Politecnico di Torino

Biomedical Engineering

A.Y. 2025/2026

Graduation Session July 2026

Quantitative analysis of contrast-enhanced ultrasound images in patients with liver metastases

Supervisors:

Prof. Kristen Mariko Meiburger
Dr. Catarina Dinis Fernandes
Dr. Roel Deckers

Candidate:

Chiara Bonizzoni

Abstract

Metastases are the most common malignancies in the liver, affecting 25% - 50% of patients with non-hematological cancers. They are formed by angiogenesis driven by the tumor, leading to an irregular and disorganized vascular network.

In this context, contrast-enhanced ultrasound (CEUS) is a useful imaging method that allows the visualization of the microvasculature in real time. Contrast agents administrated intravenously enhance lesional detection and characterization of the vascular pattern. Metastatic liver lesions on sequences obtained from CEUS are usually hypoechoic and present rim enhancement. CEUS data are typically analysed through qualitative and semi-quantitative approaches. Time-intensity curve (TICs) analysis can be used to provide semi-quantitative parameters such as peak and mean transit times, commonly used for lesion characterization and treatment monitoring. However, enhancement spatial heterogeneity is still mainly assessed visually, making it subjective and operator-dependent.

To address this limitation, this work investigates contrast-ultrasound dispersion imaging (CUDI) as a quantitative framework for spatial characterization of enhancement heterogeneity. Pixel-wise TICs are modeled using a convective dispersion model that describes the contrast dynamics, while local spatiotemporal behavior is evaluated by comparing each pixel with its neighborhood in time and frequency. Spatial similarity metrics, such as coherence and correlation, are inversely related to dispersion, and they are a surrogate of vascular organization and angiogenic activity. While previously applied to other cancer types, this framework has not yet been employed in liver metastases and may enable a detailed characterization of lesion microvasculature, supporting personalized treatment strategies.

One of the most important requirement for this analysis is that the levels of intensity in the pixels must be consistent across time. However, in liver imaging, this is usually affected by motion. Therefore, motion correction (MC) was performed to guarantee reliable temporal alignment of CEUS sequences before CUDI analysis. The parameters of motion correction and CUDI-derived parameters were then statistically analyzed using the Wilcoxon signed-rank test.

This study included multiple CEUS sequences from 11 patients with colorectal liver metastases acquired at UMC Utrecht. No MC strategies were applied during data acquisition. Postprocessing MC approaches were explored and the best-performing in terms of similarity metrics was integrated prior to the CUDI analysis. The final MC pipeline consisted of a breath-wise two-tier rigid registration followed by a deformable registration improving image alignment, with increases in normalized cross-correlation (0.64→0.76), mutual information (0.43→0.69) and structural similarity (0.36→0.48) ($p < 0.001$).

To study spatial variations in perfusion and structure, extracted CUDI parameters were evaluated using a radial analysis centered on the tumor. Statistically significant differences were observed between outer, inner and core regions. Orthogonal coherence discriminated across all regions ($p < 0.01$), while peak time and mean transit time were significantly different between outer and tumoral areas ($p < 0.05$). The results showed increased vascular activity on the rim and distinct dynamics in the core, consistent with spatial heterogeneity.

In conclusion, motion corrected CUDI-derived parameters enable a quantitative and region-specific assessment of CEUS in liver metastases, supporting a more robust and less operator-dependent interpretation of CEUS imaging and a more personalized evaluation of tumor microperfusion.

Declaration on the Use of Artificial Intelligence Tools

Artificial Intelligence tools were used during the writing of this thesis. In particular, Microsoft Copilot, based on the GPT-5 chat model, was used for language editing and translation into academic English. ChatGPT, based on the GPT-5.5 model, was used for the generation of some figures included in this manuscript. These tools were not used to generate experimental results, perform data analysis, develop the proposed methodology, interpret the findings or formulate the conclusions of this work.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Problem statement

Irregular microvascular networks often develop during tumor angiogenesis, which may include increased tortuosity, structural disorganization and heterogeneity. These vascular abnormalities affect the perfusion of tissues and may influence the transport and distribution of therapeutic agents within the tumor. For this reason, quantification of tumor vascularity has become very important to understand treatment response and patient outcome, in addition to the evaluation of liver lesions.

Techniques that can image these vascular characteristics play an important role in the characterization of liver lesions. Of these, contrast enhanced ultrasound (CEUS) is a useful technique that allows real-time imaging of blood flow on both macro and microvasculature. After the injection of an ultrasound contrast agent (UCA), the movement of the microbubbles through the blood can be continually observed, providing information on the blood flow patterns within the tissues.

Currently, CEUS is mainly used for the detection and characterization of lesions. The enhancement pattern, however, during CEUS may also provide quantitative data of the microvascular architecture. Spatial heterogeneity of contrast enhancement could reflect regional variations in vascular density, tortuosity of the vessels, efficiency of vessel perfusion, and potentially drug delivery in the lesion.

While these enhancement patterns are usually evaluated qualitatively in clinical practice, resulting in subjective and operator-dependent methods, quantitative methods are required to objectively characterize their spatial heterogeneity. Such approaches may provide imaging biomarkers capable of describing tumor vascular architecture in a reproducible manner and may ultimately contribute to the investigation of relationships between vascular heterogeneity, therapeutic drug distribution, treatment response and clinical outcome.

1.2 Research questions

In this context, the overall aim of this thesis is to investigate whether contrast-enhanced ultrasound (CEUS) and Contrast Ultrasound Dispersion Imaging (CUDI) can be used to provide a quantitative characterization of the microvascular heterogeneity of liver metastases.

To achieve this objective, several research questions are addressed. Firstly, this thesis investigates how respiratory and probe-induced motion can be compensated to enable reliable extraction of time–intensity curves (TICs) and quantitative CEUS parameters. Secondly, the usefulness of CUDI in clinical CEUS examinations of liver metastases is investigated. Special emphasis is given to the identification of whether these parameters derived from the perfusion and spatiotemporal similarity metrics, such as temporal correlation, spectral coherence and mutual information, can provide meaningful information on the vascular architecture and the local contrast transport dynamics inside the lesion. Lastly, the study examines if a radial analysis framework can uncover regional variations in vascular heterogeneity.

Finally, the goal of this work is to provide objective and reproducible quantitative characterization of this vascular heterogeneity of liver metastases, as a basis for future studies that could examine the correlation between vascular heterogeneity and drug concentration, to monitor response to treatment and clinical outcome.

1.3 Technical and clinical background

1.3.1 Basis on ultrasound imaging modalities

Ultrasound (US) plays an important role in medical diagnosis. Initially limited to simple measurements of anatomical dimensions, US imaging has progressively evolved, enabling many different applications such as fetal abnormalities screening, blood flow assessment, and the evaluation of tissue texture and characteristics [1]. Ultrasound refers to high-frequency acoustic waves that can propagate through human tissues. The propagation of these waves in biological tissues is commonly described by their relationship between frequency f, wavelength λ, and propagation speed c:

$$c = f\lambda \tag{1.1}$$

where c is the speed of sound in the medium, which depends on its mechanical properties. When US waves interact with tissue interfaces, they are partially reflected and scattered back toward their source, creating diagnostic images of organs and structures within the body [2]. Figure 1.1 provides a simplified illustration of the principles of ultrasound wave emission and reflection.

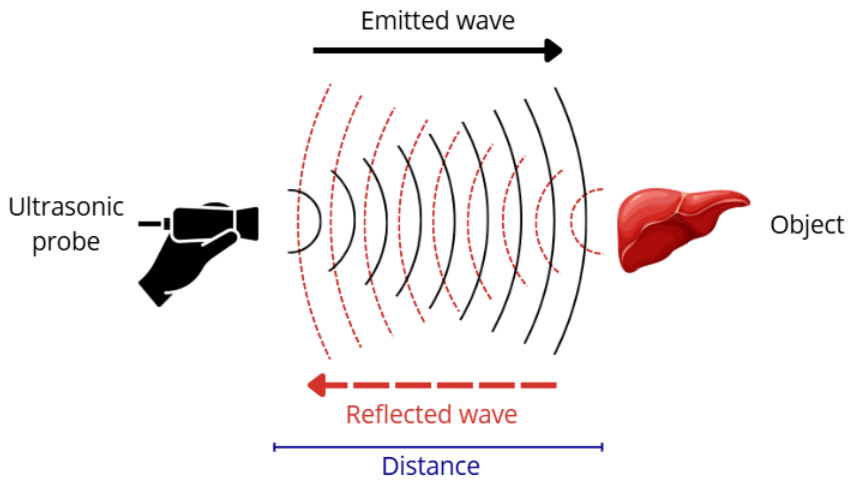

Figure 1.1: Simple scheme of US emission and reflection.

Several US imaging modalities exist. The most used is the brightness mode (B-mode) imaging, which provides a two-dimensional representation of tissues and organs. In B-mode imaging, the brightness of each pixel is proportional to the amplitude of the received echo originating from a certain spatial location. To acquire a B-mode image, a source of US waves is necessary.

For this purpose, a transducer is located on the patient's skin and it emits short US pulses into the body. After penetrating the body tissues, the signals generate echo signals once they come across interfaces of different acoustic impedance. These echoes are detected and processed to form the final B-mode image [1].

Based on the type of transducer, reported in Figure 1.2, image geometry can vary. Linear transducers produce rectangular images with parallel scanlines. Convex probes produce sector shaped images with diverging scanlines. Phased transducers also produce sector images, but this type of transducer is typically used for applications requiring a smaller acoustic window.

In addition to image geometry, intrinsic features of US images need to be considered. One critical aspect of US imaging is the presence of a characteristic granular appearence known as speckle. This phenomenon is not merely random noise, but rather the physical result of constructive and destructive interference among echoes backscattered by tissue structures that are smaller than the system resolution limits. The spatial characteristics of speckle are related to the imaging system resolution and beam geometry. In particular, when convex probes are used, the radial divergence of the scanlines causes a progressive degradation of lateral resolution with increasing depth, while axial resolution remains comparatively stable. This imbalance results in a spatially anisotropic pattern: at greater depths, speckle grains appear laterally elongated, introducing a depth-dependent variability

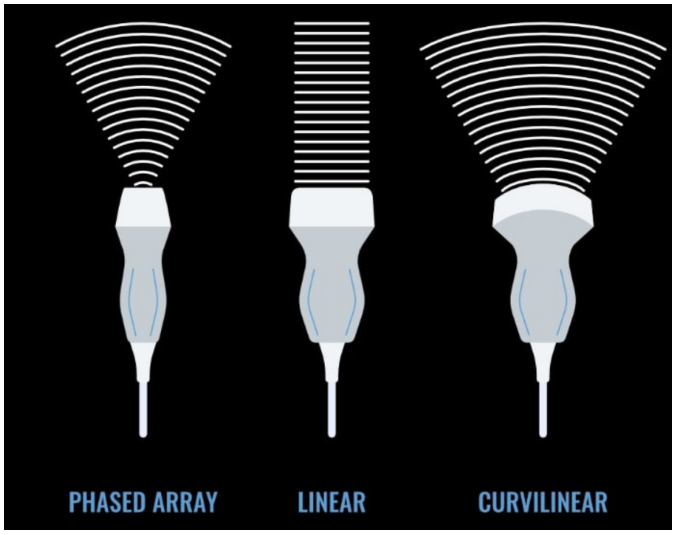

Figure 1.2: Different types of US probes.

across the image.

An alternative to standard B-mode imaging is contrast-enhanced ultrasound (CEUS), on which this project focuses. A schematic picture on CEUS imaging is illustrated in Figure 1.3.

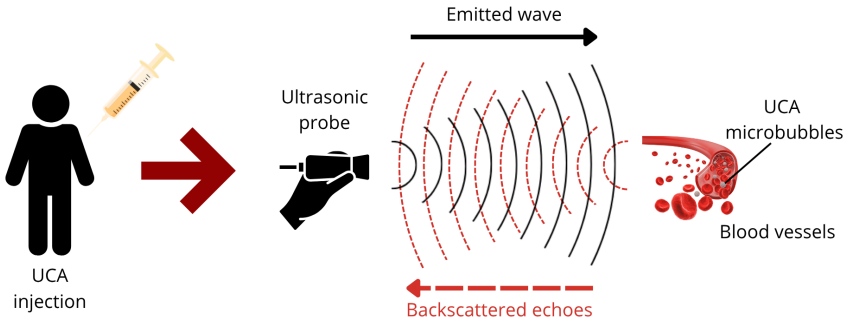

Figure 1.3: Simple scheme of how CEUS imaging modality works.

CEUS uses microbubbles as ultrasound contrast agents (UCAs) to enhance vascular imaging and allow real-time visualization of blood flow and microvascular perfusion. By analysing wash-in and wash-out dynamics of the microbubbles, CEUS

allows the assessment of tissue perfusion characteristics. The principle of CEUS lies in the similar echogenicity of blood and surrounding soft tissues in conventional US imaging; by intravenously administrate microbubbles, acoustic scattering within the vascular compartment increases, enhancing image contrast without exposing patients to ionizing radiation or toxic agents [3].

Figure 1.4 shows an example of a CEUS and a B-mode acquisition, from one of the patients whom sequences were analysed in this study. As shown, the two images

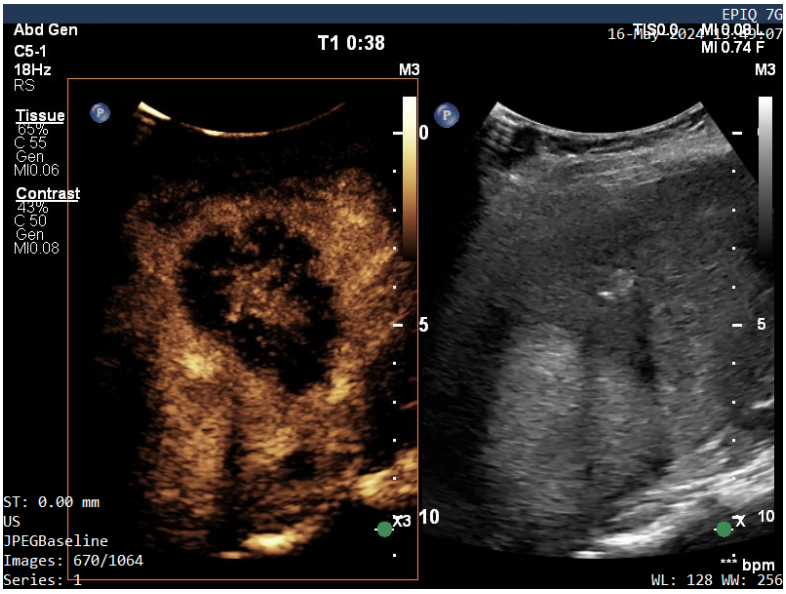

Figure 1.4: Example of a CEUS (on the left) and a B-mode (on the right) acquisitions from RASTRIC dataset.

provide complementary information. CEUS sequences highlight the presence of microbubbles in the vascular system, resulting in increased signal intensity and allowing the visualization of blood flow dynamics. In this way, CEUS enables the functional evaluation of tissue perfusion. In contrast, B-mode images provide the display of different tissues in grayscale according to their echogenicity, giving a structural representation of the organ that is useful to assess anatomical features and morphology.

