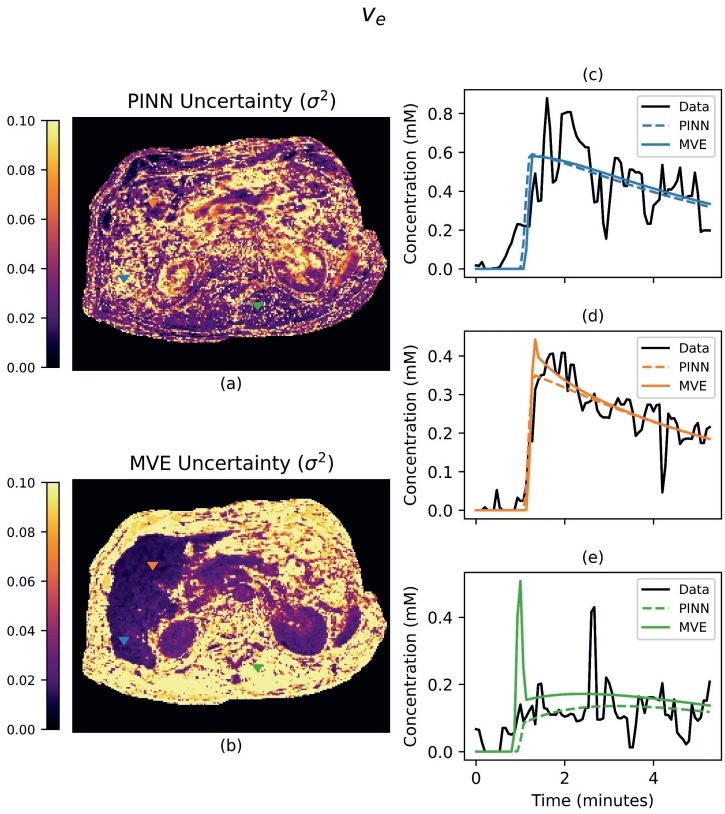

Figure 5.E.4: The predicted uncertainty in the v_e parameter compared between a PINN and MVE model (a-b). On the right, we visualised three input curves and the reconstructions matching the indicators on the slices. Figure (c) shows a curve in the liver where the $\sigma^2_{\text{PINN}} > \sigma^2_{\text{MVE}}$ $(7.7 \cdot 10^{-1} > 5.1 \cdot 10^{-4})$, (d) shows a curve in the liver where $\sigma^2_{\text{PINN}} \approx \sigma^2_{\text{MVE}}$ $(7.3 \cdot 10^{-4} \approx 3.0 \cdot 10^{-4})$, and (e) is a curve in muscle where $\sigma^2_{\text{PINN}} < \sigma^2_{\text{MVE}}$ $(2.8 \cdot 10^{-4} < 1.5 \cdot 10^{-2})$.

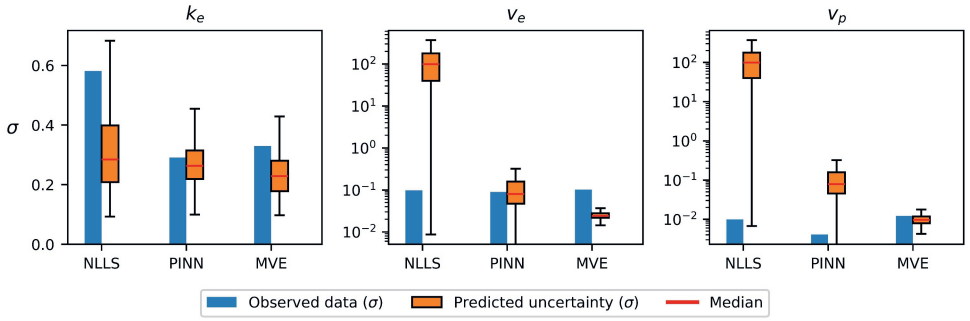

Figure 5.E.5: Full scale predicted standard deviations in the liver ROI. Note that the y-axes of the second and third plot are in log-scale.

Bibliography

[1] M. Leach, B. Morgan, P. Tofts, D. Buckley, W. Huang, M. Horsfield, T. Chenevert, D. Collins, A. Jackson, D. Lomas *et al.*, "Imaging vascular function for early stage clinical trials using dynamic contrast-enhanced magnetic resonance imaging," *European radiology*, vol. 22, pp. 1451–1464, 2012.

[2] L. Bokacheva, H. Rusinek, J. L. Zhang, and V. S. Lee, "Assessment of renal function with dynamic contrast-enhanced mr imaging," *Magnetic resonance imaging clinics of North America*, vol. 16, no. 4, pp. 597–611, 2008.

[3] K. Villringer, B. E. Sanz Cuesta, A.-C. Ostwaldt, U. Grittner, P. Brunecker, A. A. Khalil, K. Schindler, O. Eisenblätter, H. Audebert, and J. B. Fiebach, "Dce-mri blood–brain barrier assessment in acute ischemic stroke," *Neurology*, vol. 88, no. 5, pp. 433–440, 2017.

[4] R. Raja, G. A. Rosenberg, and A. Caprihan, "Mri measurements of blood-brain barrier function in dementia: a review of recent studies," *Neuropharmacology*, vol. 134, pp. 259–271, 2018.

[5] S. P. Sourbron and D. L. Buckley, "Classic models for dynamic contrast-enhanced mri," *NMR in Biomedicine*, vol. 26, no. 8, pp. 1004–1027, 2013.

[6] R. Klaassen, O. J. Gurney-Champion, J. W. Wilmink, M. G. Besselink, M. R. Engelbrecht, J. Stoker, A. J. Nederveen, and H. W. van Laarhoven, "Repeatability and correlations of dynamic contrast enhanced and t2* mri in patients with advanced pancreatic ductal adenocarcinoma," *Magnetic resonance imaging*, vol. 50, pp. 1–9, 2018.

[7] T. Ottens, S. Barbieri, M. R. Orton, R. Klaassen, H. W. van Laarhoven, H. Crezee, A. J. Nederveen, X. Zhen, and O. J. Gurney-Champion, "Deep learning dce-mri parameter estimation: Application in pancreatic cancer," *Medical image analysis*, vol. 80, p. 102512, 2022.

[8] Y. Bliesener, J. Acharya, and K. S. Nayak, "Efficient dce-mri parameter and uncertainty estimation using a neural network," *IEEE transactions on medical imaging*, vol. 39, no. 5, pp. 1712–1723, 2019.

[9] E. M. Dejene, W. Brenner, M. R. Makowski, and C. Kolbitsch, "Unified bayesian network for uncertainty quantification of physiological parameters in dynamic contrast enhanced (dce) mri of the liver," *Physics in Medicine & Biology*, vol. 68, no. 21, p. 215018, 2023.

[10] J. Kettelkamp and S. G. Lingala, "Arterial input function and tracer kinetic model-driven network for rapid inference of kinetic maps in dynamic contrast-enhanced mri (aif-tk-net)," in *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2020, pp. 1450–1453.

[11] C. Ulas, D. Das, M. J. Thrippleton, M. d. C. Valdés Hernández, P. A. Armitage, S. D. Makin, J. M. Wardlaw, and B. H. Menze, "Convolutional neural networks for direct inference of pharmacokinetic parameters: application to stroke dynamic contrast-enhanced mri," *Frontiers in neurology*, vol. 9, p. 1147, 2019.

[12] J. Zou, J. M. Balter, and Y. Cao, "Estimation of pharmacokinetic parameters from dce-mri by extracting long and short time-dependent features using an lstm network," *Medical physics*, vol. 47, no. 8, pp. 3447–3457, 2020.

[13] N. G. Gyori, M. Palombo, C. A. Clark, H. Zhang, and D. C. Alexander, "Training data distribution significantly impacts the estimation of tissue microstructure with machine learning," *Magnetic resonance in medicine*, vol. 87, no. 2, pp. 932–947, 2022.

[14] A. Nguyen, J. Yosinski, and J. Clune, "Deep neural networks are easily fooled: High confidence predictions for unrecognizable images," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 427–436.

[15] E. Hüllermeier and W. Waegeman, "Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods," *Machine learning*, vol. 110, no. 3, pp. 457–506, 2021.

[16] K. W. Vugrin, L. P. Swiler, R. M. Roberts, N. J. Stucky-Mack, and S. P. Sullivan, "Confidence region estimation techniques for nonlinear regression in groundwater flow: Three case studies," *Water Resources Research*, vol. 43, no. 3, 2007.

[17] S.-F. Shih, S. G. Kafali, T. Armstrong, X. Zhong, K. L. Calkins, and H. H. Wu, "Deep learning-based parameter mapping with uncertainty estimation for fat quantification using accelerated free-breathing radial mri," in *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2021, pp. 433–437.

[18] F. Glang, A. Deshmane, S. Prokudin, F. Martin, K. Herz, T. Lindig, B. Bender, K. Scheffler, and M. Zaiss, "Deepcest 3t: Robust mri parameter determination and uncertainty quantification with neural networks—application to cest imaging of the human brain at 3t," *Magnetic resonance in medicine*, vol. 84, no. 1, pp. 450–466, 2020.

[19] M. Jallais and M. Palombo, "Introducing μguide for quantitative imaging via generalized uncertainty-driven inference using deep learning," *Elife*, vol. 13, p. RP101069, 2024.

[20] L. Zhang, V. Vishnevskiy, A. Jakab, and O. Goksel, "Implicit modeling with uncertainty estimation for intravoxel incoherent motion imaging," in *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*. IEEE, 2019, pp. 1003–1007.

[21] P. S. Tofts, "Modeling tracer kinetics in dynamic gd-dtpa mr imaging," *Journal of magnetic resonance imaging*, vol. 7, no. 1, pp. 91–101, 1997.