Another important aspect of US imaging involves the interpretation of displayed gray levels. In clinical US systems, these pixel values do not directly correspond to the raw backscattered acoustic intensity. Instead, the received signal is processed in several steps, including amplification and logarithmic compression, to map the vast dynamic range of acoustic echoes into a limited grayscale range suitable for visualization. As a result, the relation between the displayed gray level and the acoustic intensity is nonlinear and it depends on the system, possibly affecting

quantitative intensity-based image analysis.

This is especially important to consider in the context of quantitative analysis in CEUS. The measured signal in applications like TIC analysis and CUDI is typically viewed as an indirect indicator of the concentration of the UCA. However, the relationship between the true acoustic response of the microbubbles and the resulting image intensity can be distorted using nonlinear signal processing. Therefore, it is important to linearize the UCA concentration and intensity relationship in order to allow quantitative assessments to be performed with the same physical consistency and reliability.

1.3.2 Liver anatomy

The liver is the biggest solid organ and largest gland in the human body. It is located in the superior right quadrant of the abdomen, beneath the diaphragm and to the right side of the stomach, and it is partially protected by the rib cage [4]. Anatomically, the liver is divided into two main lobes, known as right and left lobe. Functionally, however, this organ can be described as a collection of eight independent segments, each characterized by its own vascular inflow, outflow and biliary drainage [4].

On the cellular scale, the liver consists of a network of blood vessels and hepatocytes (liver cells) that form layers around a central vein. Blood flows through the sinusoids from the portal regions towards the central veins, resulting in gradients of oxygen, nutrients and metabolites in the liver. These gradients result in region-specific functional roles, a process called metabolic zonation [5][6]. In addition to hepatocytes, the liver is composed of other cells that are also associated with the liver and contribute to its vascular function, which may this affect the contrast dynamics in the liver tissue.

The liver is involved in several vital functions, including metabolism, detoxification, synthesis, storage and immunity [7]. The functional importance of this organ is closely related to the characteristics of its vascular system: the liver receives blood from two distinct sources, the portal vein and the hepatic artery. Its complex vascular system is shown in Figure 1.5. The portal vein carries about 70-75% of the liver blood volume, with blood that is rich in nutrients, coming from the gastrointestinal tract. The hepatic artery, on the other hand, provides oxygenated blood, accounting for the remaining 25–30% of the total blood supply. These two inflows converge within small vessels, known as sinusoids, where the interaction processes between the blood and hepatocytes occur [4][5].

Therefore, the correlation between the vascular system and the cellular structure of the liver plays an important role in the behavior of the organ and directly impacts the blood flow and contrast agent dynamics [6].

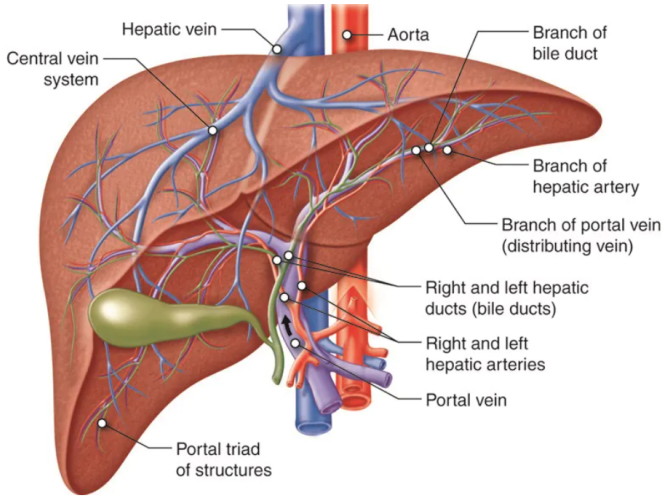

Figure 1.5: Liver and its dual blood supply. From [7].

.

1.3.3 Liver metastases

A tumor represents a localized tissue overgrowth characterized by uncontrolled expansion and invasive and aggressive behavior. Liver metastases are a hallmark of the aggressive dissemination of numerous extrahepatic cancers and represent the most common malignant lesions of the liver, with an incidence nearly 40-fold higher than primary liver tumors [8]. Indeed, between 25% and 50% of patients with non hematological malignancies develop liver metastases when diagnosed [9], making them a leading cause of cancer related mortality [8].

Metastatic progression reflects the ability of tumor cells to overcome multiple host dependent constraints, including stromal remodeling, vascular and neural infiltration, immune evasion at the primary site, and immune reprogramming at metastatic locations [8]. The high frequency of liver metastases is further explained by the liver's susceptibility to secondary colonization [8], as metastases can originate from a wide range of primary tumor sites [9]. The most common ones are:

- Breast cancer
- Lung cancer
- Genitourinary system cancers
- Gastrointestinal tract carcinomas
- Melanomas
- Sarcomas

In particular, colorectal carcinoma (CRC) represents the most common primary source of liver metastases. At the time of diagnosis, approximately 15–25% of patients with CRC already present with metastatic involvement of the liver [9].

In clinical settings, the time interval between the diagnosis of a primary malignancy and the detection of liver metastases differs markedly depending on the tumor type. In addition, hepatic metastases themselves differ in their dissemination pathways, intrahepatic growth patterns and the extent of liver injury they cause. Despite these differences, when viewed from the perspective of the target organ, metastatic tumors arising from diverse primary sites share broadly similar biological behaviors within the liver and induce overlapping hepatic responses. Consequently, research and therapeutic approaches to liver metastases should be grounded in an integrated framework that embraces tumor specific heterogeneity while also focusing on the common mechanisms governing metastatic establishment and progression in the liver [8].

These biological differences are reflected in diagnostic imaging findings. On the grayscale B-mode ultrasound imaging, the appearance of hepatic metastases varies widely depending on their primary origin and histological characteristics. Hypoechoic metastases are most frequent, especially in lung, breast and pancreatic cancer [9]. As the word hypoechoic suggests, they generate relatively few echoes, therefore they don't bounce many sound waves and they appear darker in an ultrasound image, a feature commonly associated with increased cellularity and reduced acoustic interfaces [10]. On the other hand, hyperechoic masses reflect many sound waves, therefore they appear brighter in an ultrasound image. They are usually less dense and contain air, fat or fluid [10]. This other type of metastases is present in colorectal, renal and neuroendocrine carcinomas [9]. Other types of liver metastases include metastases with peripheral halo, calcified metastases, cystic metastases and infiltrative metastases [9]. Due to this wide spectrum of grayscale appearances and the substantial overlap with benign hepatic lesions, B-mode is not the optimal imaging modality for the reliable detection and characterization of liver metastases, highlighting the need for advanced imaging techniques.

1.3.4 CEUS in liver metastases

Contrast enhanced ultrasound (CEUS) plays an important role in the detection and characterization of liver metastases, with a sensitivity ranging from 86% to 94% [11]. This high diagnostic performance is primarily related to the distinct perfusion dynamics of ultrasound contrast agents (UCAs) within metastatic lesions compared to the normal liver parenchyma. Following UCA administration, CEUS imaging allows the evaluation of three well defined vascular phases that are crucial for the identification of liver metastases [12].

The arterial phase corresponds to hepatic artefial perfusion. It starts 10-20

seconds after the injection of UCA and lasts 30-45 seconds. The portal venous phase begins approximately 20-45 seconds after the injection and lasts approximately 120 seconds, or until the portal vein is enhanced. Lastly, in the late phase, which occurs after about 120 seconds, the contrast agent is uniformly distributed in the normal liver parenchyma [12]. In Figure 1.6 there is an example of the mentioned three liver metastases phases during CEUS acquisitions.

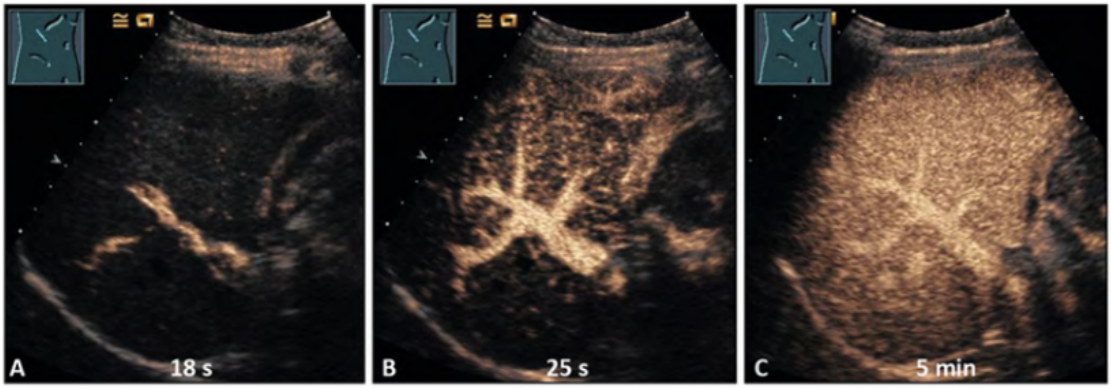

Figure 1.6: Three liver metastases phases on CEUS acquisitions. From [12].

Liver metastases are usually classified as hypovascular or hypervascular based on the vascularization of the metastases. The most common hypovascular metastases have only a few arteries and are usually found in adenocarcinomas and squamous cell carcinomas, such as colorectal, gastric, and pancreatic cancers [13]. In the arterial phase, these lesions shows either non-enhancing, hypoenhancing or peripherally rim enhancing [13][14]. Hypervascular metastases, on the other hand, show considerable arterial perfusion and they are seen as metastases with diffuse or heterogeneous hyperenhancement in the arterial phase. This is a common feature of metastases from melanomas, neuroendocrine tumors, breast cancers and renal cell carcinoma [13].

Despite their different arterial enhancement patterns, both hypo and hypervascular liver metastases have different perfusion patterns in the arteries but, characteristically, show progressive hypoechogenicity in the portal venous and late contrast phases, with a progressive contrast wash-out. This phenomenon reflects fundamental differences in vascular supply: while normal liver tissue receives approximately 70–75% of its blood flow from the portal vein and 25–30% from the hepatic artery, metastatic lesions rely exclusively on arterial neovascularization and lack portal venous perfusion [9][12]. As a consequence, UCAs are cleared more rapidly from metastatic tissue resulting in wash out, a key CEUS feature indicative of non hepatic tissue origin [12]. In Figure 1.7 a summary of how metastases appear in the three CEUS phases is shown.

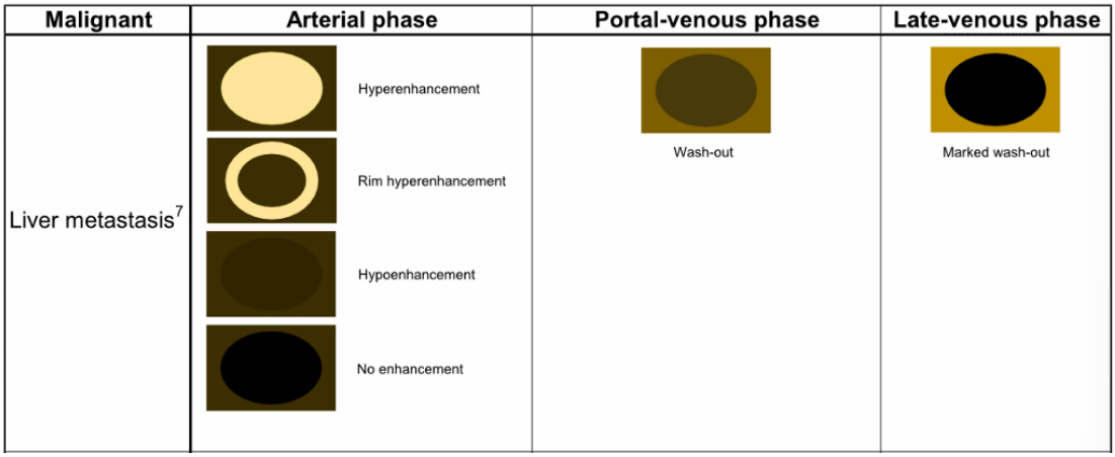

Figure 1.7: Different enhancement patterns of how liver metastases can appear during the three CEUS phases. From [14].

Thanks to these characteristic enhancement and wash-out patterns observed across the three CEUS phases, CEUS represents a highly sensitive and reliable imaging modality for the detection and diagnosis of liver metastases.

1.3.5 Qualitative and quantitative analysis of CEUS images

CEUS images can be analysed using both qualitative and quantitative analysis. Qualitative analysis refers to visual assessment of perfusion patterns. In particular, microbubble wash-in and wash-out phases can characterize the lesion, enabling the assessment of whether the lesion enhances before, after or at the same time as the surrounding tissue. One interesting parameter to visually evaluate is the lesion echogenicity, which can be compared to the surrounding healthy tissue in order to assess if the lesion is hyper-, iso- or hypo-enhanced with respect to the healthy portion of the tissue. Enhancement homogeneity can also be important to determine if the UCA distributes equally or not in the tissue.

On the other hand, quantitative analysis aims to reduce subjective interpretation by extracting numerical descriptors from the recorded CEUS sequences. Typically, signal intensity is measured over time to generate TICs, from which perfusion-related parameters can be derived. More advanced methods also examine how the contrast agent moves through the lesion, in both space and time, to give data on the vascular structure. Thus, quantitative CEUS analysis allows to objectively describe tumor vascular properties, which will decrease dependence of observer and allow to obtain reproducible parameters that could be useful for lesion characterization, follow-up, and study of the relationship between vascular structure and therapeutic response.

1.3.6 Motion correction in liver

Motion correction (MC) is the first fundamental aspect of liver acquisitions. Despite efforts to minimize motion during image acquisitions, liver imaging is always affected by both periodic and non-periodic movements. Periodic movements primarly come from diaphragmatic displacement due to respiration while non-periodic movements may be the result of intentional or unintentional probe adjustments, as well as patient movement. Specific acquisition protocols have been proposed to reduce motion artifacts. For example, probe positioning could be optimized by aligning the imaging plane with the predominant direction of breathing movements, in order to minimize out-of-plane (OOP) displacements and restricting motion mainly to the imaging plane. However, conducting acquisitions in this way is often difficult, due to the anatomical position of the liver. Indeed, the rib cage limits the available acoustic window, making it difficult to maintain the ideal probe orientation while still achieving adequate visualization of the target region [15].

Nevertheless, respiratory motion remains a problem. Conventional 2D ultrasound imaging can only acquire one imaging plane: during respiration the anatomical structures may move in and out of the field of view (FOV) part of the time and thus create motion out of plane that cannot be corrected by conventional image registration alone. This restriction is mostly attributed to the use of 2D imaging, the most prevalent imaging technique used in liver CEUS. In fact, 3D US imaging offers volumetric information and allows the tracking of tissue structures, potentially minimizing the effect of OOP motion.

Moreover, quantitative analysis requires temporal consistency of pixel intensities across the acquisition. Therefore, pixel intensities must remain as stable as possible over time to accurately reflect contrast enhancement dynamics.

Numerous MC methods for liver CEUS sequences have been suggested. First methods focused principally on rigid image registration, where motion is modeled and compensated using affine or translation models. Correlation-based methods and block matching algorithms were commonly used to estimate frame-to-frame movements and improve spatial alignment of the liver structures during the acquisition [16].

For this reason, several strategies that take breathing into account have been proposed. These techniques use the temporal periodicity or respiration to to detect frames recorded during similar phases of respiration. Zhang et al. designed an automatic respiratory motion correction system using the reference frame selection and respiratory-cycle identification. The method automatically chooses the optimal reference image using multiple learning and clustering methods and then keeps the frames of similar breathing states. It proved to be a great increase in temporal consistency and minimization of fluctuations in TICs [17].

Ta et al. introduced a two-tier in-plane motion correction scheme with adaptive

OOP motion filtering, which was a significant improvement. Within each breathing cycle, intermediate sub-reference frames (SRFs) are chosen and registered consecutively with respect to a main reference frame (MRF) in this framework, in order to reduce the accumulation error that is present in long image sequences. Furthermore, an adaptive motion filtering procedure is able to detect and eliminate any acquired frame that is out of plane to the imaging plane, which enhances the reliability of quantitative CEUS analysis [18].

More recently, Tiyarattanachai et al. introduced the Iterative Local Search Algorithm (ILSA) which is a comprehensive MC framework that performs motion compensation in the presence of periodic respiratory motion, irregular breathing and probe motion and OOP displacements. In contrast to purely respiratory-based methods, ILSA tracks the lesions by iterative local correlation search, and automatically discards OOP moving frames. The method was validated on a large, multicenter clinical dataset and showed major improvements in the quality of the TIC fitting and image consistency [15].

In general, current methods for liver CEUS motion correction include image registration techniques with techniques for respiratory motion and OOP motion management. However, obtaining strong corrections when there are complicated breathing components and long acquisitions is still a field of research.

1.3.7 Contrast Ultrasound Dispersion Imaging

Contrast Ultrasound Dispersion Imaging (CUDI) is a method proposed by Kuenen et al. [19] as an approach for quantitative characterization of microvascular architecture using CEUS data. In this framework, the transport of an UCA is modeled as a convective-dispersive process, where the temporal evolution of contrast concentration reflects the distribution of transit times through the vascular network.

The traditional perfusion-based analysis focuses on the amount and rate of blood flow. This kind of analysis can be strongly affected by other factors, such as irregular diameters and high tortuosity of the microvessels that cause flow resistance, and extravascular leakage that can lead to increasing interstitial pressure.

In contrast, CUDI aims to capture the spatiotemporal spreading of the contrast bolus, which is strongly influenced by structural features such as vascular density, tortuosity or the presence of multiple flow pathways. In fact, a complex microvascular network generates dispersion primarily because of the presence of multipath trajectories and heterogeneous flow conditions within the tissue, which makes dispersion more closely related to tissue architecture than to perfusion parameters. Ultimately, the goal of this characterization is to develop a more precise tool for the non-invasive identification of cancer lesions, supporting clinical diagnosis in biopsy targeting and guiding focal therapy.

The importance of CUDI methodologically is the extraction and analysis of time–intensity curves at the pixel level. The assumption that ultrasound signal intensity is proportional to the contrast concentration makes it possible to interpret TICs as indicator dilution curves (IDC) of the local contrast concentration as a function of time. These curves can be analysed either with a model-based method (e.g., by fitting the Local Density Random Walk (LDRW) model in order to estimate the local dispersion parameters) or with model-free methods (e.g., by calculating a spatiotemporal similarity between neighbouring TICs). In the latter case, similarity between neighboring TICs is inversely related to local dispersion. This leads to high similarity scores in angiogenesis areas where the pathways of contrast transport are correlated locally, although the microvascular network is tortuous and structurally disorganized.

An important practical limitation highlighted in the original work is that CUDI relies on the assumption that measured acoustic intensity can be related to contrast agent concentration, which requires appropriate calibration of the imaging chain. In particular, linearization of the signal is required to compensate for scanner-side logarithmic compression, ensuring that TICs accurately reflect contrast dynamics rather than system-dependent distortions. Therefore, a dedicated calibration experiment is essential to guarantee the physical consistency of the estimated parameters and to allow meaningful comparison across different acquisitions. This requirement highlights the sensitivity of the dispersion-based analysis to imaging conditions, which is an important step towards a robust and reproducible implementation. Before performing CUDI analysis, another crucial preprocessing step is speckle size regularization. Since CUDI evaluates spatiotemporal similarity using spatial kernels with fixed geometry, it assumes that the speckle pattern is spatially homogeneous and isotropic throughout the image. In practice, however, ultrasound acquisitions often deviate from these assumptions, causing similarity measurements to be influenced by image formation characteristics in addition to the underlying contrast dynamics. Consequently, a speckle size regularization procedure is required to reduce these effects and to improve consistency with the assumptions of the CUDI framework [19] [20].

Chapter 2

Materials and Methods

This chapter describes the methodological pipeline adopted in this study. The workflow can be summarized as follows:

1. Motion correction was performed to correct for in-plane and out-of-plane (OOP) movement, in order to ensure temporal consistency for pixel-wise analysis.

2. Time-intensity curves were fitted and extracted from selected regions of interest.

3. Spatial similarity analysis on CUDI framework was carried out in order to characterize dispersion patterns within the lesion.

4. Quantitative TIC and CUDI analysis results were evaluated and compared with each other.

All the processing and analyses of the dataset were performed using MATLAB (versions R2022a [21] and R2025b [22]).

2.1 Clinical Dataset

The dataset used in this study was acquired at UMC Utrecht and it includes CEUS acquisitions from 14 patients, all diagnosed with liver metastases originating from colorectal cancer (CRC). Each acquisition consists of both B-mode and contrast-enhanced ultrasound (CEUS) imaging.

For eight of the fourteen patients, data were collected at two different time points, separated by an interval of approximately two to three weeks. The other six patients were examined at one single time point. During each acquisition session, two different injections of the ultrasound contrast agent (UCA) were administered,

allowing for two CEUS sequences per session. These sequences were useful in detecting the wash-in and wash-out of the contrast agent in the liver metastases.

All acquisitions were performed using a Philips EPIQ 7G ultrasound system. Most of the datasets were acquired with a C5-1 probe, but a subset of the acquisitions was performed using alternative probes: C9-2, L12-3, L12-5, and eL18-4. As the CUDI framework requires a consistent speckle pattern to enable reliable spatiotemporal analysis, additional preprocessing was necessary to standardize the speckle size. Preliminary experiments on speckle regularization were conducted using only data acquired with the C5-1 probe. Consequently, to ensure consistency across the analysis, all acquisitions obtained with probes other than the C5-1 were excluded from the final dataset. Additional acquisitions were excluded due to excessive motion, that made them incompatible with both preprocessing and subsequent analysis. Moreover, one acquisition presented flashed frames generated by a US system setting that enables the operator to destroy the microbubbles and observe the tissue replenishment. This acquisition was also excluded from the analysis.

Therefore, the final dataset used in this project includes a total of 28 acquisitions from 11 patients. A schematic representation of the exclusion criterion is illustrated in Figure 2.1.

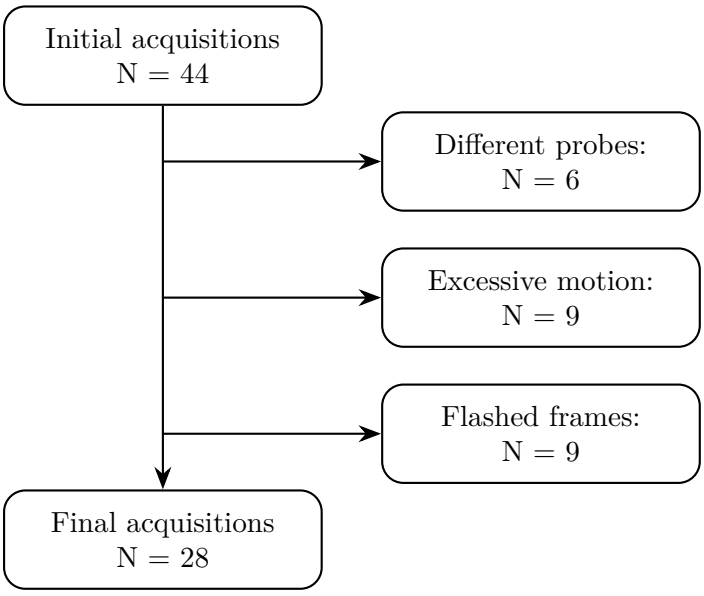

Figure 2.1: Schematic workflow of the exclusion criterion.

A summary of the dataset characteristics is reported in Table 2.1.

Table 2.1: Summary of the clinical dataset used in this study.

Parameter	Value
Initial number of patients	14
Final number of patients	11
Pathology	Liver metastases from CRC
Patients with two acquisition days	8
Patients with single acquisition day	6
CEUS sequences per session	2
Ultrasound system	Philips EPIQ 7G
Main probe	C5-1
Excluded probes	C9-2, L12-3, L12-5, eL18-4
Final number of useful acquisitions	28

2.2 Motion Correction

2.2.1 Implemented algorithm

The proposed motion correction framework was developed to address motion artifacts in CEUS acquisition of the liver both in-plane and out-of-plane (OOP) without compromising the physiologically relevant contrast enhancement dynamics. All processing steps were done in a user-defined ROI centered on the lesion, with the aim to have motion estimation mainly driven by the tumor motion and its surroundings without being influenced by highly distinctive anatomical features. Figure 2.2 shows an example of the selected ROI for RASTRIC patient 45.

The pipeline has five principal steps. The first step is the selection of a representative main reference frame (MRF), and temporal analysis is limited to the frames that include meaningful information of the contrast enhancement. Secondly, an OOP motion detection process is performed at the beginning to detect the low similarity frames with the MRF. Complementary CEUS information is then used to re-evaluate potentially relevant frames to prevent unnecessary data being excluded. Third, the respiratory motion is analyzed, to find breathing cycles and to group the acquisition into a temporally coherent set of frames. This breathing-aware organization allows the implementation of a two-level registration approach which minimizes the effect of respiratory motion. Fourthly, motion correction is carried out by rigid and deformable image registration. Motion estimation is performed on B-mode images and the transformation obtained is then applied to B-mode and CEUS images. Last, quality control measures are used to identify possible non-physiologic frames and to discard noisy registrations. Several image similarity metrics are then computed before and after the MC pipeline is performed, and the

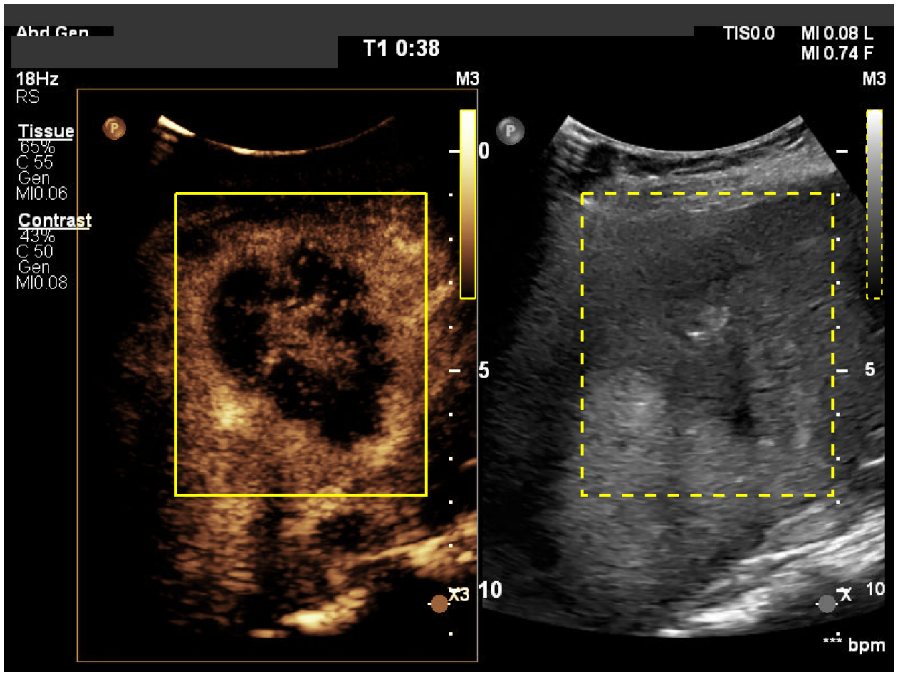

Figure 2.2: ROI selected on RASTRIC patient 45, on the CEUS (left) and the corresponding B-mode (right) image.

performance of the pipeline is quantitatively assessed. MC was then integrated prior to the CUDI framework in order to increase the temporal consistency of pixel intensities, thus enabling a more robust and reliable TIC fitting as well as a more precise spatial similarity analysis. The proposed workflow is presented in a schematic overview in Figure 2.3. The next sections detail each phase of the pipeline and the alternate techniques that were explored during its construction.

Choice of MRF and pre-processing

The selection of a representative MRF is a crucial step in the MC pipeline, as all subsequent registrations and transformations are estimated with respect to this frame. An inappropriate choice of the MRF may propagate errors and negatively affect registration accuracy and quantitative analysis.

Before selecting the MRF, a temporal pre-selection of frames was performed to focus the analysis on the contrast-enhanced portion of the acquisition. A TIC was obtained in the ROI and automatic detection of the start of contrast enhancement (wash-in) was performed using intensity- and slope-based criteria. Frames acquired before wash-in were generally excluded from further analysis, since they are characterized by low signal intensity and provide limited information for