[22] P. S. Tofts, G. Brix, D. L. Buckley, J. L. Evelhoch, E. Henderson, M. V. Knopp, H. B. Larsson, T.-Y. Lee, N. A. Mayr, G. J. Parker *et al.*, "Estimating kinetic parameters from dynamic contrast-enhanced t1-weighted mri of a diffusable tracer: standardized quantities and symbols," *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 10, no. 3, pp. 223–232, 1999.

[23] X. Yang and M. V. Knopp, "Quantifying tumor vascular heterogeneity with dynamic contrast-enhanced magnetic resonance imaging: a review," *BioMed Research International*, vol. 2011, no. 1, p. 732848, 2011.

[24] M. R. Orton, J. A. d'Arcy, S. Walker-Samuel, D. J. Hawkes, D. Atkinson, D. J. Collins, and M. O. Leach, "Computationally efficient vascular input function models for quantitative kinetic modelling using dce-mri," *Physics in Medicine & Biology*, vol. 53, no. 5, p. 1225, 2008.

[25] T. S. Ahearn, R. T. Staff, T. W. Redpath, and S. I. K. Semple, "The use of the levenberg–marquardt curve-fitting algorithm in pharmacokinetic modelling of dce-mri data," *Physics in Medicine & Biology*, vol. 50, no. 9, p. N85, 2005.

[26] D. L. Buckley, "Uncertainty in the analysis of tracer kinetics using dynamic contrast-enhanced t1-weighted mri," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 47, no. 3, pp. 601–606, 2002.

[27] Y. Guo, S. G. Lingala, Y. Bliesener, R. M. Lebel, Y. Zhu, and K. S. Nayak, "Joint arterial input function and tracer kinetic parameter estimation from undersampled dynamic contrast-enhanced mri using a model consistency constraint," *Magnetic resonance in medicine*, vol. 79, no. 5, pp. 2804–2815, 2018.

[28] E. Henderson, J. Sykes, D. Drost, H.-J. Weinmann, B. K. Rutt, and T.-Y. Lee, "Simultaneous mri measurement of blood flow, blood volume, and capillary permeability in mammary tumors using two different contrast agents," *Journal of Magnetic Resonance Imaging*, vol. 12, no. 6, pp. 991–1003, 2000.

[29] K. Murase, "Efficient method for calculating kinetic parameters using t1-weighted dynamic contrast-enhanced magnetic resonance imaging," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 51, no. 4, pp. 858–862, 2004.

[30] D. A. Nix and A. S. Weigend, "Estimating the mean and variance of the target probability distribution," in *Proceedings of 1994 ieee international conference on neural networks (ICNN'94)*, vol. 1. IEEE, 1994, pp. 55–60.

[31] L. Sluijterman, E. Cator, and T. Heskes, "Optimal training of mean variance estimation neural networks," *arXiv preprint arXiv:2302.08875*, 2023.

[32] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," *Advances in neural information processing systems*, vol. 30, 2017.

[33] G. J. Parker, C. Roberts, A. Macdonald, G. A. Buonaccorsi, S. Cheung, D. L. Buckley, A. Jackson, Y. Watson, K. Davies, and G. C. Jayson, "Experimentally-derived functional form for a population-averaged high-temporal-resolution arterial input function for dynamic contrast-enhanced mri," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 56, no. 5, pp. 993–1000, 2006.

[34] P. J. van Houdt, S. Ragunathan, M. Berks, Z. Ahmed, L. E. Kershaw, O. J. Gurney-Champion, S. Tadimalla, J. Arvidsson, Y. Sun, J. Kallehauge *et al.*, "Contrast-agent-based perfusion mri code repository and testing framework: Ismrm open science initiative for perfusion imaging (osipi)," *Magnetic resonance in medicine*, vol. 91, no. 5, pp. 1774–1786, 2024.

[35] M. D. Holland, A. Morales, S. Simmons, B. Smith, S. R. Misko, X. Jiang, D. A. Hormuth, C. Christenson, R. P. Koomullil, D. E. Morgan *et al.*, "Disposable point-of-care portable

perfusion phantom for quantitative dce-mri," *Medical physics*, vol. 49, no. 1, pp. 271–281, 2022.

[36] C. M. De Bazelaire, G. D. Duhamel, N. M. Rofsky, and D. C. Alsop, "Mr imaging relaxation times of abdominal and pelvic tissues measured in vivo at 3.0 t: preliminary results," *Radiology*, vol. 230, no. 3, pp. 652–659, 2004.

[37] E. S. Shalom, H. Kim, R. A. van Der Heijden, Z. Ahmed, R. Patel, D. A. Hormuth, J. C. DiCarlo, T. E. Yankeelov, N. J. Sisco, R. D. Dortch *et al.*, "The ismrm open science initiative for perfusion imaging (osipi): Results from the osipi–dynamic contrast-enhanced challenge," *Magnetic resonance in medicine*, vol. 91, no. 5, pp. 1803–1821, 2024.

[38] M.-H. Laves, S. Ihler, J. F. Fast, L. A. Kahrs, and T. Ortmaier, "Recalibration of aleatoric and epistemic regression uncertainty in medical imaging," *arXiv preprint arXiv:2104.12376*, 2021.

[39] F. Musil, M. J. Willatt, M. A. Langovoy, and M. Ceriotti, "Fast and accurate uncertainty estimation in chemical machine learning," *Journal of chemical theory and computation*, vol. 15, no. 2, pp. 906–915, 2019.

[40] F. Isensee, M. Rokuss, L. Krämer, S. Dinkelacker, A. Ravindran, F. Stritzke, B. Hamm, T. Wald, M. Langenberg, C. Ulrich *et al.*, "nninteractive: Redefining 3d promptable segmentation," *arXiv preprint arXiv:2503.08373*, 2025.

[41] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *international conference on machine learning*. PMLR, 2016, pp. 1050–1059.

[42] F. Verdoja and V. Kyrki, "Notes on the behavior of mc dropout," *arXiv preprint arXiv:2008.02627*, 2020.

[43] A. Der Kiureghian and O. Ditlevsen, "Aleatory or epistemic? does it matter?" *Structural safety*, vol. 31, no. 2, pp. 105–112, 2009.

[44] B. Mucsányi, M. Kirchhof, and S. J. Oh, "Benchmarking uncertainty disentanglement: Specialized uncertainties for specialized tasks," *Advances in neural information processing systems*, vol. 37, pp. 50 972–51 038, 2024.

[45] D. Kuppens, S. Barbieri, D. v. d. Berg, P. Schouten, H. C. Thoeny, M. Wennen, and O. J. Gurney-Champion, "Acquisition-independent deep learning for quantitative mri parameter estimation using neural controlled differential equations," *arXiv preprint arXiv:2412.20844*, 2024.

[46] C. Duan, J. F. Kallehauge, G. L. Bretthorst, K. Tanderup, J. J. Ackerman, and J. R. Garbow, "Are complex dce-mri models supported by clinical data?" *Magnetic resonance in medicine*, vol. 77, no. 3, pp. 1329–1339, 2017.

[47] H. Bagher-Ebadian, S. L. Brown, M. M. Ghassemi, P. C. Acharya, I. J. Chetty, B. Movsas, J. R. Ewing, and K. Thind, "Probabilistic nested model selection in pharmacokinetic analysis of dce-mri data in animal model of cerebral tumor," *Scientific Reports*, vol. 15, no. 1, p. 1786, 2025.

[48] M. Rata, D. J. Collins, J. Darcy, C. Messiou, N. Tunariu, N. Desouza, H. Young, M. O. Leach, and M. R. Orton, "Assessment of repeatability and treatment response in early phase clinical trials using dce-mri: comparison of parametric analysis using mr-and ct-derived arterial input functions," *European radiology*, vol. 26, pp. 1991–1998, 2016.

[49] M. A. Bernstein, K. F. King, and X. J. Zhou, *Handbook of MRI pulse sequences*. Elsevier, 2004.

[50] C. S. Zee, H. D. Segall, M. R. Terk, S. Destian, J. Ahmadi, J. R. Gober, and P. M. Colletti, "Spir mri in spinal diseases," *J Comput Assist Tomogr*, vol. 16, no. 3, pp. 356–360, 1992.

[51] L. Feng, L. Axel, H. Chandarana, K. T. Block, D. K. Sodickson, and R. Otazo, "Xd-grasp: golden-angle radial mri with reconstruction of extra motion-state dimensions using compressed sensing," *Magnetic resonance in medicine*, vol. 75, no. 2, pp. 775–788, 2016.

[52] H. H. Ku *et al.*, "Notes on the use of propagation of error formulas," *Journal of Research of the National Bureau of Standards*, vol. 70, no. 4, 1966.

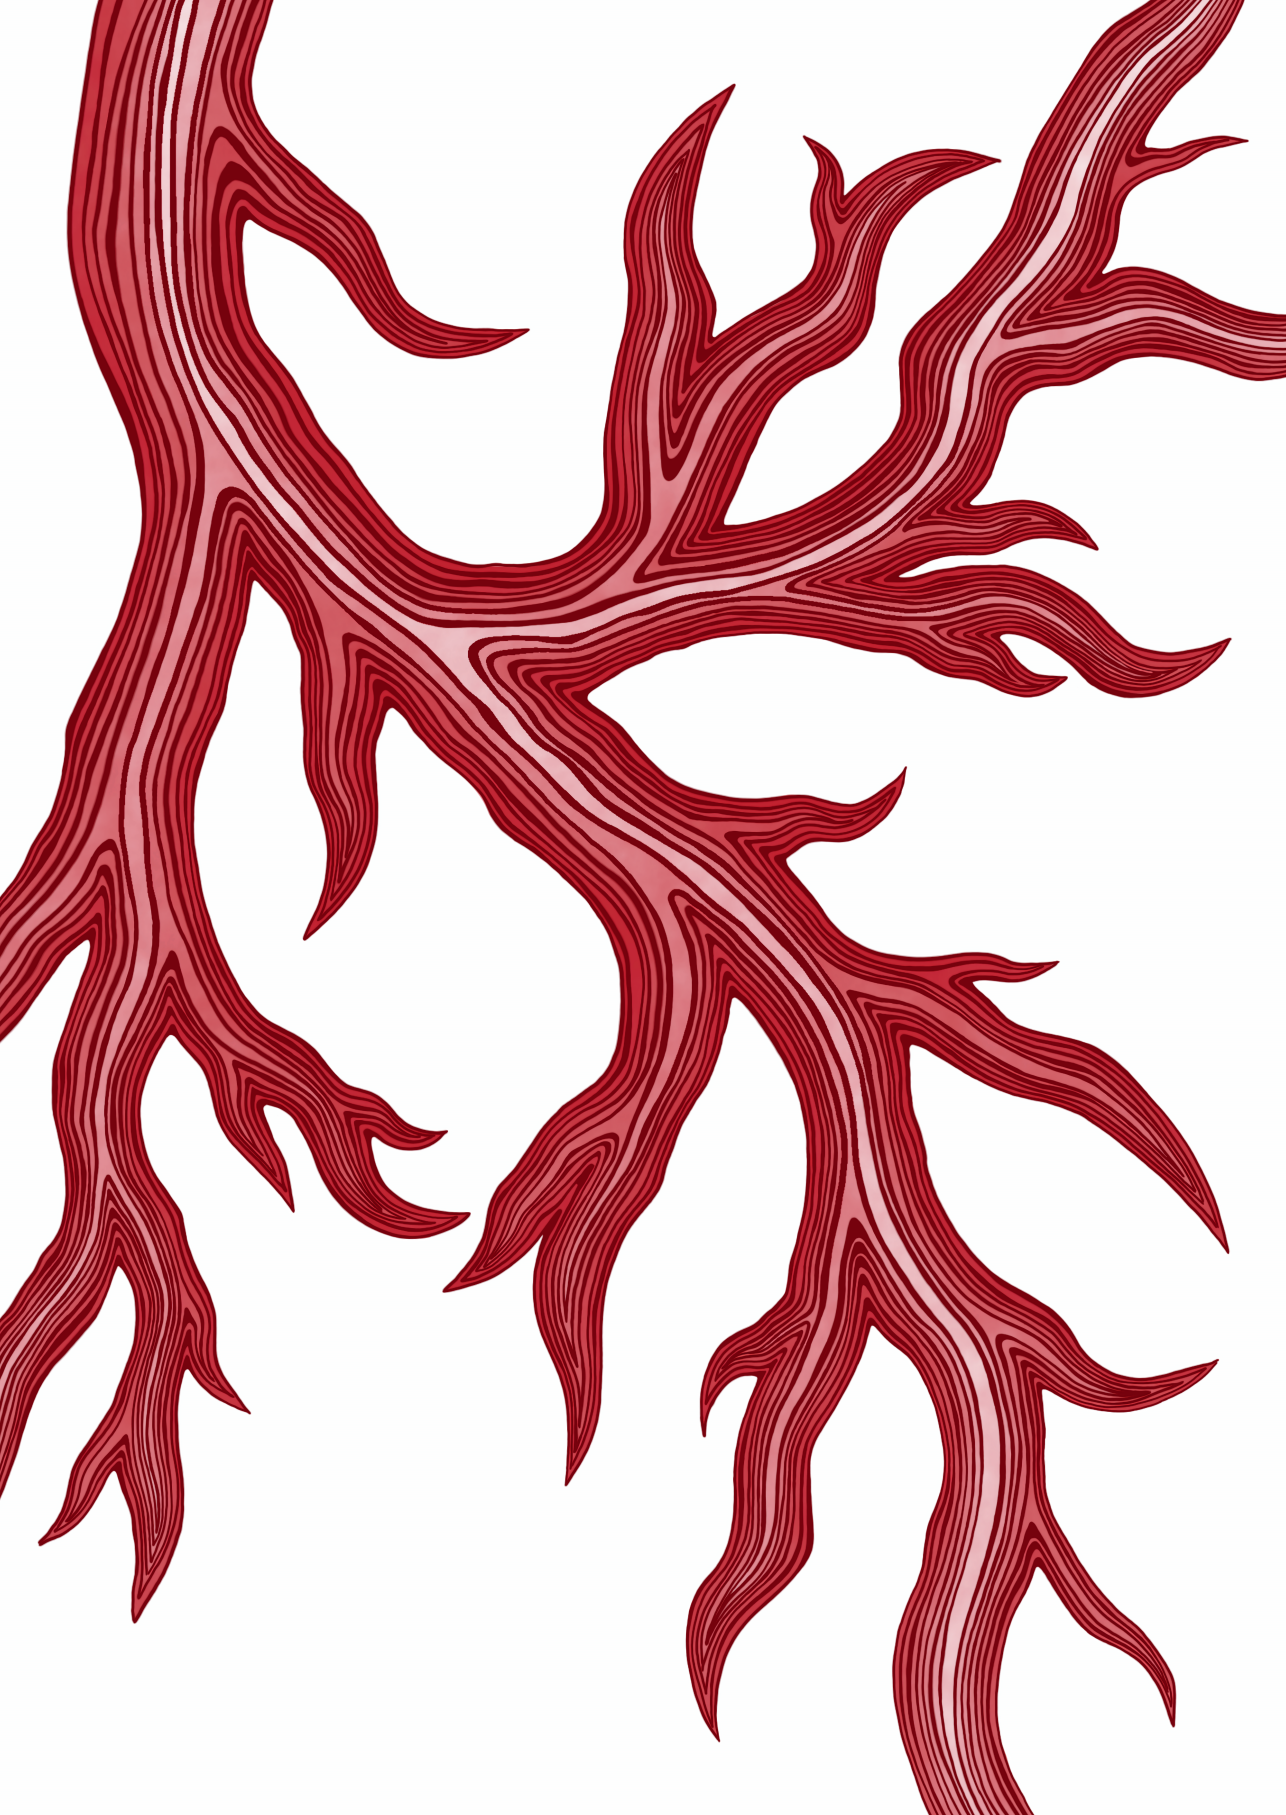

We do not need magic to change the world, we carry all the power we need inside ourselves already: we have the power to imagine better.

J. K. Rowling

6

General discussion

Quantitative DCE-MRI provides patient-specific information on tissue perfusion and vascular permeability, offering valuable insight into disease state, progression, and response to therapy. Despite its considerable potential, quantitative DCE-MRI remains largely a research-oriented technique due to numerous methodological and practical challenges. This thesis addresses several of these challenges, with a particular focus on image reconstruction, modelling, and post-processing. This general discussion builds on the technical work presented in this thesis, while also extending beyond it to consider the broader set of requirements for clinical translation of quantitative DCE-MRI, including image quality, reproducibility, and clinical validation.

6.1 Image quality

Image quality is a fundamental requirement for the successful clinical translation of quantitative DCE-MRI. It encompasses not only visually appealing images but, more importantly, the accuracy and precision of the derived quantitative parameters. Without sufficient image quality, it is not possible to demonstrate added clinical value, and clinicians are unlikely to adopt quantitative DCE-MRI as a reliable tool for clinical decision-making.

6.1.1 Spatial and temporal resolution

Vascular heterogeneity within tumors underlies the requirement for fine spatial resolution in DCE-MRI, as variations in microvascular density and permeability can occur over small spatial scales [1, 2]. Nevertheless, even at fine spatial resolution, intra-voxel heterogeneity often persists, particularly in highly heterogeneous tumors. Although it may seem intuitive to assume that the contrast agent concentration–time curve of a voxel represents a linear superposition of curves from homogeneous sub-voxel compartments, this assumption does not generally hold. Due to the nonlinear, exponential relationship between contrast agent concentration $C(t)$ and pharmacokinetic rate constants such as K^{trans} and k_{ep}, the voxel-averaged signal does not correspond to the average of the underlying kinetic parameters [3]. As a result, intra-voxel heterogeneity can introduce systematic bias into parameter estimates.

To minimize these effects, DCE-MRI data should be acquired with the highest spatial resolution that is practically achievable. However, spatial resolution is intrinsically linked to temporal resolution, which is constrained by pulse sequence parameters, anatomical coverage, and SNR. For liver imaging, QIBA recommends a temporal resolution of approximately 3–5 s/frame [4]. Ideally, even faster temporal sampling would be preferred, as the rapid contrast dynamics observed in AIF require temporal resolutions on the order of 1 s/frame for accurate characterization [5]. One potential strategy to mitigate this spatio-temporal trade-off is the direct estimation of pharmacokinetic parameters from undersampled k-space data using model-based reconstruction, as explored in **Chapter 4**. The feasibility and robustness of this approach, however, strongly depend on the complexity of the pharmacokinetic model and the acceleration factor imposed during data acquisition.

6.1.2 Motion

Because DCE-MRI acquisitions typically extend over several minutes, patient motion is inevitable. In abdominal and thoracic imaging, breath-holding is generally not feasible for the entire duration of the scan, leading to respiratory motion, while additional sources of motion include cardiac pulsation, peristaltic activity, and bulk patient movement. When combined with the high spatial and temporal resolution requirements of quantitative DCE-MRI, uncorrected motion can significantly distort the measured signal–time curves (tissues move between dynamic images) and introduce image artifacts (tissues move within acquisition of one image), such as blurring (signal averaging over moving structures), aliasing (misregistration of spatial frequency component), and ghosting (replicas of moving objects along the phase-encoding direction) [6]. These artifacts directly propagate into phar-

macokinetic modeling, leading to biased or unstable parameter estimates [7, 8]. Consequently, accurate detection and correction of motion is a critical prerequisite for achieving reliable and reproducible quantitative DCE-MRI.