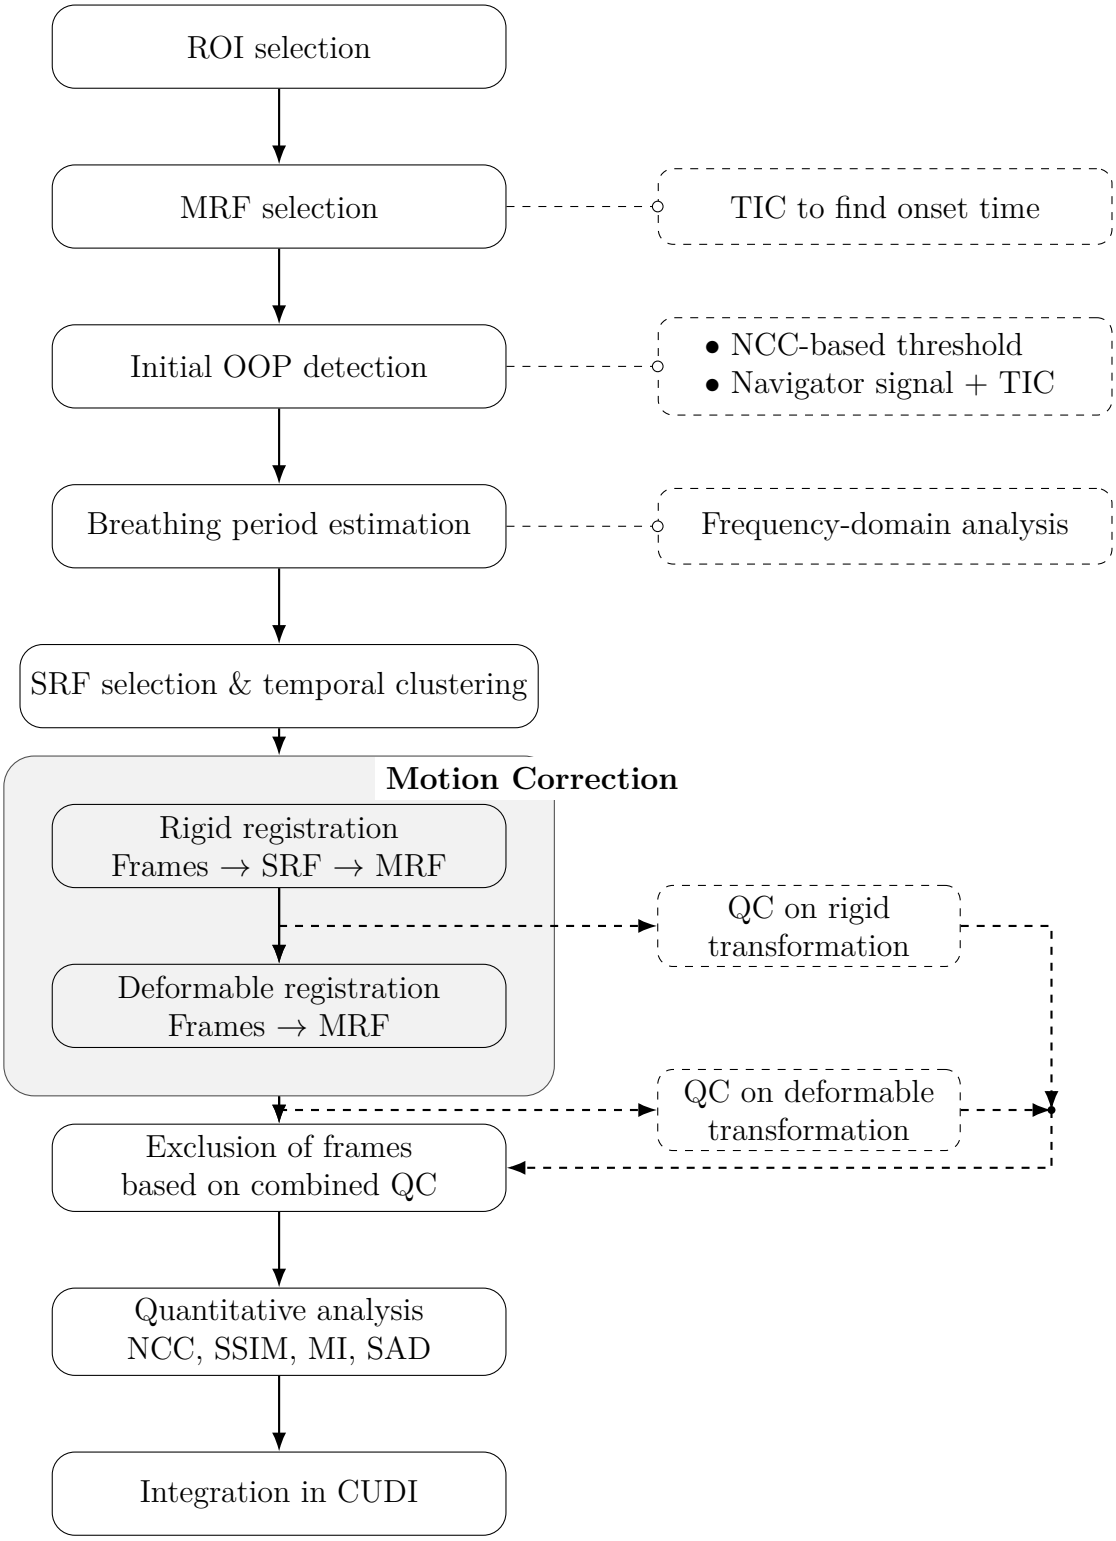

Figure 2.3: Workflow of the proposed motion correction framework.

similarity assessment. However, a short baseline segment preceding wash-in was retained, with its duration adaptively determined according to the sequence length. This ensured that relevant pre-contrast information was preserved while preventing a large number of non-enhanced frames from biasing the MRF selection. The TIC was then employed in subsequent steps along the pipeline, for the re-evaluation of the OOPs and to choose the final analysis range.

Several strategies for MRF selection were investigated during the development of the framework. A first simple approach consisted of manually selecting the MRF from the CEUS acquisition. During the arterial phase, contrast enhancement allows the tumor to be clearly visible, enabling the user to identify frames with optimal metastasis delineation. Even if this method is intuitive and can lead to good results, it is subjective and user-dependent. A second strategy followed the approach proposed by Ta et al. [18]. The MRF was selected from the B-mode acquisitions as the frame exhibiting the highest NCC with respect to the temporal mean frame. The main advantage of this method lies in its fully automatic nature. However, it relies on a fictitious reference image rather than a physically acquired one, potentially affecting the reliability of the similarity assessment. To overcome these limitations, a third approach combined the previous two methods. In this case, MRF was defined as the B-mode frame showing the highest NCC with respect to the B-mode frame that corresponds to a user-selected CEUS frame known to clearly show the tumor. Although this strategy links the reference selection to a real and clinically meaningful frame, it still requires user interaction and therefore does not completely eliminate subjectivity. A final and selected method was developed aimed at obtaining a fully user-independent MRF selection while preserving clinical relevance for CEUS analysis. The NCC was computed between all pairs of frames within the CEUS acquisition, and the MRF was identified as the frame exhibiting the highest overall similarity with respect to all other frames. Selecting the MRF directly from the CEUS sequence increases the likelihood that the reference corresponds to a time point in which the tumor is clearly visible due to contrast enhancement. In this way, the selected MRF represents the most temporally consistent and spatially stable configuration of the sequence, improving scalability and reproducibility across different patients and imaging conditions. MC is subsequently performed on B-mode sequence to ensure robustness against contrast-induced intensity variations, while the resulting transformations were applied to both B-mode and CEUS sequences.

For all strategies, similarity measures were computed within a ROI initially defined by the user, rather than over the entire image. Restricting the analysis to the ROI allows the MC to focus on the relevant area and reduces the influence of background structures. The same ROI is subsequently used in all following steps of the pipeline.

After MRF selection, a pre-processing step is performed to all of the frames. A

median filter with a 3×3 pixel neighborhood is applied to reduce speckle noise while preserving edges and anatomical boundaries. Furthermore, intensity normalization is done to compensate for global intensity variations and enhance the robustness of similarity metrics in future processing steps.

Thresholding for OOP frames

After the selection of the MRF, an initial rough OOP frame detection is performed. The NCC is calculated between the MRF and each of the remaining frames of the B-mode sequence. If the similarity between two frames is low with respect to the MRF, the frame may be affected by OOP motion and thus it could be a candidate to excluded. Various NCC-based thresholding techniques are examined to account for the different severity of motion and preferences of the user. These thresholds provide the user with options for either a "harder" or a "softer" OOP frame removal. In particular, percentile-based thresholds are used to identify frames with the lowest similarity values. The 5th percentile is a more conservative criterion, discarding only the frame which has the lowest NCC possible, whereas the 10th percentile is a more aggressive criterion that discards larger proportion of frames with low similarities. In addition, a robust statistics-based threshold defined as the median minus 1.5 times the MAD is adopted to reduce sensitivity to outliers and non-Gaussian distributions of NCC values. Figure 2.4 illustrates the NCC values calculated for RASTRIC patient number 45, together with the possible threshold values that can be selected by the user to determine if a frame is classified as OOP or not.

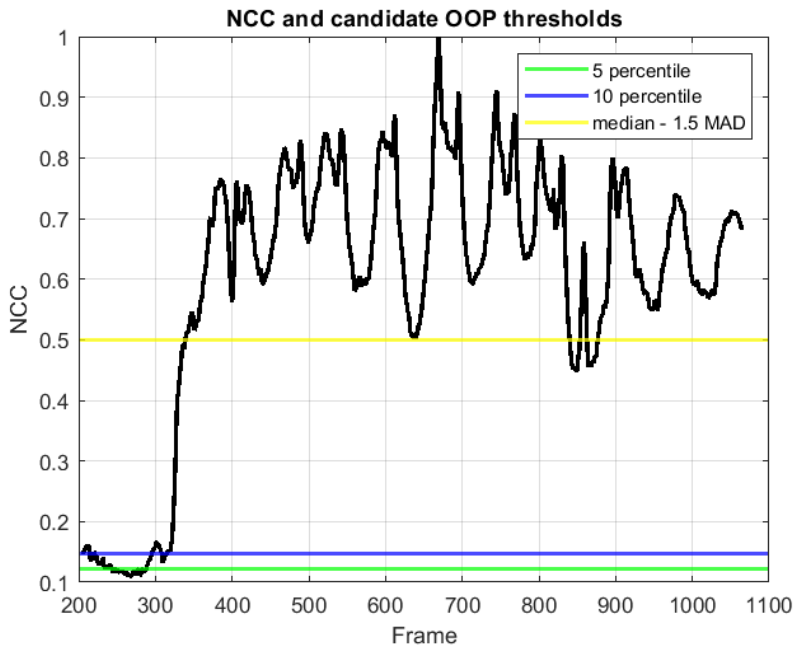

Figure 2.4: NCC and possible thresholds for initial detection of OOP frames.

Once the user selects the desired threshold, frames from the sequence with NCC values lower than the selected threshold are considered to be OOP frames and they are temporarily removed from the sequence. In fact, low values of NCC mean that the structure is less similar to the reference frame, and typically in US acquisitions this is due to OOP motion or decorrelation of speckle patterns.

Reinsertion of OOP frames

A CEUS-based navigator signal is added after the threshold selection, in order to support the re-evaluation of frames that were not included. The navigator consists of a single line intersecting the tumor, defined as the average of five input lines manually selected by the user. This signal is plotted over time, allowing visualization of the tumor motion after UCA enhancement. On the CEUS navigator plot, candidate OOP frames identified in Section 2.2.1 are highlighted in red, enabling a direct visual assessment of their temporal correspondence with the UCA enhancement pattern. In addition to the navigator signal, the corresponding TIC is displayed. An example of this combined visualization is shown in Figure 2.5. CEUS navigator and TIC together allows the user to identify frames presenting low similarity that may be incorrectly excluded because they correspond to physiological contrast dynamics, such as wash-in or wash-out phase of the UCA.

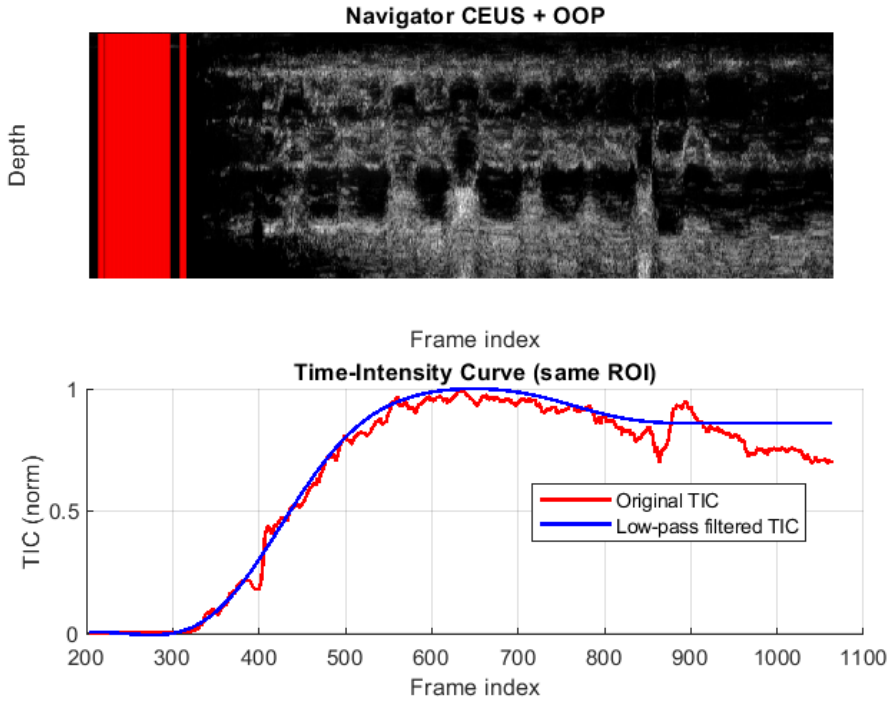

Figure 2.5: Plots used to decide whether the initially detected OOP frames should be reinserted. Top: navigator-based CEUS signal, with OOP frames identified during the initial detection highlighted in red. Bottom: averaged TIC within the ROI to visualize the UCA dynamics.

Based on this evaluation, the user may decide to reinsert some or all the selected frames into the sequence if they are clinically and/or temporally relevant (for e.g. if they are located within the wash-in period). This step introduces a supervised validation stage that complements the automatic thresholding, reducing the risk of excluding clinically relevant frames.

Breathing period, SRFs and clusters

Once the final set of in-plane frames has been defined after OOP frames exclusion and reinsertion, the frames are divided into clusters in preparation for a two-tier registration strategy, inspired by the work of Ta et al. [18]. To this end, respiratory motion periodicity is first estimated in order to enable a breathing-aware registration [17]. The NCC is computed on B-mode acquisition after removal of OOP frames. The fast Fourier transform (FFT) of the NCC signal is then calculated and the dominant frequency within a physiological respiratory range is identified in order to estimate the breathing cycle period, following the approach proposed by Zhang

et al. [17]. The frequency search range is set between 0.1 and 0.7 Hz, which correspond approximately to 6-42 breathings per minute.

After the identification of the dominant frequency, the corresponding breathing period is derived as the reciprocal of the dominant frequency. Peaks in the NCC signal are then detected using the estimated breathing period as the minimum peak distance, scaled by a factor of 0.8 to allow tolerance to respiratory variability. These peaks are chosen to be SRFs and they serve as intermediate reference points during the two-stage registration approach [18]. Specifically, one SRF per breathing cycle is identified.

All other remaining frames are then assigned to clusters that are temporally close to each of the SRFs. Therefore, each cluster is associated with a correspondant SRF, which guarantees that the frames are in a breathing cycle. In this way, registration accuracy is stabilized and improved. An example of the resulting clustering is shown in Figure 2.6. This clustering strategy assumes a relatively regular respiratory cycle. In reality, respirations can be asymmetric or vary over time, which can lead to suboptimal frame assignments. However, the method offers a useful approximation for classifying frames into a similar respiratory state.

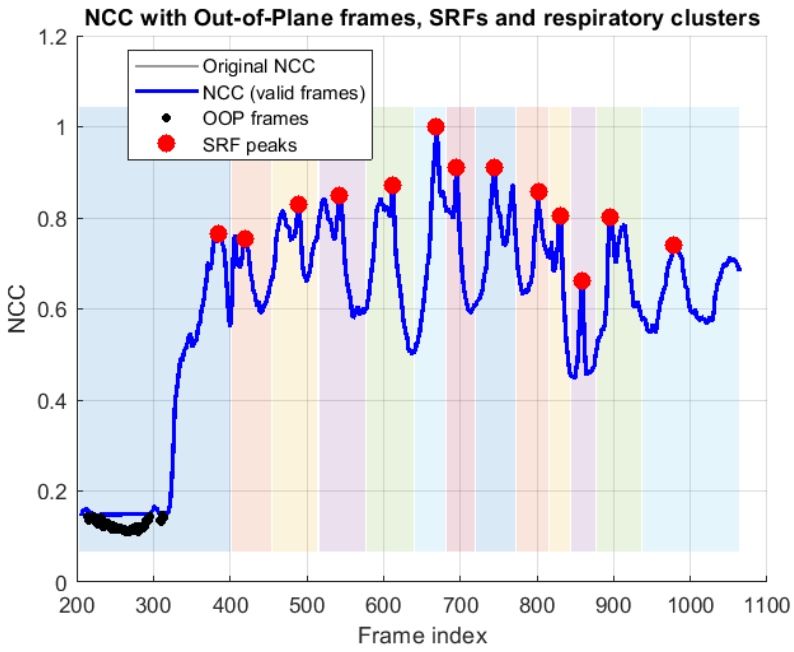

Figure 2.6: NCC values (blue), OOP frames (black dots), SRFs (red markers), clusters assigned to each SRF shown in different colors.