Non-Cartesian imaging has become clinically established for abdominal DCE-MRI due to its inherent robustness to motion, particularly with radial sampling schemes [9]. Nevertheless, respiratory motion can still lead to residual blurring and incoherent artifacts, especially in the absence of explicit motion correction [10]. Because organs such as the liver, pancreas, and heart undergo complex non-rigid deformations during free breathing, non-rigid motion correction is often required for most body imaging applications [6]. To address these deformation patterns, image registration techniques can be employed [11, 12]. However, multi-contrast registration often struggles because residual motion persists when registration fails, particularly between images with varying contrast [13]. Motion-corrected reconstruction may provide a more reliable solution, as it avoids performing registration on images with differing contrast and explicitly accounts for intra-scan motion. This is the approach we implemented and evaluated in **Chapter 3**.

6.1.3 Future outlook

A logical next step in advancing quantitative DCE-MRI is to integrate model-based reconstruction (**Chapter 4**) and motion-corrected reconstruction (**Chapter 3**) into a unified reconstruction framework that directly estimates pharmacokinetic parameters while compensating for motion. In such a framework, motion fields would be incorporated together with the pharmacokinetic signal model into the forward operator that maps k-space data to quantitative parameters. Beyond compensating for motion, defining the motion fields as optimized parameters for each time frame may further enhance reconstruction quality by reducing reliance on assumptions of idealized periodic respiratory motion and better capturing patient-specific deformation patterns. However, such a combined approach introduces significant challenges. Most importantly, the number of free parameters increases dramatically, encompassing both motion fields and pharmacokinetic variables, which complicates the inverse problem and may require sophisticated regularization and optimization strategies to avoid instability and overfitting. Additionally, because these methods generally rely on iterative nonlinear optimization, computational costs can be substantial, often orders of magnitude higher than conventional reconstruction. This increased computational burden may present a considerable barrier to clinical translation.

Both of these problems can potentially be alleviated by deep learning–based methods. First, deep learning models are trained on collections of examples, allowing them to learn

complex priors of the output distribution. This implicitly regularizes the reconstruction problem and reduces reliance on handcrafted regularizers and sensitive initialisations that iterative methods require, helping deep learning-based approaches better deal with uncertainty and non-uniqueness of solutions inherent in highly ill-posed inverse problems [14]. In addition, convolutional neural networks act as implicit regularizers, learning to produce reconstructions that avoid physiologically implausible parameter values by exploiting patterns learned from training data [15]. While learning data priors might be undesirable in some situations (e.g., if the training distribution does not reflect the diversity of clinical cases), **Chapter 5** demonstrated that it is possible to quantify uncertainty associated with deep learning predictions, which provides a measure of confidence. Second, once a deep learning model is trained, inference is much faster than traditional iterative optimization, enabling near-real-time reconstruction and parameter estimation, which is more feasible in a clinical setting. This trend is supported by the increasing adoption of dedicated GPUs and AI-focused computing directly on MRI scanners by vendors.

6.2 Repeatability and reproducibility

The central barrier to the clinical adoption of quantitative DCE-MRI is the lack of repeatability and reproducibility. For example, measurements of the parameter K^{trans} obtained from the same healthy volunteer across multiple abdominal organs showed variations of up to 20% when the experiment was performed on two different 3 T MRI scanners [16]. The reproducibility of K^{trans} across four commercially available software packages in a cohort of 15 patients was limited, with within-subject coefficients of variation reaching up to 74% [17].

6.2.1 Scanner-induced variations

An important first step toward harmonizing DCE-MRI acquisition protocols has been taken by the Quantitative Imaging Biomarkers Alliance (QIBA), which has made substantial efforts to publish recommended guidelines for data acquisition [4]. However, even the recommended use of a 3D fast spoiled gradient echo sequence presents challenges, as this sequence is not fully standardized across vendors and can result in differences in measured signals. While the overall design of vendor-specific acquisition sequences can generally be accessed, limited flexibility for modification and the presence of proprietary or undocumented components hinder direct comparison and harmonization of protocols across platforms.

One way of approaching this challenge is through the use of vendor-neutral pulse sequence software such as gammaSTAR [18] and Pulseq [19]. These frameworks enable the programmatic description of MRI pulse sequences in a vendor-independent manner and allow their execution on compatible scanners. As an initial proof-of-concept, a simple dynamic 3D fast spoiled gradient-echo sequence could be implemented using these tools and evaluated in a multi-vendor study. At this stage, the administration of contrast agent is not required, as the primary objective would be to assess cross-vendor consistency of pre-contrast signal behaviour and derived baseline T_1 estimates.

A further step toward harmonization is ensuring that the measured contrast agent–induced relaxation accurately reflects the true contrast agent concentration. Because such validation requires known ground-truth parameters, MRI phantoms provide a practical and controlled experimental framework. While conventional static phantoms with reference T_1 values are useful for baseline calibration and validation, they are insufficient for validating the accuracy of contrast agent–induced signal changes in DCE-MRI, as they offer only a limited range of relaxation values and cannot replicate the dynamic signal changes observed during contrast uptake and washout. Thus, perfusion-specific phantoms are required. Initial studies using point-of-care perfusion phantoms have already demonstrated improved reproducibility of contrast enhancement curves and increased intraclass correlation coefficients of K^{trans} across different scanners [16, 20]. Recently, a more advanced phantom design has been proposed, which models capillary exchange using 3D printed porous gel structures [21]. Such phantoms enable the generation of DCE-MRI datasets with known pharmacokinetic parameters and in vivo–like contrast agent dynamics, thereby facilitating a more accurate assessment of the relationship between measured MRI signals and estimated quantitative parameters.

6.2.2 Software-induced variations

The choice of pharmacokinetic model also influences the repeatability and reproducibility of quantitative DCE-MRI parameters. On the one hand, simpler models generally provide more stable parameter estimation, are less prone to overfitting noise, and are less likely to converge to local minima during optimization [22]. On the other hand, more complex models that incorporate additional physiological compartments or exchange processes may be necessary to adequately describe tissues with heterogeneous vascular properties or complex contrast dynamics. However, the increased number of free parameters in these models leads to increased sensitivity to noise, a stronger dependence on the fitting strategy, which can negatively affect reproducibility. As proposed in [23], a pragmatic approach is to first restrict the space of candidate models using prior knowledge about tissue vascularity

and physiology, and then to select the most appropriate model based on both goodness-of-fit criteria and model complexity. Adopting such a strategy could support the development of general guidelines for pharmacokinetic model selection across healthy tissues and pathological conditions. This, in turn, would reduce unnecessary methodological variability and help prevent repeated ad hoc model selection when developing or evaluating new DCE-MRI acquisition, reconstruction, or analysis techniques. At the same time, standardization should not preclude methodological innovation when technological advances challenge existing model assumptions. For example, in **Chapter 2**, we proposed a correction to the extended Tofts-Kety model to account for modern non-Cartesian k-space sampling schemes used in DCE-MRI.

Even after selecting a specific pharmacokinetic model, quantitative DCE-MRI results may vary substantially due to differences in algorithmic implementations, optimization strategies, and numerical assumptions [24,25]. Variations in parameter initialization, constraints, temporal interpolation, noise handling, and convergence criteria can all lead to discrepancies in the estimated pharmacokinetic parameters. An important step toward harmonizing post-processing software has been initiated by the Open Science Initiative for Perfusion Imaging (OSIPI) [26], which established an open, community-driven repository for pharmacokinetic fitting algorithms that are systematically evaluated and benchmarked. While this resource enhances transparency and reproducibility, it stops short of providing consensus recommendations on which approaches to adopt in practice, highlighting the need for further work toward standardized guidelines. Although some degree of variability is likely to remain due to unavoidable modeling assumptions and methodological choices, further comparative and validation studies should aim to achieve acceptable levels of reproducibility that support reliable clinical interpretation and multi-center studies.

6.3 Clinical translation

In addition to technical validation and optimisation, quantitative DCE-MRI requires clinical validation before it can be adopted as a reliable imaging biomarker. A first step in this process is to establish a clear relationship between pharmacokinetic parameters and underlying biological processes and/or clinically relevant endpoints [27]. For example, in patients with resectable tumours, such biological validation can be achieved by comparing DCE-MRI parameters with histopathology data obtained from resected tumour specimens. Several studies have reported correlations between the volume transfer constant (K^{trans}) and microvessel density, a histological marker of tumour angiogenesis [28, 29].

However, while increased microvessel density is often associated with more aggressive disease and poorer prognosis [30], the exact biological meaning of K^{trans} and its direct relevance for clinical decision-making remain incompletely understood. Consequently, more comprehensive studies are required to generate and refine hypotheses regarding the role of DCE-MRI parameters in disease activity. In this context, preclinical studies may be particularly valuable, as they allow well-controlled experimental conditions and direct correlation with biological reference standards. The subsequent and essential step in clinical qualification is prospective validation, in which the proposed imaging biomarker is tested in well-designed clinical studies with sufficiently long follow-up periods (e.g., 3–5 years) [31]. These studies are necessary to determine whether the biomarker provides independent and clinically meaningful information about patient outcomes.

Although quantitative DCE-MRI has not yet crossed translational gaps, qualitative DCE-MRI is well established and routinely used in the clinic. Interestingly, in recent years, several MRI vendors have introduced non-Cartesian k-space acquisition trajectories, which are often used in quantitative DCE-MRI, into clinically deployed DCE sequences. For example, a radial stack-of-stars acquisition is implemented as StarVIBE on Siemens systems [32], 3D VANE XD on Philips systems [33], and LAVA Star on GE systems [34]. In clinical practice, these sequences are typically optimized for high SNR and motion robustness rather than for high temporal resolution. As a result, only a limited number of contrast phases (e.g., pre-contrast, arterial, and post-contrast phases) are reconstructed, with each phase incorporating a relatively large number of k-space readouts. If the raw k-space data underlying these acquisitions were stored in addition to the reconstructed images, such data could potentially be repurposed for quantitative MRI research. This reuse of clinically acquired data could reduce the financial burden of dedicated research scans, decrease the number of additional gadolinium administrations performed solely for research purposes, and enable larger studies through access to more extensive datasets.

6.4 Environmental impact

There are several environmental concerns associated with DCE-MRI. The first one is related to gadolinium-based contrast agents, which, after administration, are excreted in urine and enter municipal sewer systems. Wastewater treatment plants do not efficiently remove these compounds, leading to their discharge into surface waters such as rivers, lakes, and coastal ecosystems [35]. While geogenic gadolinium (from natural geological sources) is naturally present in surface waters, anthropogenic gadolinium from medical

imaging has also been detected in freshwater and even tap water [35]. In fact, among 32 European cities and 28 cities outside Europe, tap water samples from Amsterdam contained the highest concentrations of anthropogenic gadolinium [35]. The potential effects of this gadolinium on human health and aquatic fauna remain incompletely understood [36]. Injecting lower doses of contrast agent and reusing and sharing existing imaging data could be practical first steps to slow the rate at which gadolinium enters the environment. For example, in [37], it was demonstrated that full-dose–equivalent DCE-MRI images can be synthesized using deep-learning methods from acquisitions performed with a ten-fold reduced gadolinium dose.

Another concern relates to the substantial computational workloads associated with quantitative DCE-MRI, particularly when using iterative optimization and deep-learning approaches that require extensive experimentation and high-performance computing resources. In the Netherlands, data centers already consume a significant amount of electricity. In 2024, their energy use was estimated at 5,100 GWh, equivalent to the annual electricity consumption of approximately two million Dutch households [38]. This demand is expected to continue growing, driven in part by the rapid expansion of artificial intelligence technologies [39]. High electricity consumption contributes to CO_2 emissions and generates heat that requires additional water use for cooling systems, further increasing environmental impact. These issues extend to medical imaging research, where compute-intensive workflows, such as hyperparameter tuning, repeated model retraining, and large-scale optimization across multiple parameter settings, are increasingly common. While some of these efforts lead to meaningful methodological advances, others yield only marginal performance gains with limited clinical relevance. This pursuit of "harder, better, faster, stronger" performance can magnify the environmental cost of research without proportionate scientific benefit. By balancing scientific rigor with environmental stewardship, the research community can help ensure that progress in medical imaging remains both technologically innovative and environmentally responsible, minimizing unnecessary strain on energy systems and water resources.

Bibliography

[1] B. Endrich, H. S. Reinhold, J. F. Gross, and M. Intaglietta, "Tissue perfusion inhomogeneity during early tumor growth in rats," *Journal of the National Cancer Institute*, vol. 62, no. 2, pp. 387–395, 1979.

[2] G. Tozer, S. Lewis, A. Michalowski, and V. Aber, "The relationship between regional variations in blood flow and histology in a transplanted rat fibrosarcoma," *British journal of cancer*, vol. 61, no. 2, pp. 250–257, 1990.

[3] J. L. Evelhoch, "Key factors in the acquisition of contrast kinetic data for oncology," *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 10, no. 3, pp. 254–259, 1999.

[4] A. Shukla-Dave, N. A. Obuchowski, T. L. Chenevert, S. Jambawalikar, L. H. Schwartz, D. Malyarenko, W. Huang, S. M. Noworolski, R. J. Young, M. S. Shiroishi *et al.*, "Quantitative imaging biomarkers alliance (qiba) recommendations for improved precision of dwi and dce-mri derived biomarkers in multicenter oncology trials," *Journal of Magnetic Resonance Imaging*, vol. 49, no. 7, pp. e101–e121, 2019.

[5] E. Henderson, B. K. Rutt, and T.-Y. Lee, "Temporal sampling requirements for the tracer kinetics modeling of breast disease," *Magnetic resonance imaging*, vol. 16, no. 9, pp. 1057–1073, 1998.

[6] E. M. Schrauben, G. Lima da Cruz, C. W. Roy, and T. Küstner, "Motion mitigation techniques for abdominal and cardiac mr imaging," *Journal of Magnetic Resonance Imaging*, 2025.

[7] A. Melbourne, J. Hipwell, M. Modat, T. Mertzanidou, H. Huisman, S. Ourselin, and D. Hawkes, "The effect of motion correction on pharmacokinetic parameter estimation in dynamic-contrast-enhanced mri," *Physics in Medicine & Biology*, vol. 56, no. 24, p. 7693, 2011.

[8] J. Tokuda, H. Mamata, R. R. Gill, N. Hata, R. Kikinis, R. F. Padera Jr, R. E. Lenkinski, D. J. Sugarbaker, and H. Hatabu, "Impact of nonrigid motion correction technique on pixel-wise pharmacokinetic analysis of free-breathing pulmonary dy-

namic contrast-enhanced mr imaging," *Journal of Magnetic Resonance Imaging*, vol. 33, no. 4, pp. 968–973, 2011.

[9] L. Feng and H. Chandarana, "Accelerated abdominal mri: A review of current methods and applications," *Journal of Magnetic Resonance Imaging*, 2025.

[10] S. D. Serai and K. Darge, *Advanced clinical MRI of the kidney: methods and protocols*. Springer, 2023.

[11] E. Hodneland, A. Lundervold, J. Rørvik, and A. Z. Munthe-Kaas, "Normalized gradient fields for nonlinear motion correction of dce-mri time series," *Computerized Medical Imaging and Graphics*, vol. 38, no. 3, pp. 202–210, 2014.

[12] F. G. Zöllner, R. Sance, P. Rogelj, M. J. Ledesma-Carbayo, J. Rørvik, A. Santos, and A. Lundervold, "Assessment of 3d dce-mri of the kidneys using non-rigid image registration and segmentation of voxel time courses," *Computerized Medical Imaging and Graphics*, vol. 33, no. 3, pp. 171–181, 2009.

[13] S. Tirunagari, N. Poh, K. Wells, M. Bober, I. Gorden, and D. Windridge, "Movement correction in dce-mri through windowed and reconstruction dynamic mode decomposition," *Machine Vision and Applications*, vol. 28, no. 3, pp. 393–407, 2017.

[14] C. M. Hyun, S. H. Baek, M. Lee, S. M. Lee, and J. K. Seo, "Deep learning-based solvability of underdetermined inverse problems in medical imaging," *Medical Image Analysis*, vol. 69, p. 101967, 2021.

[15] W. Bian and Y. K. Tamilselvam, "A review of optimization-based deep learning models for mri reconstruction," *AppliedMath*, vol. 4, no. 3, pp. 1098–1127, 2024.

[16] H. Kim, M. Mousa, P. Schexnailder, R. Hergenrother, M. Bolding, B. Ntsikoussalabongui, V. Thomas, and D. E. Morgan, "Portable perfusion phantom for quantitative dce-mri of the abdomen," *Medical physics*, vol. 44, no. 10, pp. 5198–5209, 2017.

[17] T. Heye, M. S. Davenport, J. J. Horvath, S. Feuerlein, S. R. Breault, M. R. Bashir, E. M. Merkle, and D. T. Boll, "Reproducibility of dynamic contrast-enhanced mr imaging. part i. perfusion characteristics in the female pelvis by using multiple computer-aided diagnosis perfusion analysis solutions," *Radiology*, vol. 266, no. 3, pp. 801–811, 2013.

[18] S. Konstandin, M. Günther, and D. C. Hoinkiss, "gammastar: A framework for the development of dynamic, real-time capable mr sequences," *Magnetic Resonance in Medicine*, 2025.

[19] K. J. Layton, S. Kroboth, F. Jia, S. Littin, H. Yu, J. Leupold, J.-F. Nielsen, T. Stöcker, and M. Zaitsev, "Pulseq: a rapid and hardware-independent pulse sequence prototyping framework," *Magnetic resonance in medicine*, vol. 77, no. 4, pp. 1544–1552, 2017.

[20] M. D. Holland, A. Morales, S. Simmons, B. Smith, S. R. Misko, X. Jiang, D. A. Hormuth, C. Christenson, R. P. Koomullil, D. E. Morgan *et al.*, "Disposable point-of-care portable perfusion phantom for quantitative dce-mri," *Medical physics*, vol. 49, no. 1, pp. 271–281, 2022.

[21] M. S. Sarwar, A. Vallatos, C. H. Lau, A. Waldman, S. Dimartino, and M. J. Thrippleton, "3d-printable phantoms for quantitative dynamic contrast-enhanced mri," *Magnetic Resonance in Medicine*, 2025.

[22] M. E. Gaens, W. H. Backes, S. Rozel, M. Lipperts, S. N. Sanders, K. Jaspers, J. P. Cleutjens, J. C. Sluimer, S. Heeneman, M. J. Daemen *et al.*, "Dynamic contrast-enhanced mr imaging of carotid atherosclerotic plaque: model selection, reproducibility, and validation," *Radiology*, vol. 266, no. 1, pp. 271–279, 2013.

[23] S. P. Sourbron and D. L. Buckley, "Classic models for dynamic contrast-enhanced mri," *NMR in Biomedicine*, vol. 26, no. 8, pp. 1004–1027, 2013.

[24] "A multi-institutional comparison of dynamic contrast-enhanced magnetic resonance imaging parameter calculations," *Scientific reports*, vol. 7, no. 1, p. 11185, 2017.

[25] L. Beuzit, P.-A. Eliat, V. Brun, J.-C. Ferré, Y. Gandon, E. Bannier, and H. Saint-Jalmes, "Dynamic contrast-enhanced mri: Study of inter-software accuracy and reproducibility using simulated and clinical data," *Journal of magnetic resonance imaging*, vol. 43, no. 6, pp. 1288–1300, 2016.

[26] P. J. van Houdt, S. Ragunathan, M. Berks, Z. Ahmed, L. E. Kershaw, O. J. Gurney-Champion, S. Tadimalla, J. Arvidsson, Y. Sun, J. Kallehauge *et al.*, "Contrast-agent-based perfusion mri code repository and testing framework: Ismrm open science initiative for perfusion imaging (osipi)," *Magnetic Resonance in Medicine*, vol. 91, no. 5, pp. 1774–1786, 2024.