Effective motion correction

Navigator-based motion correction Before choosing a final efficient MC pipeline, several other strategies were investigated. A first, simple approach relied on the CEUS-based navigator signal to estimate and compensate motion. Since the navigator provides a one-dimensional temporal trajectory of the tumor, this approach was initially used to estimate vertical translational motion by identifying a sinusoidal pattern associated with respiratory displacement, as seen in Figure 2.7.

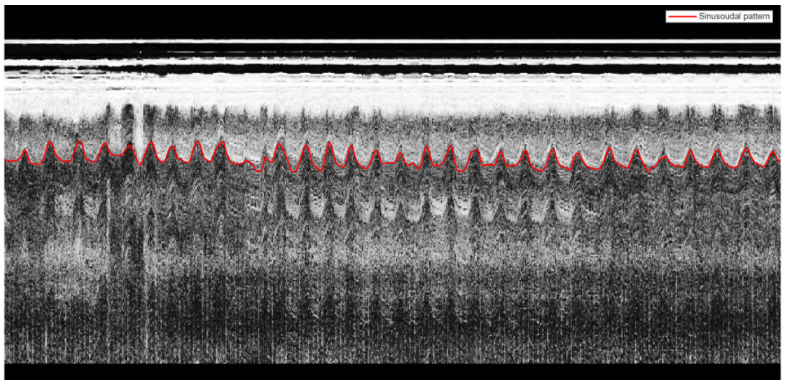

Figure 2.7: Example of the sinusoidal pattern extracted from the navigator signal of the RASTRIC patient 38.

The tumor trajectory was extracted by applying morphological operations to the navigator signal, estimating the displacement of the tumor position for each frame. A frame-wise vertical translation was then applied in order to reverse this displacement and realign the tumor to a fixed position over time. The application of this correction resulted in stable tumor positioning at the expense of vertical motion in the surrounding image. Although this approach proved effective in cases dominated by vertical respiratory motion, it failed to handle more complex motion patterns, such as horizontal or oblique displacements. Therefore, due to its limitation to a single direction of motion, this approach was not retained as part of the final motion correction pipeline and was instead used as a qualitative tool to support OOP frame identification and the interpretation of contrast dynamics.

Explored alternatives and discarded approaches To overcome the limitation of navigator-based compensation, which is inherently constrained to a vertical direction, alternative strategies were investigated to estimate motion direction directly from image data. Among these, Principal Component Analysis (PCA) was evaluated as a potential approach to identify the dominant direction of motion in a data-driven manner to later compensate the acquisition using the navigator in that principal direction. PCA identifies orthogonal directions that maximize variance in

the data [23]. In this context, it was assumed that respiratory motion represents the dominant source of variance in the image sequence, so that the first principal component would be aligned with the main direction of respiratory displacement and could be compensated using a single directional correction. However, in CEUS acquisitions, total variance is strongly influenced by additional factors, including speckle noise, contrast enhancement dynamics and local tissue deformation. As a result, PCA did not lead to a clear separation between motion-related components and other sources of variability. Multiple components exhibited comparable variance contributions, preventing the identification of a unique dominant motion direction. Consequently, PCA-based motion estimation proved unreliable and insufficiently discriminative for robust MC in CEUS liver acquisitions and was therefore not pursued further.

Since PCA can be formulated through the singular value decomposition (SVD) of the centered data matrix, SVD-based analysis was also explored to confirm whether motion-related energy could be isolated into a small number of dominant modes. However, similarly to PCA, motion-related information was distributed across multiple singular components with comparable energy contributions, preventing reliable motion characterization. Consequently, SVD-based motion analysis did not provide additional advantages and was not pursued further.

Block matching techniques were also evaluated to estimate a displacement field for each frame, with the aim of identifying the predominant motion direction and magnitude within each acquisition. This approach involved monitoring local image blocks between adjacent frames and computing a median displacement vector to represent global motion. This method, however, was unsuccessful for US imaging. CEUS images present high speckle noise, low signal-to-noise ratio and quick decorrelation of intensity patterns from frame to frame. These properties significantly impair reliable block correspondence, leading to unreliable and incoherent displacement vectors. As a result, the resulting displacement fields were found to be highly inhomogeneous and produced no meaningful and consistent quiver plot. Basing MC on these noisy displacement fields led to worse alignment and generation of artifacts. Therefore, block matching-based displacement field correction proved inappropriate for MC in CEUS liver acquisitions and it was not included in the final MC pipeline.

Registration framework Following the preliminary investigations, MC was formulated as an image registration problem, as these methods provide a more general and robust framework to estimate motion directly from spatial information. Two alternatives were tested to find the best framework for implementation: the open-source image registration software *Elastix* and the MATLAB built-in functions.

Elastix is an intensity-based registration framework that offers a library of algorithms to perform image registration by optimizing registration parameters in

an iterative manner. The registration workflow is controlled by an easily configured parameter set that defines key registration elements like the similarity metric, the optimizer, the interpolator and the transformation model. The framework operates by defining a fixed image and a moving image: the fixed image serves as the reference towards which the registration is performed, while the moving image is iteratively transformed to achieve alignment [24].

In parallel, MATLAB built-in functions, including *imregtform* and *imregdemons*, were evaluated. The function *imregtform* estimates global geometric transformations (e.g., rigid or affine) between a moving and a fixed image by optimizing a similarity metric. The resulting transformation is represented by a 3×3 matrix encoding translation and rotation in the rigid case, and scaling and shear in the affine case. The function *imregdemons* instead estimates a dense displacement field that aligns a moving image to a fixed reference by computing a local displacement vector for each pixel, allowing local non-rigid alignment between images.

All registration procedures were implemented within the ROI previously defined, ensuring that motion estimation was driven by the lesion and its surrounding region. A preliminary comparison was done by directly registering all frames to the MRF using various transformation models, such as rigid, affine and deformable registration. This simple configuration enabled a comparison of the two frameworks under very similar conditions.

In this context, the navigator based correction was also compared to registration-based methods. As a result, the navigator did not offer any improvement on motion compensation over rigid image registration. Intensity-based registration methods, in contrast, were more successful at capturing the dominant motion components without the need of a predefined direction of motion.

The evaluated transformation models demonstrated to have significant advantage in frame alignment, with rigid registration being the best one. Affine transformations, which extend the rigid model by including scaling and shear components, showed only marginal additional benefits while significantly increasing computational cost. For this reason, affine registration was not retained in the final pipeline. Deformable registration was further considered to compensate for residual local motion not addressed by rigid alignment. Both *Elastix* and MATLAB-based implementations were tested. Both frameworks gave similar results, so practical considerations were done for the final choice. Built-in functions of MATLAB enabled simpler integration into the existing processing workflow, and reduced implementation complexity. Therefore, for the final MC pipeline, MATLAB built-in functions were chosen as the final one.

Pipeline configuration and selection After selecting the optimal framework and transformation models for registration, various pipeline configurations were tested to identify the most effective approach to MC.

Initially, rigid registration was carried out by directly registering all the frames to the MRF. However, this approach was often resulted in aligning frames from different breathing cycle due to the presence of respiratory motion, increasing misalignment and temporal inconsistency.

In order to address this limitation, two alternative pipeline configurations were evaluated. In the first configuration (Configuration *A*), a two-tier method was adopted to perform rigid registration. Each frame was aligned to its respective SRF, which was subsequently registered to the MRF. The resulting transformations were then composed to form a set of transformations that would map all frames to the MRF, minimizing the effect of inter-cycle differences. Each frame was then directly registered to MRF using deformable registration.

In the second configuration (Configuration *B*), rigid registration was done between the individual frames and the MRF without an intermediate SRF. Deformable registration, instead, followed a two-tier strategy, where each frame was first aligned to the SRF and the resulting deformation was subsequently propagated towards the MRF. These configurations were designed to evaluate whether the two-tier strategy is more effective when applied to global (rigid) or local (deformable) motion components. Figure 2.8 summarizes the difference between them.

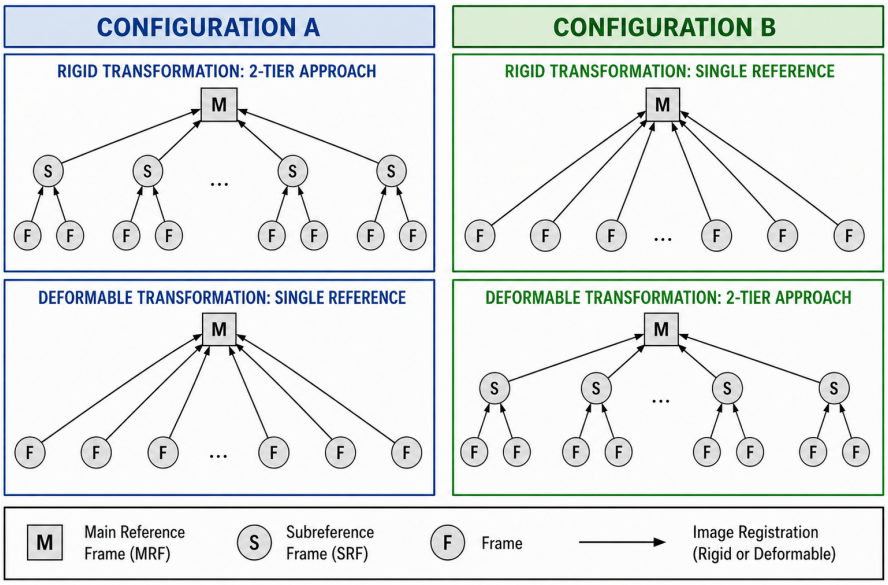

Figure 2.8: Comparison between different MC pipeline configurations. Schematic workflows are taken from [18].

Motion estimation was performed on B-mode data within the defined ROI, and the resulting transformations were applied identically to both B-mode and CEUS

sequences, ensuring consistent alignment across modalities. A subset of 8 representative acquisitions was used to compare the candidate pipeline configurations. The subset included one acquisition per patient and it covered different levels of motion, signal quality, and lesion characteristics. The selected configuration was subsequently applied to the full dataset of 28 acquisitions for the final analysis. The two configurations were quantitatively compared using the similarity metrics NCC, SAD, SSIM and MI, defined in Section 2.2.1, computed before and after MC. Configuration *A* demonstrated superior robustness and reduced accumulation of local deformation errors in all the evaluated sequences. By implementing direct deformable registration to the MRF, consistent spatial reference was ensured, avoiding progressive distorsions introduced by multiple transformation steps. Based on this comparison, Configuration *A* was adopted as the final motion correction pipeline. The detailed quantitative results of this comparison are reported in Section 3.1.1, while a more deepen discussion of these results is presented in Section 4.1.1.

Quality control

A dedicated quality control (QC) procedure was conducted separately for each registration step to identify potentially unreliable transformations and residual OOP frames. The QC strategy was based on the analysis of the transformation parameters returned by the MATLAB registration functions.

For rigid registration, the translation magnitude and rotation angle were extracted for each frame. The frames where the transformation values exceeded the 95th percentile of the distributions for the associated parameters were marked as possible OOP frames. For deformable registration, QC was based on the analysis of the estimated displacement field. The same approach was used for frames that contained deformation magnitudes that exceeded the 95th percentile of the distribution of deformations within the ROI.

The results of the rigid and deformable QC were subsequently combined. Frames exceeding either criterion were considered potentially affected by OOP motion or registration failure and were excluded from further analysis in both B-mode and CEUS data. The motivation of this strategy is the expected characteristics of physiological liver motion. Despite the presence of motion throughout the acquisition, it is mostly driven by respiration, therefore relatively smooth and limited in magnitude. Transformations which significantly diverge from the sequence may indicate that they are unlikely to be physiological motion and may indicate the presence of residual OOP motion or registration failure. On the basis of this, frames that were found to have unusually large rigid or deformable transformation were removed.

Similarity metrics

To ensure a fair comparison between pre- and post-motion corrected data and to avoid bias introduced by edge artifacts resulting from the registration process, similarity metrics were computed within a smaller ROI focusing on the tumor region. This ROI was used for both the pre- and post-corrected data, in order to compare the same spatial support.

Several similarity measure parameters such as Normalized Cross-Correlation (NCC), Sum of Absolute Differences (SAD), Structural Similarity Index Measure (SSIM) and Mutual Information (MI) were used to measure the performance of the MC pipeline. These metrics are complementary and give information about the quality of global and local alignment between frames.

The NCC between two images I_1 and I_2 is defined as:

$$NCC = \frac{\sum(I_1 - \bar{I}_1)(I_2 - \bar{I}_2)}{\sqrt{\sum(I_1 - \bar{I}_1)^2 \sum(I_2 - \bar{I}_2)^2}} \tag{2.1}$$

where $\bar{I}_1$ and $\bar{I}_2$ represent the mean intensities. NCC measures the linear correlation between images, with higher values indicating stronger similarity.

The SAD is defined as:

$$SAD = \sum |I_1 - I_2| \tag{2.2}$$

This metric quantifies the absolute intensity difference between images and it is sensitive to misalignment.

The SSIM evaluates perceptual similarity by combining luminance, contrast, and structural information:

$$SSIM = \frac{(2\mu_1\mu_2 + C_1)(2\sigma_{12} + C_2)}{(\mu_1^2 + \mu_2^2 + C_1)(\sigma_1^2 + \sigma_2^2 + C_2)} \tag{2.3}$$

where μ, σ and σ_{12} represent mean, variance and covariance, respectively, and C_1, C_2 are stabilizing constants.

MI measures the statistical dependency between images and is defined as:

$$MI = \sum_{i,j} p(i,j) \log\left(\frac{p(i,j)}{p(i)p(j)}\right) \tag{2.4}$$

where $p(i,j)$ is the joint probability distribution and $p(i)$, $p(j)$ are the marginal distributions. MI is robust to intensity differences and reflects the amount of shared information between images.

For each metric, values were computed with respect to the MRF both before and after MC. Summary statistics were reported using the median and the Median Absolute Deviation (MAD) to ensure robustness to outliers. Results are presented in Section 3.1.2.

2.3 Speckle Size Regularization Experiment

As previously discussed in Section 1.3.1, US images acquired with convex probes present a depth-depedent anisotropic speckle pattern, due to divergent scanlines. However, spatiotemporal similarity analysis employed in CUDI assumes a homogeneous and isotropic spatial structure, since it relies on spatial kernel with fixed geometry. As a result, the presence of anisotropic and depth-dependent speckle may introduce bias in the estimation of similarity metrics and dispersion parameters. This discrepancy between model assumptions and acquisition physics motivates the need for speckle size regularization. For this reason, a dedicated experiment was designed to characterize and regularize the speckle size under controlled conditions, in order to enforce spatial homogeneity and reduce depth dependency [20].