[27] D. Hernando, "Development, validation, qualification and dissemination of quantitative mr methods: Overview and recommendations by the ismrm quantitative mr study group."

[28] S. H. Kim, H. S. Lee, B. J. Kang, B. J. Song, H.-B. Kim, H. Lee, M.-S. Jin, and A. Lee, "Dynamic contrast-enhanced mri perfusion parameters as imaging biomarkers of angiogenesis," *PLoS One*, vol. 11, no. 12, p. e0168632, 2016.

[29] N. Karabay, H. M. Bülbül, E. Doğan, A. Ö. İkiz, G. Bülbül, and S. Sarioğlu, "The correlations between dynamic contrast enhanced magnetic resonance imaging and immunohistochemical data in head and neck squamous cell carcinomas," *Turkish journal of medical sciences*, vol. 52, no. 6, pp. 1950–1957, 2022.

[30] J. Zhang, L. Zhang, Q. Lin, W. Ren, and G. Xu, "Prognostic value of endoglin-assessed microvessel density in cancer patients: A systematic review and meta-analysis," *Oncotarget*, vol. 9, no. 7, p. 7660, 2017.

[31] J. P. O'connor, E. O. Aboagye, J. E. Adams, H. J. Aerts, S. F. Barrington, A. J. Beer, R. Boellaard, S. E. Bohndiek, M. Brady, G. Brown *et al.*, "Imaging biomarker roadmap for cancer studies," *Nature reviews Clinical oncology*, vol. 14, no. 3, pp. 169–186, 2017.

[32] N. Jin, "Breathe easy with freezeit—new imaging techniques to improve body mr imaging."

[33] T. Kaga, Y. Noda, N. Kawai, K. Kajita, Y. Ueda, M. Honda, F. Hyodo, H. Kato, and M. Matsuo, "Free-breathing gadoxetic acid-enhanced hepatobiliary phase imaging using stack-of-stars radial sampling and compressed sense."

[34] S. Ichikawa, U. Motosugi, M.-L. Kromrey, D. Tamada, T. Wakayama, K. Wang, T. A. Cashen, A. Ersoz, and H. Onishi, "Utility of stack-of-stars acquisition for hepatobiliary phase imaging without breath-holding," *Magnetic Resonance in Medical Sciences*, vol. 19, no. 2, pp. 99–107, 2020.

[35] J.-A. Barrat, G. Bayon, R. Tripier, M.-L. Rouget, Y. Germain, and D. B. Salem, "Rare earth element abundances and gadolinium contamination in tap water worldwide," *Chemosphere*, vol. 394, p. 144812, 2026.

[36] L. A. Souza, R. M. Pedreira, M. Miró, and V. Hatje, "Evidence of high bioaccessibility of gadolinium-contrast agents in natural waters after human oral uptake," *Science of The Total Environment*, vol. 793, p. 148506, 2021.

[37] E. Gong, J. M. Pauly, M. Wintermark, and G. Zaharchuk, "Deep learning enables reduced gadolinium dose for contrast-enhanced brain mri," *Journal of magnetic resonance imaging*, vol. 48, no. 2, pp. 330–340, 2018.

[38] C. B. voor de Statistiek, "Datacenters verbruiken 4,6 procent van de elektriciteit," Dec 2025. [Online]. Available: https://www.cbs.nl/nl-nl/nieuws/2025/51/datacenters-verbruiken-4-6-procent-van-de-elektriciteit?utm_campaign=data-centres-in-the-netherlands-now-use-as-much-power-as-2-million-homes&utm_medium=referral&utm_source=www.dutchbrief.com

[39] Dec 2025. [Online]. Available: https://www.dutchdatacenters.nl/cijfers/duurzaamheid-energie/

Summary

Dynamic contrast-enhanced MRI can quantify tissue perfusion and vascular permeability, offering information that can support disease characterization and guide personalized treatment strategies. This thesis focuses on technical developments in quantitative DCE-MRI, in particular in image reconstruction, modelling, and post-processing, with the aim of improving the accuracy, robustness, and clinical feasibility of quantitative parameter estimation.

CHAPTER 2 investigates the impact of time-averaged signals in DCE-MRI acquisitions that oversample the center of k-space, such as spiral, radial, and PROPELLER sequences, on the accuracy of pharmacokinetic modeling. A modified extended Tofts model was developed to account for the temporal averaging inherent to these acquisitions and was evaluated using Monte Carlo simulations and in vivo pseudo-spiral variable-density MRI scans of pancreatic cancer patients reconstructed at multiple temporal resolutions. The proposed model demonstrated improved accuracy and precision compared with conventional pharmacokinetic modeling approaches, particularly for longer acquisition times, while also recovering information related to high temporal signal behavior. In vivo validation further showed more stable estimates of pharmacokinetic parameters across sampling frequencies, supporting the use of the corrected model for variable-density dynamic contrast-enhanced MRI acquisitions, especially at temporal resolutions of 8 seconds or longer.

CHAPTER 3 presents a respiratory motion–corrected reconstruction framework for free-breathing abdominal DCE-MRI aimed at improving image quality and the reliability of quantitative pharmacokinetic analysis. The proposed method retrospectively sorts free-breathing pseudo-spiral DCE-MRI data into different respiratory states, estimates respiratory motion using deformable image registration, and incorporates the resulting motion information directly into a compressed sensing reconstruction framework with weighted view sharing. Evaluation in abdominal DCE-MRI data from four patients demonstrated substantial improvements in image sharpness and anatomical delineation

compared with conventional compressed sensing reconstruction, while preserving temporal enhancement behavior and quantitative pharmacokinetic parameter estimates. The results indicate that integrating respiratory motion correction directly into image reconstruction enhances both spatial and temporal fidelity without introducing measurable bias into quantitative analysis, supporting its application in free-breathing abdominal DCE-MRI for clinical and research purposes.

CHAPTER 4 presents a model-based reconstruction framework for quantitative DCE-MRI that directly estimates pharmacokinetic parameters from raw k-space data, thereby avoiding intermediate image reconstruction and potentially reducing the trade-off between temporal resolution and spatial resolution. The framework integrates the extended Tofts pharmacokinetic model into the PyQMRI platform, which was then evaluated using both simulated abdominal phantom data and in vivo liver DCE-MRI data from a healthy volunteer. Compared with a conventional image-based non-linear least squares fitting approach, the proposed MBR method produced more precise pharmacokinetic maps while maintaining similar or improved accuracy, particularly enhancing the delineation of fine anatomical structures such as blood vessels. These improvements remained consistent across a wide range of temporal resolutions and were especially pronounced at the shortest tested frame rate. Although in vivo gains were less substantial, likely due to physiological complexity and motion-related degradation, the study demonstrates that direct parameter estimation from k-space data is a promising strategy for improving the robustness and accuracy of quantitative DCE-MRI under challenging imaging conditions.

CHAPTER 5 investigates uncertainty estimation in quantitative DCE-MRI, where accurate perfusion analysis is hindered by noisy data and complex pharmacokinetic modeling. The chapter introduces an ensemble of mean-variance estimation (MVE) neural networks designed to estimate perfusion parameters together with aleatoric (data-related) and epistemic (model-related) uncertainties. The proposed approach was compared with conventional non-linear least squares (NLLS) fitting and a physics-informed neural network (PINN), both relying on covariance matrix–based uncertainty estimation. Simulation experiments demonstrated that MVE achieved the highest accuracy for both parameter estimation and uncertainty quantification, with aleatoric uncertainty closely matching true reconstruction errors, while NLLS and PINN generally overestimated uncertainty. Both MVE and PINN ensembles also showed increased epistemic uncertainty for out-of-distribution data, indicating sensitivity to unfamiliar inputs. In vivo liver DCE-MRI experiments further showed that MVE generated smoother and more reliable uncertainty maps, whereas NLLS and PINN produced noisier maps with substantial outliers. Overall, the chapter demonstrates that MVE-based uncertainty estimation can substantially

improve the reliability and interpretability of AI-driven quantitative DCE-MRI analysis, supporting its potential clinical translation.

CHAPTER 6 provides a general discussion of the main challenges and future directions in quantitative DCE-MRI, building on the work presented in this thesis. It places the technical developments in a broader context and reflects on the requirements for clinical translation, including limitations in image quality, sources of variability, and issues related to reproducibility and standardization. The chapter also discusses methodological directions for improving reconstruction, modelling, and analysis approaches, as well as the potential role of emerging techniques such as deep learning. In addition, it considers the steps needed for clinical validation and adoption, and briefly reflects on wider practical and environmental considerations associated with quantitative DCE-MRI.

Nederlandse samenvatting

dynamic contrast-enhanced MRI (DCE-MRI) kan weefselperfusie en vasculaire permeabiliteit kwantificeren, wat informatie oplevert die kan bijdragen aan de karakterisering van ziekten en het ontwikkelen van gepersonaliseerde behandelstrategieën. Deze thesis richt zich op technische ontwikkelingen in kwantitatieve DCE-MRI, met name op het gebied van beeldreconstructie, modellering en nabewerking, met als doel de nauwkeurigheid, robuustheid en klinische toepasbaarheid van kwantitatieve parameterbepaling te verbeteren.

Hoofdstuk 2 onderzoekt de impact van tijdsgemiddelde signalen in DCE-MRI-acquisities die het centrum van de k-space oversampled, zoals spiraal-, radiaal- en propellervormige sequenties, op de nauwkeurigheid van farmacokinetische modellering. Er werd een aangepast extended Tofts-model ontwikkeld om rekening te houden met de temporele middeling die inherent is aan deze acquisities. Dit model werd geëvalueerd met Monte Carlo-simulaties en in vivo MRI-scans van patiënten met alvleesklierkanker, verkregen met een pseudo-spiraalvormige acquisitie met variabele dichtheid en gereconstrueerd op meerdere temporele resoluties. Het voorgestelde model vertoonde een verbeterde nauwkeurigheid en precisie in vergelijking met conventionele farmacokinetische modelleringsmethoden, met name bij langere acquisitietijden, en herstelde tevens informatie met betrekking tot hoogfrequent temporeel signaalgedrag. In vivo-validatie toonde bovendien stabielere schattingen van farmacokinetische parameters over verschillende sampling frequenties, wat het gebruik van het gecorrigeerde model ondersteunt voor DCE-MRI acquisities met variabele dichtheid k-space sampling, vooral bij temporele resoluties van 8 seconden of langer.

In **Hoofdstuk 3** wordt een reconstructiekader met correctie voor ademhalingsbewegingen voor DCE-MRI van de buik tijdens vrije ademhaling gepresenteerd. Dit hoofdstuk is gericht op het verbeteren van de beeldkwaliteit en de betrouwbaarheid van kwantitatieve farmacokinetische analyses. De voorgestelde methode sorteert retrospectief pseudo-spiraalvormige DCE-MRI data tijdens vrije ademhaling in verschillende

ademhalingstoestanden, schat de ademhalingsbeweging met behulp van deformeerbare beeldregistratie af en integreert de resulterende bewegingsinformatie direct in een reconstructiekader met compressed sensing en weighted view sharing. Evaluatie van abdominale DCE-MRI data van vier patiënten toonde aanzienlijke verbeteringen in beeldscherpte en anatomische afbakening in vergelijking met conventionele reconstructie met compressed sensing, terwijl het temporele contrastopnamegedrag en de kwantitatieve schattingen van farmacokinetische parameters behouden bleven. De resultaten geven aan dat het direct integreren van ademhalingsbewegingscorrectie in de beeldreconstructie zowel de spatiale als temporele nauwkeurigheid verbetert zonder meetbare vertekening in de kwantitatieve analyse te introduceren.

Hoofdstuk 4 presenteert een modelgebaseerd reconstructiekader (MBR) voor kwantitatieve DCE-MRI dat farmacokinetische parameters direct schat uit ruwe k-space data, waardoor tussentijdse beeldreconstructie wordt vermeden en de afweging tussen temporele en spatiale resolutie mogelijk kan worden verminderd. Het kader integreert het uitgebreide Tofts-farmacokinetische model in het PyQMRI-platform, dat vervolgens werd geëvalueerd met behulp van zowel gesimuleerde abdominale fantoom data als in vivo lever-DCE-MRI data van een gezonde vrijwilliger. Vergeleken met een conventionele methode, welke is gebaseerd op non-linear least-squared (NLLS) fit, produceerde de voorgestelde MBR-methode nauwkeurigere farmacokinetische maps met behoud van een vergelijkbare of verbeterde nauwkeurigheid, waarbij met name de afbakening van fijne anatomische structuren zoals bloedvaten werd verbeterd. Deze verbeteringen bleven consistent over een breed scala aan temporele resoluties en waren vooral uitgesproken bij de kortste geteste framesnelheid. Hoewel de in vivo-winsten minder substantieel waren, waarschijnlijk vanwege fysiologische complexiteit en bewegingsgerelateerde degradatie, toont de studie aan dat directe parameter afschatting uit k-space data een veelbelovende strategie is om de robuustheid en nauwkeurigheid van kwantitatieve DCE-MRI te verbeteren tijdens lastige acquisitieomstandigheden.

Hoofdstuk 5 onderzoekt onzekerheidsmarge in kwantitatieve DCE-MRI, waar nauwkeurige perfusieanalyse wordt belemmerd door ruis in de data en complexe farmacokinetische modellering. Dit hoofdstuk introduceert een ensemble van neurale netwerken voor mean-variance estimation (MVE), ontworpen om perfusieparameters af te schatten, samen met aleatorische (data-gerelateerde) en epistemische (model-gerelateerde) onzekerheden. De voorgestelde methode werd vergeleken met conventionele NLLS fit en een physics-informed neuraal netwerk (PINN), die beide gebruikmakten van onzekerheidsmarge op basis van covariantiematrices. Simulatie-experimenten toonden aan dat MVE de hoogste nauwkeurigheid behaalde voor zowel

parameterschatting als onzekerheidskwantificatie, waarbij de aleatorische onzekerheid nauw aansloot bij de werkelijke reconstructiefouten, terwijl NLLS en PINN de onzekerheid over het algemeen overschatten. Zowel de MVE- als de PINN-ensembles vertoonden ook een verhoogde epistemische onzekerheid voor out-of-distribution-data, wat wijst op gevoeligheid voor onbekende input.

Hoofdstuk 6 biedt een algemene beschouwing van de belangrijkste uitdagingen en toekomstige richtingen binnen kwantitatieve DCE-MRI, voortbouwend op het werk dat in deze thesis wordt gepresenteerd. Het plaatst de technische ontwikkelingen in een bredere context en reflecteert op de vereisten voor klinische implementatie, waaronder beperkingen in beeldkwaliteit, bronnen van variabiliteit en vraagstukken rond reproduceerbaarheid en standaardisatie. Het hoofdstuk bespreekt daarnaast methodologische richtingen voor het verbeteren van reconstructie-, modellerings- en analysemethoden, evenals de mogelijke rol van opkomende technieken zoals deep learning. Bovendien worden de stappen richting klinische validatie en implementatie behandeld, en wordt kort gereflecteerd op bredere praktische en omgevingsgerelateerde aspecten van kwantitatieve DCE-MRI.

Portfolio

Name PhD student	Natalia Korobova
PhD period	01.11.2021 – 01.01.2026
PhD supervisors	Dr. O. J. Gurney-Champion
	Prof. Dr. Ir. A. J. Nederveen

Activity	**Year**	**ECTS**
1. PhD training		
Courses		
Radiation Oncology	2022	1.5
In vivo NMR	2022	1.4
3T MRI scan qualification	2022	0.2
NT2 beginners 1+2	2022	5.0
Ethics and Integrity	2023	2.0
Practical Biostatistics	2024	1.4
Be your best journey	2025	0.8
Seminars, workshops, and master classes		
Weekly MR physics seminars	2021–2025	4.0
Weekly abdominal seminars	2021–2025	4.0
Monthly Bernstein book club (co-organiser)	2022–2024	2.0
Oral presentations		
ISMRM = International Society for Magnetic Resonance in Medicine		
ESMRM = European Society for Magnetic Resonance in Medicine and Biology		
ISMRM annual meeting, Toronto	2023	1.5
ESMRM annual meeting, Basel	2023	1.5
ISMRM Benelux chapter, 's Hertogenbosch	2024	1.5

Activity	Year	ECTS
ISMRM Benelux chapter, Hilversum	2025	1.5
(Inter)national conferences		
ISMRM Benelux chapter, Maastricht	2022	0.2
ISMRM Benelux chapter (co-organizer), Brussels	2023	2.0
ISMRM motion correction workshop, Quebec	2024	0.2
ISMRM Benelux chapter, Antwerp	2026	0.2
Other		
Cancer Center Amsterdam retreat, Noordwijkerhout	2022	0.2
RNG Research day, Amsterdam	2023	0.2
Cancer Center Amsterdam retreat, Noordwijkerhout	2024	0.2
RNG Research day, Amsterdam	2025	0.2
2. Teaching		
Courses		
Medische beelden en signalen	2022	0.5
Matlab course	2022	0.5
Matlab course	2023	0.5
Advanced MRI	2024	0.5
Supervising		
Max Bongers, Master's thesis	2023	1.0
Jonas van Elburg, Master's thesis	2024	1.0
3. Parameters of Esteem		
Awards and Prizes		
ESMRMB best abstract award	2023	

Acknowledgements

First of all, dear **committee**, thank you for the time and effort you have devoted to evaluating and commenting on this thesis.

My promoters, **Oliver** and **Aart**, I am deeply grateful for your support and guidance throughout this journey. **Oliver**, you have always been there whenever I needed you, ready with thoughtful feedback, reassuring words at the right moments, and always able to make time for spontaneous discussions despite your busy schedule. Thank you for being such an inspiring example, not only as an excellent researcher, but even more importantly, as a genuinely kind and thoughtful person. Dear **Aart**, thank you for your guidance during my PhD. Although we met less frequently, your perspective often helped me see the bigger picture and reflect on my progress more clearly. I also very much enjoyed our conversations during Z0-uitjes.

My paranymphs, **Reagan**, **Diogo** and **Wilhelm**, I am happy to have gotten to know you! And to have you by my side on the day of the defence. **Reagan**, our chats about everything, serious or silly, made my PhD journey so much lighter. From enjoying the summer warmth in Quebec to surviving that very alternative underground theatre show, and from sangria spills to endless laughter, every moment with you was easy and fun. I can't wait to see where our friendship goes next. **Diogo**, thank you for being such a good listener when I needed this. And thank you for bringing your positive (and wonderfully wild) energy to Z0. The office would not have been the same without you. **Wilhelm**, I'm really glad we got to know each other better over time, especially during my last year at AMC. Whether we were puzzling over RF pulses, untangling the mysteries of Dutch word order, or going out for a drink, every moment was a treasure. Thank you!