Experimental setup

The experiment was conducted at UMC Utrecht, aiming at reproducing a controlled acoustic environment together with a homogeneous distribution of UCA microbubbles. The complete setup is visible in Figure **??** A clinical ultrasound system (Philips EPIQ 7G) equipped with a C5-1 convex probe operating in CEUS mode was used. A water basin was filled with degassed water to minimize cavitation and unwanted acoustic reflections. To guarantee a homogeneous distribution of the microbubbles in the medium, a magnetic stirrer (IKA RCT basic) equipped with a magnetic bar submerged into the basin was used. The ultrasound probe was placed slightly below the water level and kept non-orthogonally to minimize specular reflections, following standard in-vitro CEUS protocols. An acoustic absorber was added to the basin to suppress reflection artifacts even further.

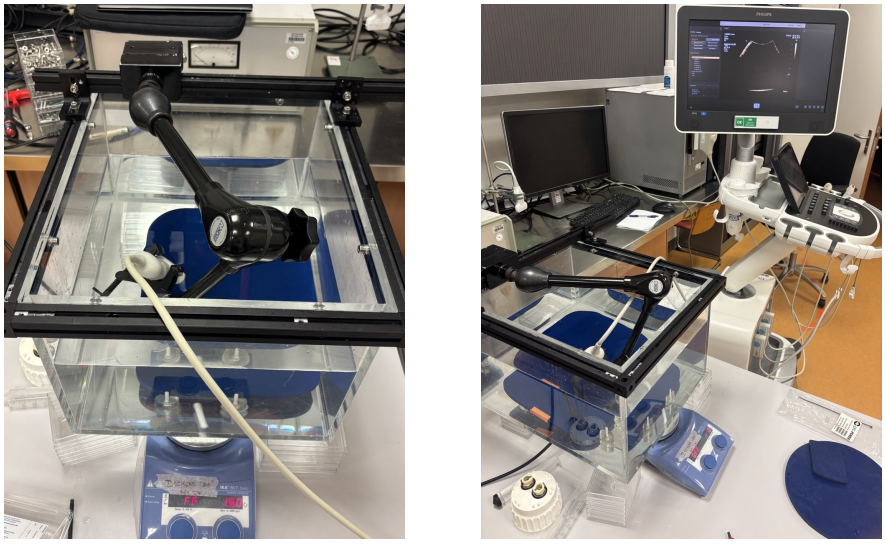

(a) **(b)**

Figure 2.9: Experimental setup. (a) Water basin filled with degassed water, probe and probe holder, magnetic stirrer and magnetic bar, acoustic absorber. (b) Complete setup and US system Philips EPIQ 7G used for the acquisitions.

Acquisition Protocol

SonoVue (25 mg lyophilized powder) was used as UCA. The solution was prepared by adding 5 mL of saline (9 mg/L concentration) to the vial, and then extracting 0.25 mL of UCA to obtain an estimated final concentration of ~ 0.05 mg/L in the basin. The chosen low concentration was used to allow the detection of individually distinguishable scatterers and to be able to estimate the speckle size accurately [20]. Each acquisition lasted 30 seconds. A baseline acquisition was first performed using only degassed water to characterize the background signal. Subsequently, the contrast agent was injected and multiple acquisitions were recorded under varying imaging conditions. The selected imaging parameters correspond to standard US system settings that influence signal strength, image contrast and image quality. A brief explanation of these parameters is provided below.

The mechanical index is a safety-related parameter that quantifies the likehood of mechanical bioeffect, such as cavitation. It depends on the acoustic pressure and the transmission frequency, and it is particularly relevant in CEUS due to its effect on microbubble stability [25]. Dynamic range (DR) controls the number of gray levels displayed in the image, determing in this way the contrast resoluton. A higher DR results in smoother grayscale transitions, while a lower DR increases contrast between structures. Output power defines the amplitude of the transmitted US waveform. Increasing this parameter improves the intensity of the received

echoes, increasing signal strength but also acoustic exposure. For this reason, it requires careful adjustment according to safety guidelines [26]. Gain refers to the overall amplification of the received echo signals and it controls the global brightness of the image. On the other hand, Time Gain Compensation (TGC) adjusts the amplification as a function of depth. In this way, it compensates for signal attenuation and it ensures a more uniform brightness of the image [26]. Persistence is a temporal filtering technique that averages consecutive frames, reducing noise while compromising temporal resolution. Increasing this parameter will smooth the image, while decreasing it generates a more speckled image [27]. The imaging depth determines the penetration of the US beam and the FOV: increasing depth allows the inspection of deeper structure, but the frame rate and spatial resolution are reduced. Finally, XRES is a proprietary image technique for speckle reduction and edge enhancement [28]. Overall, these parameters were selected and controlled to ensure a consistent acquisition protocol.

The following parameters were kept constant:

- Mechanical Index: 0.08
- Dynamic Range: 50 dB
- Output Power: -27.5 dB
- TGC: centered
- Persistence: OFF

The following parameters were varied:

- Gain: 40%, 50%
- Depth: 4, 7, 10 cm
- XRES: ON/OFF

The variations of certain parameters were introduced to evaluate their effect on speckle appearance, considering that image processing and acquisition settings directly influence the measured intensity distribution and, therefore, subsequent quantitative analysis.

In the end, the magnet rotation speed was primarily set to 150 RPM to maintain homogeneous bubble distribution, with an additional acquisition performed at 100 RPM to assess the effect of reduced motion of the magnet.

2.3.1 Acquisition Overview

A total of 12 acquisitions were performed, including both baseline and contrast-enhanced conditions (see Table 2.2). Most of the acquisitions were performed in CEUS mode. Two acquisitions were performed in dual mode to allow comparison between contrast-specific and conventional imaging modalities.

Table 2.2: Summary of acquisition parameters for the speckle size experiment.

Condition	Acquisition	Gain (%)	Depth (cm)	XRES	Magnet Speed (RPM)
Without contrast	1	50	10	ON	150
	2	50	7	ON	150
	3	50	4	ON	150
With contrast (0.25 mL SonoVue)	4	50	10	ON	150
	5	40	10	ON	150
	6	50	10	OFF	150
	7	40	10	OFF	150
	8	40	7	OFF	150
	9	40	4	OFF	150
	10 (dual)	40	10	OFF	150
	11 (dual)	50	10	OFF	150
	12	40	10	OFF	100

2.3.2 Analysis

Speckle size characterization was conducted through a detailed 2D spatial autocovariance analysis of the CEUS signal.

Local regions were extracted for each frame using a window sliding method. The 2D autocovariance function was calculated and then averaged across all frames to minimize temporal variation for each window position. The autocovariance function was averaged with the assumption of a gaussian shape of the speckle pattern. The lateral and axial dimensions of the speckle were then estimated by fitting the elliptical model to the averaged autocovariance function. The resulting values of the speckle sizes were therefore converted to physical units (mm) using the spatial calibration parameters derived from the DICOM metadata. Unrealistic values were excluded from the data based on predefined thresholds, and orientation inconsistencies were corrected to maintain alignment with the expected pulse propagation direction. Finally, the relationship between speckle size and depth was investigated by fitting the estimated speckle dimensions as a function of the distance

from the probe using a linear model, giving a clear mathematical description of how speckle characteristics change with the distance from the probe [20].

2.4 Calibration experiment

The aim of the calibration experiment is to verify the relationship between UCA concentration and backscattered acoustic intensity. For low concentrations of UCA and low mechanical index, the relationship between the concentration and the acoustic intensity is expected to be linear [29]. This correlation, however, is not directly seen on the image intensity that is observed, since the image intensity displayed from the ultrasound scanner is logarithmically compressed.

Since the compression function used by the scanner is proprietary and unknown, a simplified logarithmic model was used to approximate between the gray level displayed and the acoustic intensity.

$$Q = a_0 \cdot \ln\left(\frac{I_{lin}}{I_{ref}}\right) \tag{2.5}$$

where I_{lin} is a quantity proportional to the acoustic intensity after compensation for logarithmic compression, and I_{ref} is an internal reference intensity defined by the scanner manufacturer and it is unknown. Since I_{ref} is constant for a given scanner configuration, it only affects the scaling of the recovered signal and can therefore be absorbed into the model parameters. Under this assumption, the displayed gray level can be related to a linearized signal intensity I_{lin}, assuming $I_{lin}/I_{ref} \approx I_{lin}$. This intensity is expected to be linearly proportional to the concentration of UCA, at low concentrations:

$$I_{lin} = a_1 \cdot C + a_2 \tag{2.6}$$

For this reason, the experiment was designed to compensate for the image compression and verify if the relationship between the UCA concentration and acoustic intensity is preserved after linearization, so that the relationship between UCA concentration and acoustic intensity can be assumed as an IDC and the parameters can be obtained from the IDC fitting analysis.

Experimental setup

The experiment was conducted at UMC Utrecht. Similarly to the speckle size regularization setup (Section 2.3), a clinical US system Philips EPIQ 7G was equipped with a C5-1 probe. The probe was fixed on a probe holder and positioned just below the water surface in a basin filled with degassed water, with an inclination of approximately 120°. An acoustic panel was used to reduce reflection artifacts.

Condom-based phantoms containing different UCA concentrations were used as imaging targets.

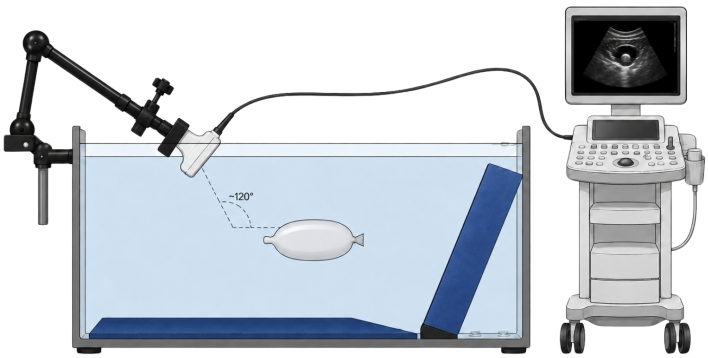

Figure 2.10: Schematic drawing of the experimental setup: C5-1 probe, connected to the Philips EPIQ 7G scanner, is fixed on a probe holder and immersed in a basin filled with degassed water. A condom filled with UCA is positioned just below the water level in the basin, while an acoustic absorber panel is used to minimize reflection artifacts. AI-generated image based on author's prompt.

Acquisition protocol

For each UCA concentration, a separate acquisition was recorded according to the protocol described below. A starting SonoVue contrast solution (5 mL of 45 mg/mL lyophilized product) was made by placing 5 mL of saline (9 mg/L concentration) into a contrast vial and gently shaking. Subsequently, 0.0111 mL of this solution was diluted into 200 mL of saline, resulting in a base concentration of 2.5 mg/L. From this base solution, multiple volumes were taken and different dilutions were generated to obtain a range of UCA concentrations. The list of all the concentrations used is reported in Table 2.3.

Table 2.3: Preparation of UCA dilutions starting from a base concentration of 2.5 mg/L (final volume = 120 mL).

C (mg/L)	Volume contrast solution (mL)	Volume saline (mL)
0.05	2.4	117.6
0.10	4.8	115.2
0.20	9.6	110.4
0.40	19.2	100.8
0.60	28.8	91.2
0.80	38.4	81.6
1.00	48.0	72.0

For each concentration, the solution was mixed and 120 mL were used to fill a condom phantom, minimizing air inclusion before closing the phantom. Each phantom was then submerged in the water basin and manually translated under the probe for approximately 30 s during a dual-mode acquisition. Sliding the phantom was important to ensure that each frame corresponded to a slightly different slice, improving the sampling of bubble distribution. All the performed acquisitions are summarized in Table 2.4. For each concentration, multiple acquisitions were performed.

Table 2.4: Acquisition protocol for the C–I calibration experiment. Acquisitions were performed at different UCA concentrations under constant imaging conditions (XRES OFF, mechanical index = 0.08, depth = 10 cm). Each concentration presents multiple acquisitions.

Concentration (mg/L)	Gain (%)	DR (dB)
0.05	47	50
0.1	47	40–60
0.4	47	40–60
0.8	47	40–60
1	47	40–60
2	47	40–60
10	47	50

Analysis

From the adopted model (Equation 2.5), the linearized intensity was computed as:

$$I_{lin} = e^{\frac{Q}{a_0}} \tag{2.7}$$

In addition, consistently with the behavior reported in the literature [29], linearized intensity can be expressed as:

$$I_{lin} = a_1 \cdot C + a_2 \tag{2.8}$$

where a_1 represent the sensitivity of the scanner, while a_2 is the baseline intensity. Therefore, knowing Q from the acquisitions and the relative UCA concentration, it is possible to compute the corresponding I_{lin} and verify whether the relationship between C and I is linear.

The analysis was performed using MATLAB (version 2025b). Prior to the analysis, the gray level Q was compensated by reversing the colormap applied by the scanner, in order to obtain a quantity proportional to the compressed signal before displaying the mapping.

Several frames were selected for each concentration and a ROI was manually defined. The phantom was tracked across the selected frames and the mean gray level was extracted for each frame. A corresponding value of I_{lin} was derived from Q. The resulting pairs (C, I_{lin}) were plotted and fitted with a linear model.

2.5 CUDI-based analysis

As briefly introduced in Section 1.3.7, CUDI-based analysis consists of two main components: IDC fitting analysis and spatiotemporal similarity analysis. Both methods require dedicated preprocessing steps before the actual analysis can be executed. Once preprocessing is completed, quantitative parameters for tumor characterization can be extracted.

The ultimate output of both analysis types is the generation of parametric maps that display the estimated quantitative parameter as color-coded values overlaid on the anatomical US image.

2.5.1 IDC fitting analysis

TICs quantify the acoustic backscattered intensity following the injection of an UCA at a given pixel or ROI in the US image as a function of time. For physiological modeling of perfusion, however, information is needed in terms of UCA concentration. This type of information is provided by the IDCs, which indicate the relative distribution of UCA over time.

The measured TICs are then used to derive the concentration curves in the IDC and fitted with a mathematical model to obtain important hemodynamic parameters for the characterization of the tumor. The conversion from TICs to IDCs is based on the linear relationship between the UCA concentration and the acoustic intensity at low concentrations and low mechanical index. For this reason,

a dedicated calibration experiment was conducted and it is reported in Section **??**. Since acoustic intensity can be assumed to be proportional to microbubble concentration under these conditions, a properly calibrated and linearized TIC can be interpreted as an IDC. This step is essential, as the mathematical models used to characterize tumor vasculature operate on concentration-based data rather than raw intensity values.

A schematic representation of the IDC fitting analysis pipeline is provided in Figure 2.11 to summarize the main processing steps.

Preprocessing

Before IDC fitting, spatial and temporal filtering are applied on raw US sequences to improve the signal-to-noise ratio and reduce the computational time.

- Spatial filter: a Gaussian filter ($\sigma = 0.5$ mm, downsampling factor 3) is used to match the typical spatial scale of microvascular structures relevant to angiogenesis imaging. This choice follows the rationale proposed in [19], assuming that microvascular heterogeneity occurs at similar spatial scales across solid tumors, including liver metastases, making an organ-independent filter design appropriate. Downsampling by a factor of three after spatial filtering reduces computation time substantially while preserving the relevant spatial information.

- Temporal filter: a zero-phase finite impulse response (FIR) low-pass filter is applied along the time axis. A cutoff frequency of 0.2 Hz is adopted, lower than the 0.5 Hz used in prostate CUDI studies [19] to better capture the slower perfusion dynamics characteristic of hepatic tissue and to limit the influence of noise and residual motion on the TIC profiles.

All subsequent processing steps are performed on the filtered signals and follow Kuenen's approach[19].

Conversion to IDC and fitting domain

After filtering and linearization, TICs are interpreted as IDCs and used for model fitting. The fitting procedure is performed in the logarithmic domain, since US intensity noise has a multiplicative nature and becomes approximately additive with more homogeneous variance after logarithmic transformation.

Recirculation removal

A critical aspect of IDC analysis is the presence of recirculation: after the first passage of the UCA bolus, subsequent passages through the same region distort

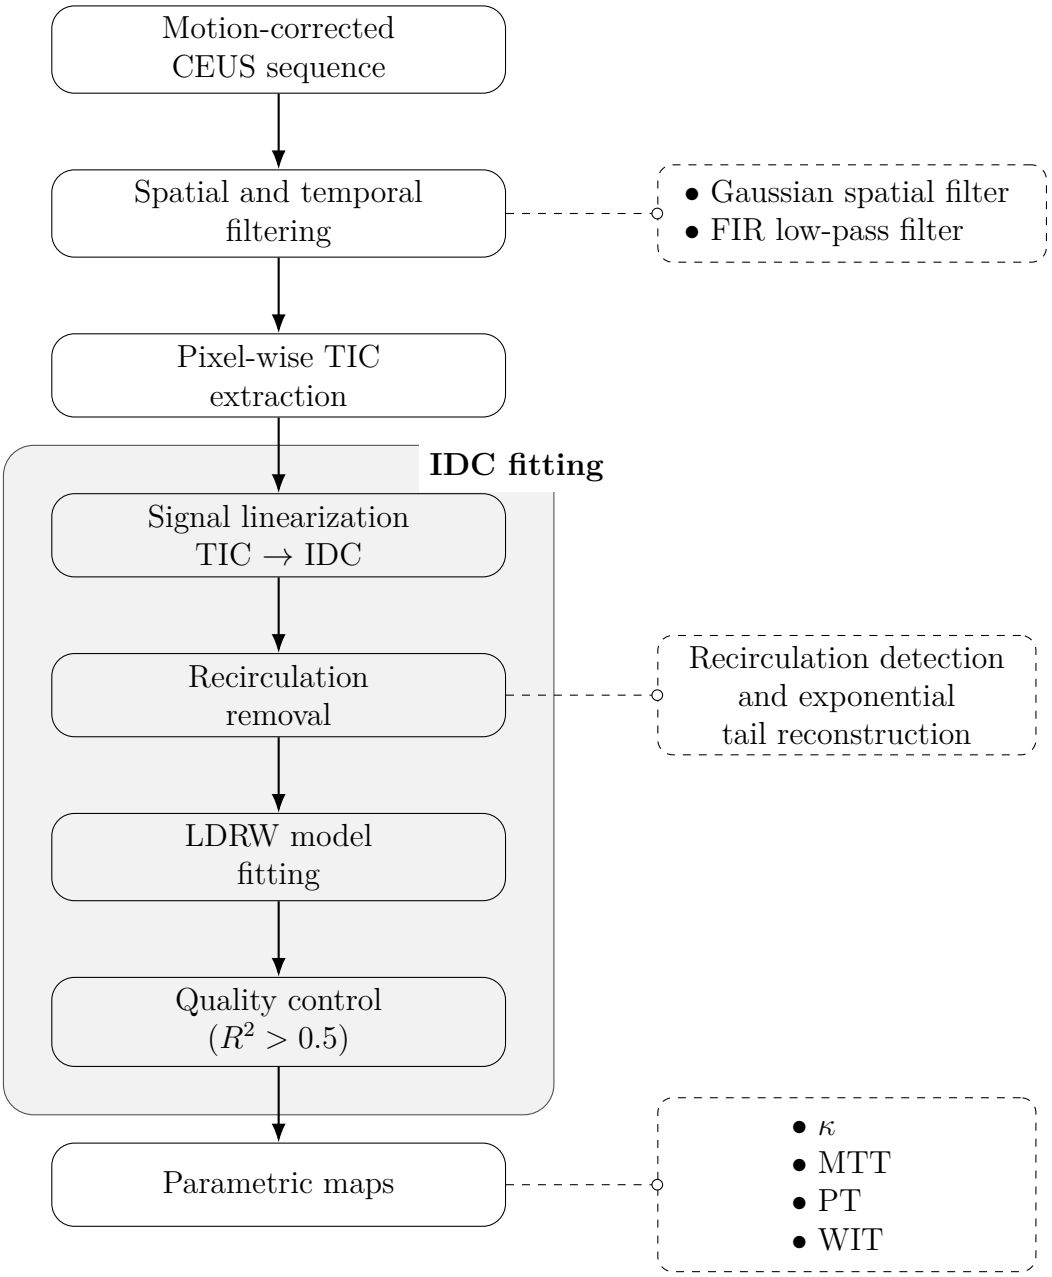

Figure 2.11: Workflow of the IDC fitting analysis. Motion-corrected CEUS sequences are filtered and converted into IDCs, which are fitted using the LDRW model to extract perfusion-related and dispersion-related parameters. The resulting parametric maps are subsequently used for radial lesion analysis.

the tail of the curve, making it no longer representative of the first-pass dynamics that the model describes. To avoid fitting the corrupted tail, the curve is truncated at a time t_{stop}, estimated by identifying the interval over which the IDC decay in the logarithmic domain is best approximated by a linear function, corresponding to a purely exponential decay in the real domain. The search for t_{stop} is restricted to the range where the IDC amplitude has decayed to between 20% and 50% of its peak value, and an additional low-pass filter at 0.1 Hz is applied beforehand to minimize effects of noise on this estimate. For the area of the curve beyond t_{stop}, an exponential extrapolation is performed to reconstruct the complete first-pass IDC tail, which is needed for the subsequent moment-based parameter initialization.

LDRW model

The physical and mathematical framework adopted to model the temporal evolution of the IDC is the Local Density Random Walk (LDRW) model. This model represents an intravascular UCA transport as a one-dimensional convective diffusion process and it represents the movement of the microbubbles and their dispersion throughout the vasculature. In particular, it offers an analytical solution to the mono-dimensional convective diffusion equation. This analytical solution is given by:

$$C(t) = AUC\sqrt{\frac{\kappa}{2\pi(t-t_0)}}exp\left(-\frac{\kappa(t-t_0-\mu)^2}{2(t-t_0)}\right) \tag{2.9}$$

where $C(t)$ represents the UCA concentration at time t, AUC denotes the area under the curve and acts as a scaling parameter related to local perfusion and contrast agent concentration, t_0 is the arrival time of the UCA in the tumor, μ stands for the mean transit time (MTT), defined as the average time needed for microbubbles to transverse the microvascular network, and finally κ is a key parameter, as it represents the dispersion parameter. It is defined as the ratio of the squared convective velocity to the dispersion coefficient, as visible in Equation 2.10.

$$\kappa = \frac{v_l^2}{2D_l} \tag{2.10}$$

where v_l is the local convective velocity of the microbubbles and D_l is the dispersion coefficient.

Physically, κ quantifies the balance between convection and dispersion. From the previous work on prostate [19], a high κ value indicates that convection dominates over dispersion. This behavior is typically associated with tumor tissue, where increased vascular tortuosity and microvascular confinement limit lateral dispersion of microbubbles.

Parameter estimation

By fitting the LDRW model to the measured IDCs, obtained after pixel-wise linearization of the raw TICs, local hemodynamic parameters can be estimated independently of the bolus dilution history. This is a significant benefit over the standard perfusion-based methods that can be affected by overall dilution dynamics. The framework to be used for the estimation of parameters is Maximum Likelihood (ML) estimation, which treats the IDC as an observed histogram of the transit time of microbubbles. This statistical approach is particularly selected for its superior robustness in the presence of multiplicative noise, typical of CEUS.

The estimation process proceeds in two distinct stages that guarantees computational efficiency and accuracy. First, a non-iterative initialization is found by calculating the statistical moments of the IDC. In this way, closed-form estimates of the model parameters are obtained requiring iterative optimization. These initial values are then given into Levenberg-Marquardt (LM) algorithm, which numerically optimizes the likelihood function in the logarithmic domain, refining the solution. As done in Kuenen work[19] at most, five iterations are carried out, which has been demonstrated to not significantly improve precision of the estimation.

A quality control step was added to ensure the reliability of the extracted parameters, based on the goodness-of-fit. Pixels whose determination coefficient (R^2) of the LDRW fit in the logarithmic domain was greater than 0.5 were used for further analysis. Pixels that did not satisfy this criterion, usually because of low SNR or motion artifacts, were discarded from the parametric maps to avoid inclusion of non-representative hemodynamic data. This study lowered the threshold for $R^2 < 0.5$ compared to Kuenen's original work on prostate tumors, in order to capture the higher noise levels observed in hepatic CEUS than in prostate imaging.

Extracted parameters

Several parameters were extracted from the fitted IDCs in order to characterize the hemodynamic properties of the lesions. These parameters can be categorized into a dispersion-related parameter, derived from the LDRW model, and conventional perfusion-related parameters used for comparative analysis.

- Dispersion parameter κ: derived from the LDRW model and already defined in this Section, representing the balance between convective and dispersive transport.

- Mean Transit Time (MTT): represented by the model parameter μ, it identifies the average time required for the UCA microbubbles to traverse the local microvascular network.

- Peak Time (PT): defined as the time elapsed between the injection time (t_0) and the time when the IDC reaches its maximum value.
- Wash-in Time (WIT): this parameter represents the time needed to increase UCA concentration from 10% to 90% of the peak concentration.

Figure 2.12 illustrates representative IDC together with the conventional perfusion-related parameters extracted from it.

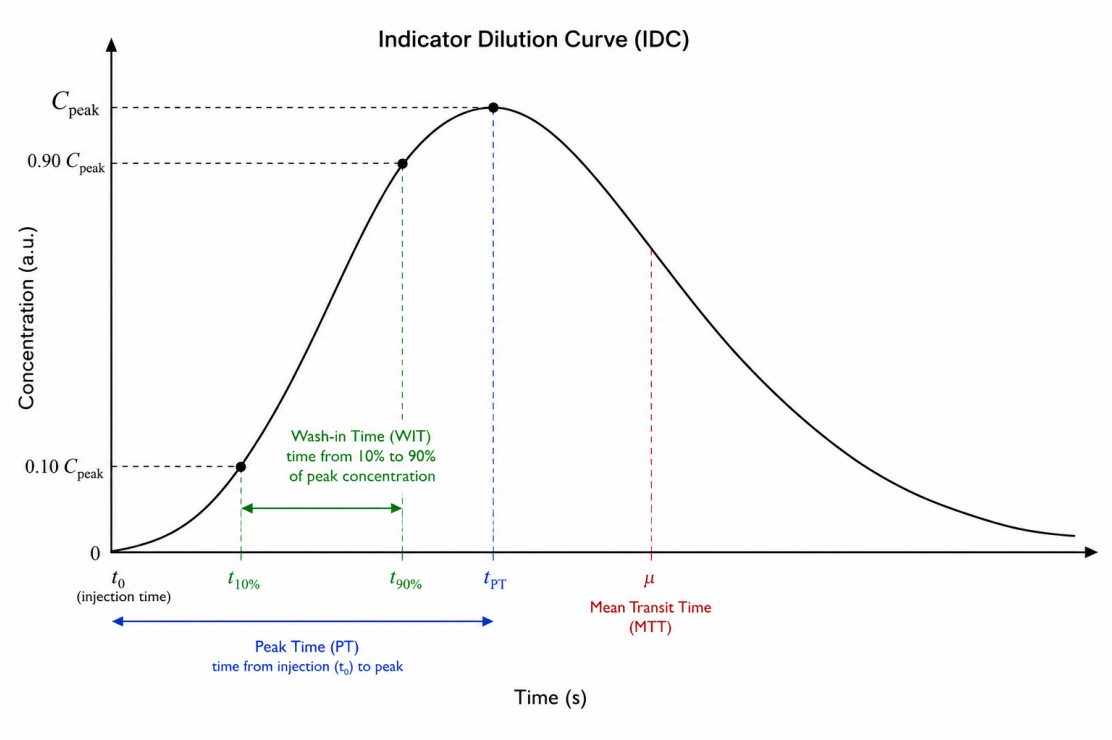

Figure 2.12: Representative IDC illustrating the perfusion-related parameters extracted in this study. AI-generated image based on author's prompt.

The four extracted parameters are subsequently used for lesion characterization and for comparison with conventional perfusion-based metrics.

2.5.2 Spatiotemporal similarity analysis

As an alternative to IDC fitting, spatiotemporal similarity analysis offers an indirect assessment of local dispersion by evaluating the relationship between IDCs at neighboring pixels. This methodology is based on the physical rationale that the chaotic and tortuous microvascular architecture of a tumor limits the spatial spreading of the contrast bolus and thus results in lower dispersion. Consequently,

in angiogenic regions, IDCs of adjacent pixels tend to exhibit a higher degree of shape similarity compared to those in healthy tissue [19].

A major advantage of this analysis is that it does not require model fitting or isolation of the bolus first-pass, making it less sensitive to artifacts such as recirculation and baseline noise.

Similarly to IDC fitting analysis, Figure 2.13 shows a workflow of the spatiotemporal similarity analysis, summarizing the main processing steps from CEUS signal preprocessing to the generation of similarity parametric maps.