To the Big **Z0+ family**, I will miss you all! This has been a wonderful journey and I am very thankful for it! A special thank you to all my colleagues who made life at Z0 so vibrant: the **BernsteinBookClub** (oh boy...), **NederlandseTaalClub**, **Padelmeisters**, **KerstBorrelBazen**, and the ladies who "Dutchified" me with klompen. You are all amazing and brought so much joy to my days. To my fellow PhDers (or almost-PhDs). **Koen** and **Sophie**, thank you for the unforgettable dance floor moments and a much-needed reminder that PhD life isn't all about papers and deadlines! **Zarah**, **Daphne**, **Janine**, **Eline**, **Anne Marijn**, your pos-

itivity, energy, and smiles made even the toughest days brighter. It meant a lot to have you as colleagues. **Nienke**, you're the kind of colleague who makes work feel less like work. Everyone feels at ease around you, and your ability to listen and support is a rare gift. It's clear how much you care, and it hasn't gone unnoticed. **Isa**, our coffee breaks and side chats were the perfect escape from the daily grind. **Renske** and **Daan B**, thank you for the many gezellige drinks, where we could relax, unwind, and talk about absolutely everything. I'm so glad I got to work with you. **Bobby**, de Bobster, you set the bar high for what it means to be a PhD student. Your discipline and hard work have inspired so many of us, me included. **Elon**, those co-working sessions in Z0-174, fueled by sweet treats and mutual determination, were a highlight. Your humour made the process not just bearable, but enjoyable. **Susi**, we spent so many meetings together figuring out how model-based reconstruction works. Thank you for all the brainstorming and for cheering me up in difficult moments. **Luc**, thank you for all the meme knowledge you shared with me. Thanks also for teaching me the basics of padel. My shots may not always land where I aim, but at least I now know how to backhand and forehand. Ты даже практиковал со мной русский язык, это было очень неожиданно и приятно! With you at Z0, I do not doubt that the future of MRI research is in excellent hands. **Slava**, хоть мы и познакомились в самом конце моего PhD, твоё чувство юмора и лёгкость обернули наше общение в радость. Желаю тебе увлекательного пути в AMC! The abdominal group, **Daan K**, **Myrte**, **Dilara**, **Eva** and **Benthe**, thank you for the insightful discussions and the lively debates that pushed our work forward. **Anne**, **Daisy**, **Fleur**, **Hugo**, **Kim**, **Liza**, **Luuk**, **Marieke**, **Roos**, thank you for the lunches that were rarely just lunches, often featuring a quiz and lively discussions. Thank you for many fun uitjes, the endless laughter, and for making sure there was always something to look forward to beyond experiments and deadlines. You made the PhD journey considerably more enjoyable (and significantly less serious). **Rick** and **Thijs**, even though we met late in my PhD, I enjoyed your company very much and was curious to learn more about the exciting research you're so passionate about. Here's to the future, may it be as bright as your enthusiasm. Senior researchers, **Anouk**, **Bram**, **Melissa**, **Ot**, **Pim**, **Sofieke**, thank you for your insightful questions during MR physics meetings, your invaluable input into this work, and, of course, for the endless fun during lunches, train rides to Utrecht, conferences and an endless supply of "coffee". **Eric**, a special thank you for your supervision on the motion chapter and for the many, many insightful advices during abdominal meetings. I have always admired your hands-on expertise with MRI and your ability to find practical solutions to almost any problem. Beyond work, I have also thoroughly enjoyed our many social adventures, the Double Dutch comedy night, the trip to Quebec, padel games, and an uncountable number of borrels. Thanks for making both the science and everything

around it so much more enjoyable. The 3T MRI team, **Paul**, thank you for your patience, especially with my numerous requests about the Luna server. Your help was invaluable! **Sandra**, thank you for gently introducing me to the MRI scanner. **Alexandra** and **Katinka**, thank you for keeping the growing group of researchers running smoothly.

Candido, although our time as colleagues has been brief, I'm so grateful our paths crossed. Your easygoing nature made every conversation, whether in the office or over a drink, a genuine delight.

Jonas, it has been a great pleasure working with you. You arrived at AMC as a bright and eager student, curious to explore the world of MRI. In this short time, you have demonstrated remarkable autonomy, proactivity, and a thirst for knowledge. I consider myself truly lucky to have had such an outstanding student like you. I hope your curiosity in life will always be with you.

Wonderful people I got to know while organizing ISMRM Benelux chapter, **Danie**, **Daniëlle**, **Donnie**, **Emma**, **Ǩasa**, **Sam**, **Sarah**, **Theresa**, thank you for the laughter, the shared stories, and the good times. Every reunion feels like a burst of sunshine. Your energy and camaraderie made the journey unforgettable!

My besties, **Inigo**, **Cristina**, **Stephanos**, and **Jarry**, we've known each other since almost day one of our master's. We've been there for each other through exams, corona, changing rooms and apartments, travelling, and now PhDs. I'm so glad to have a group of friends where we can all just be ourselves, share so much, and always find support.

Катя и **Антон**, я рада, что смогу разделить этот день с вами, с тем, с кем когда-то начинала учиться.

Paint and sip (aka paint and panic) group, **Djennifer**, **Giulia**, **Charlotte**, and **Cristina**, our painting sessions have been my secret weapon against stress, and I'm so glad we turned this idea into a tradition. Your company and creativity made every brushstroke (and occasional mishap) a joy. I am looking forward to many more masterpieces, and the stories behind them!

Cynical scientists at AVL-NKI, **Celia**, **Georgios**, **Mar**, **Osman**, **Roque**, **Stijn**, although we only got acquainted while I was polishing the final chapters of this dissertation, you made me feel welcome from day one. Thank you for giving me such a warm welcome and making the transition to the new place so smooth. I am very much looking forward to working with all of you, and, hopefully, to everything that happens outside the office as well. **Stijn**, a special thank you for making the Nederlandse Samenvatting actually readable.

Dunya, you've been the purr-fect companion I never knew I needed. Your cuteness is so overwhelming that even Mike and I can't resist, though, let's be honest, neither can anyone else. With your love, even the biggest problems shrink to the size of a mouse.

Oma en **Opa van den Donker** bedankt voor alle gezellige familiefeesten in het Chinese restaurant en voor de fijne momenten die ik met jullie heb mogen meemaken.

Nicole and **Jerry**, ik ben dankbaar dat jullie deel uitmaken van mijn leven en dat we zoveel bijzondere familiemomenten met elkaar mogen delen. Ik hoop dat we dit nog lang zo kunnen blijven doen.

Lieve **Tineke** en **Ronald**, ik wil jullie van harte danken voor het warme welkom in jullie familie. Jullie hebben me vanaf het begin laten voelen dat ik erbij hoor. Of het nu ging om het wennen aan een nieuw land, het geduldig leren van de taal, of gewoon het delen van mooie en uitdagende momenten, jullie stonden altijd voor me klaar. Ik had me geen betere schoonouders kunnen wensen. Dankzij jullie voelt deze reis als thuis, en ik ben jullie oneindig dankbaar voor jullie onvoorwaardelijke steun en vriendelijkheid.

Дина, Оскар, Булат, Володя, Ира, Паша, Андрюша, хорошо, что мы есть друг у друга. Я знаю, что всегда могу положиться на вас в сложный момент, а вы на меня.

Бабуля, дедуля, дәү әни, дәү әти, я знаю, что вы всегда рядом со мной и очень гордитесь.

Софа, у меня не могло быть лучшей сестрёнки! Хоть и невербально, но ты каждый раз даешь мне знать, как сильно ты меня любишь и ждешь в гости. Это даёт мне больше сил, чем любые слова.

Мама и папа, ваш вкад в эту диссертацию сложно переоценить. **Мама**, отношения между мамой и дочкой это сплошная путаница. Но я хочу, чтобы ты знала, что эта работа и её защита стали возможны благодаря тебе. С детства я знала, что девочки могут всё. Что у них есть право учиться, что они могут совмещать работу с семьёй, и что они могут выбирать свою дорогу. Мне повезло, этому меня научила ты. Я бы хотела, чтобы у каждой девочки была такая мама, которая показала бы ей, что она способна на всё. **Папа**, я с детства восхищалась тобой. Все знают тебя как профессионала, доктора от Бога. Мне повезло видеть, как коллеги ценят твои знания, а пациенты благодарны за помощь. Но ты не только отличный врач, ты просто хороший человек. Спасибо, что всегда был рядом, рисовал за меня картинки для уроков ИЗО, шил костюмы для школьных спектаклей и поддерживал во всём.

Mike, you have been there for me every step of the way during this PhD and beyond. In good times and bad, you find exactly the right words that reach me. You have taught me so much about people, myself, work, and life in general. You have also developed an impressive tolerance for conversations about things that probably seemed fascinating only to me. I am grateful for this, thank you! I am happy we are on this journey together and I am looking forward to the new adventures ahead. I love you ♡



Curriculum Vitae



Natalia Vladimirovna Korobova was born in Kazan, Russia, on November 3rd 1995 as the first child of Vladimir Viktorovich Korobov and Liliya Rafgatovna Safiullina. She grew up in Kazan with her sister Sofya. After graduating from high school in 2013, where she developed a passion for mathematics and first learned to program in Pascal, she moved to Moscow to pursue a Bachelor's degree in Applied Mathematics at the Higher School of Economics.

She graduated cum laude in 2017 and continued her academic journey, earning a Master's degree in Applied Mathematics (cum laude) in 2019. During her studies, Natalia became deeply interested in interdisciplinary research, particularly the intersection of technical fields like mathematics with natural sciences such as biology. This curiosity led her to pursue a second Master's degree in Medical Imaging at Utrecht University, where she was admitted to the program and awarded the Utrecht Excellence Scholarship for non-EU students. During her studies, she discovered the principles of medical image acquisition and processing, as well as gained insight into upcoming developments in the field. She graduated from Utrecht University in 2021.

Inspired by the engineering behind imaging modalities and the potential of advanced algorithms for image processing, Natalia began her PhD in 2021 at Amsterdam University Medical Center, under the supervision of Oliver Gurney-Champion and Aart Nederveen. Her research focused on developing new reconstruction and processing techniques for quantitative dynamic contrast-enhanced magnetic resonance imaging, with the potential to contribute to future applications such as personalized cancer treatment planning. During her PhD, she supervised students, assisted in teaching courses, and presented her work at multiple (inter)national conferences.

Currently, Natalia works as a Clinical Scientist in the Department of Radiotherapy at the Netherlands Cancer Institute. She lives in Houten with her partner Mike and their cat Dunya.