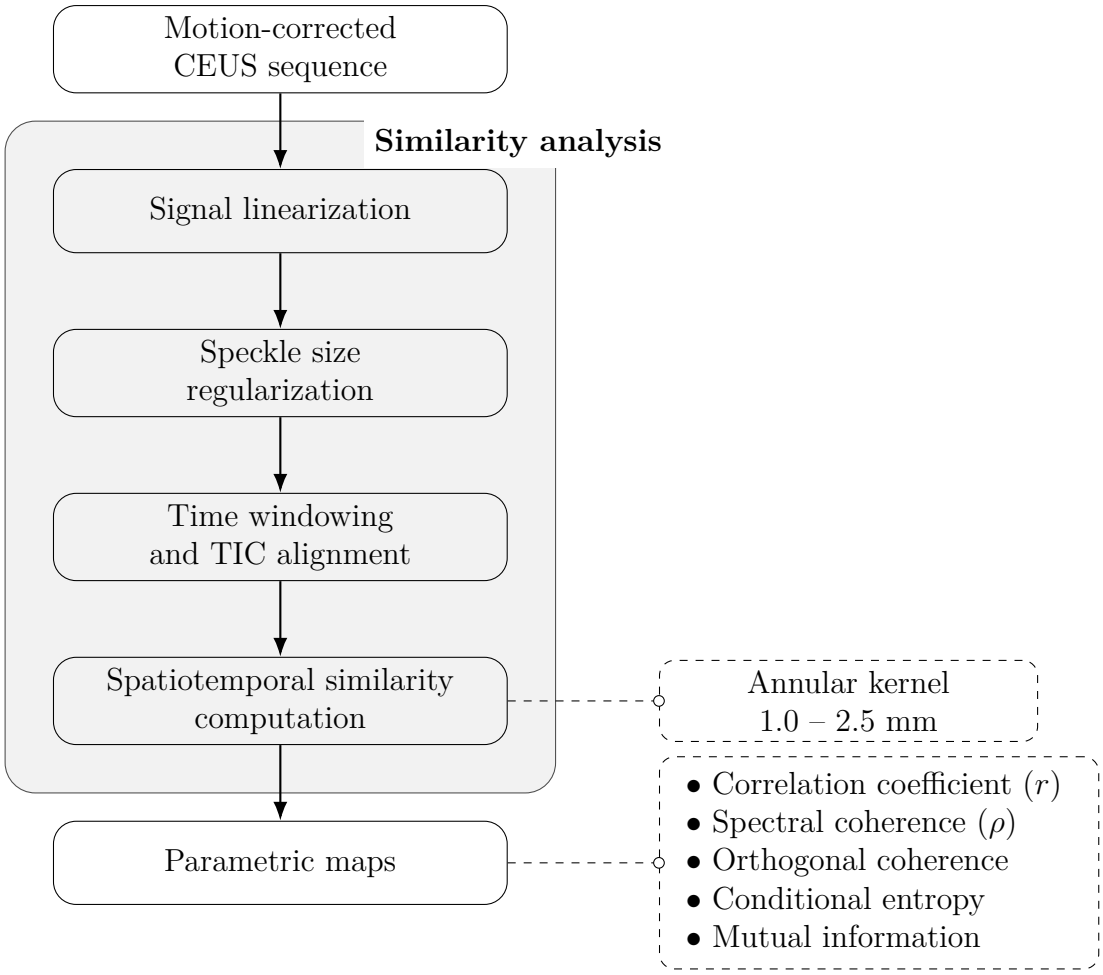

Figure 2.13: Workflow of the spatiotemporal similarity analysis. Motion-corrected CEUS sequences are linearized, regularized, and temporally aligned before computing similarity metrics within an annular spatial kernel. The resulting similarity maps provide information on local contrast transport dynamics and vascular heterogeneity.

Preprocessing

To ensure that similarity estimates reflect true hemodynamics rather than imaging artifacts, the following preprocessing steps are implemented:

- Data linearization: as in IDC fitting analysis, raw TICs are linearized to compensate for scanner-side logarithmic compression. This step ensures that the resulting data is proportional to the actual concentration of microbubbles.

- Speckle size regularization: as explained in Section 2.3, US system resolution is anisotropic due to divergent scan lines and depth dependency. In order not to bias similarity results, a spatial filtering step based on Wiener deconvolution is applied to regularize the speckle size to an approximately isotropic value, ensuring that similarity measures present homogeneous spatial resolution in the entire image and is not influenced by direction-dependent resolution differences.

- Time windowing: a time window of 35-40 s is applied to each TIC, starting at the estimated local appearance time (t_{app}), in order to focus the analysis on the most informative portion of the bolus passage and to minimize the influence of background noise. In this way, the signals are aligned and not distorted by bolus arrival time differences when performing similarity calculations.

Spatial kernel and parametric mapping

A ring-shaped (annular) kernel is used to compute similarity at each pixel, in order to ensure the statistical independence of the similarity values and not introduce artifacts from the resolution limitations of the system. This kernel compares the central pixel with those located at a distance between 1.0 mm and 2.5 mm. The inner radius allows to not consider pixels that otherwise would be too similar just because of US beam width. On the other hand, the outer radius is selected to match the physical scale of angiogenic networks. This kernel design also ensures that the analysis remains independent of the dominant flow direction.

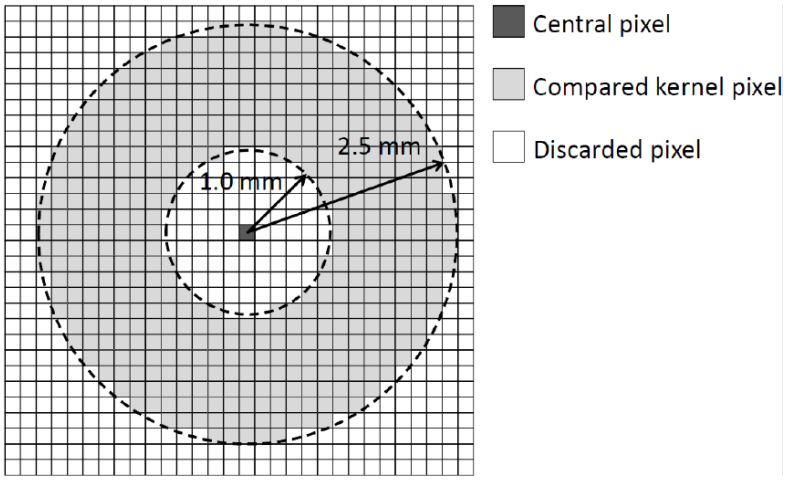

Figure 2.14: Annular kernel used when computing spatiotemporal similarity analysis. From [19].

The final output is a parametric similarity map, where color-coded values of the extracted parameter are overlaid on the anatomical images, highlighting regions of reduced dispersion.

Extracted similarity parameters

The similarity can be quantified using different parameters:

- Spectral coherence (ρ): it is calculated in the frequency domain (usually within a 0–0.5 Hz bandwidth) and it measures the correlation between the magnitude spectra of neighboring curves [19]. The main advantage of this parameter is that it does not take into account phase shifts, but it evaluates the shape of the curves regardless differences in contrast arrival times.
- Correlation coefficient (r): it is calculated in the time domain. Even if it is inherently sensitive to time offsets, the use of a window-based realignment allows r to be used effectively. The main advantage of r is that it keeps phase information, capturing additional kinetic details useful for tumor characterization.
- Orthogonal coherence: it is calculated in the frequency domain and it quantifies how similar neighboring TICs are after subtracting the common spectral component. Therefore, it emphasizes local variations in the transport of contrast and it is more sensitive to spatial heterogeneity.
- Conditional entropy: it measures the uncertainty in the neighboring TICs after taking into account the information in the current TIC [30]. Higher values

correspond to more variability and complexity in the dynamics of local contrast enhancement over time, while lower values correspond to more predictable dynamics over time.

- Mutual information: it measures statistical dependency of neighboring pixels without the assumption of a linear relation. This method describes nonlinear similarities in the UCA transport kinetics which may not be detected by traditional correlation or coherence techniques.

2.5.3 Radial analysis of tumor regions

For the extraction of the parameters of interest described in the previous paragraphs, a region-based analysis was performed to investigate the spatial heterogeneity of tumor perfusion.

For each CEUS sequence, the contour was manually delineated. In most cases, the lesion exhibited a rim-enhancing pattern presenting a hyperenhanced ring surrounding a hypoechoic or necrotic core. Based on this manually defined contour, a set of concentric regions was automatically generated. In particular, three external rings and three internal rings were constructed around the original contour, all geometrically concentric. In this way, multiple radial region were defined, enabling the characterization of perfusion properties from the outer area toward the inner tumor regions. The core of the tumor, not belonging to any of the defined internal rims, was analysed as a third separated region. In Figure 2.15 an example of these regions is present.

For each region, parameter values were extracted from the corresponding parametric maps and summarized using the median. Median was chosen to provide a more robust estimate of central tendency, reducing the influence of noise, the presence of outliers and inaccuracy of local fitting. Therefore, to highlight global trends, the median values of the three external rims were combined across all the datasets to obtain a representative value for the outer tumoral region. The same was done for the three inner rims. Median was also computed on core values to obtain global values base on all the datasets. Regions with a number of pixels below a certain defined threshold ($Th = 100$) were not considered for the global analysis, in order to ensure robust statistical estimates and to avoid potential biased results from poorly sampled rims. Regions with a limited number of pixels, in fact, are more likely influenced by noise, local fitting inaccuracies and outlier values. Implementing this threshold, resulting measures of median and dispersion are derived from significant sample size, providing a more reliable characterization of the tumor.

In the end, these descriptors were compared to investigate radial patterns of perfusion and to assess tumor heterogeneity. This analysis allows the identification of characteristics spatial trends in the extracted parameters, supporting the

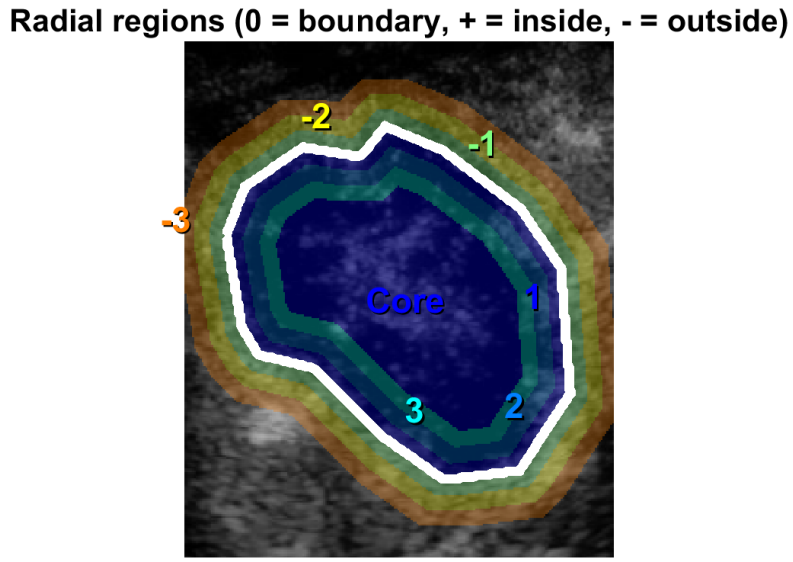

Figure 2.15: Representation of rims for the radial analysis on one patient. The white line corresponds to the manually delinated tumor contour; negative rims are the outer rims; positive rims indicate the inner region of the tumor; the central area is the core of the tumor.

differentiation between tumor regions and surrounding tissue.

2.5.4 Statistical analysis

To establish whether the considered parameters were able to discriminate between different tumor regions, a non-parametric statistical analysis was performed. In particular, the Wilcoxon signed-rank test was computed using the MATLAB in-built function *signrank*. It was used to determine paired comparisons between regions (outer, inner, core). This test was selected because the data did not necessarily satisfy the assumptions of normality required for parametric tests, and because measurements from different regions were derived from the same dataset, constituting paired observations. The Wilcoxon signed-rank test is used to compare the difference between two observations, to rank the magnitudes of the differences and to determine whether there is symmetry of sign of difference around zero. This enables the test to determine whether there is one region in which the parameters tend to have higher or lower values than in the other or not, without taking any assumption about the distribution assumed by the parameters. In this case, it was used for comparing:

- outer vs inner regions, to evaluate differences between surrounding parenchyma and tumor rim;

- inner vs core regions, to investigate intratumoral heterogeneity;
- outer vs core regions, to compare overall contrast between healthy tissue and necrotic areas.

A p-value < 0.05 was considered as statistically significant for each comparison. The Wilcoxon test gave a powerful structure to test for systematic differences in the distribution of parameter values between regions, because of the paired structure of data and the possibility of potential outliers.

Chapter 3

Results

This chapter presents the results obtained at the different stages of the proposed analysis pipeline. First, the performance of the MC approach is evaluated. Then, the results of the speckle size regularization and calibration experiments are reported. Finally, the outcomes of the quantitative analysis of CEUS data are presented.

3.1 Motion Correction

3.1.1 Choice of best pipeline

In order to select the most suitable MC pipeline configuration, a quantitative evaluation of performance was carried out by comparing similarity metrics (MI, NCC, SSIM, SAD) before and after correction for both configurations A and B on a subset of 8 representative sequences.

Figure 3.1 reports the distribution of the absolute improvements obtained for each metric. As shown in Figure 3.1, pipeline A achieved larger median improvements than pipeline B for all of the investigated metrics. The higher differences were observed for MI and SAD. The paired Wilcoxon signed-rank test indicated statistically significant differences between the two pipelines for all metrics ($p = 0.0078$).

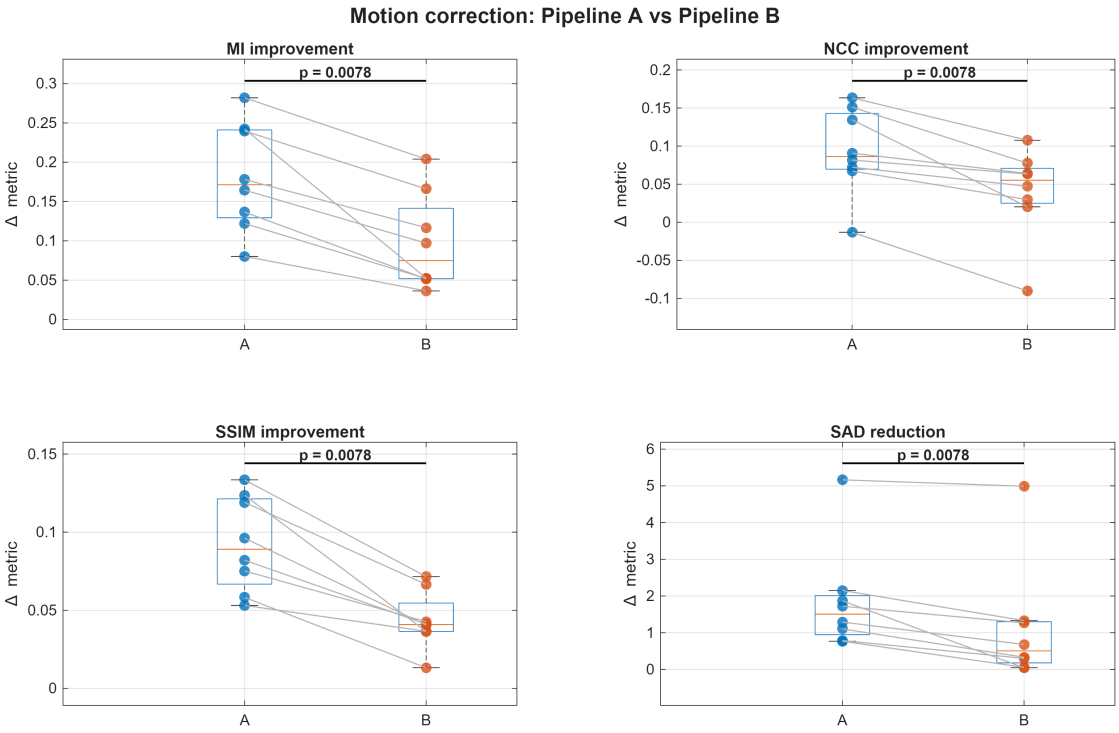

Figure 3.1: Boxplots of improvement for NCC, SSIM, MI, and SAD comparing A and B MC. Each box represents the distribution across all datasets. Gray lines connect paired observation from the same dataset. P-values correspond to paired Wilcoxon signed-rank tests.

Table 3.1 reports a quantitative summary of the comparison between pipelines. In addition to median improvements, the table presents also the relative ΔMAD values (defined as post-MC MAD – pre-MC MAD) quantifying the variation of temporal variability. Negative values indicate a reduction in temporal variability after MC, whereas positive values indicate a slight increase.

Table 3.1: Comparison between motion-correction pipelines A and B on the analysed datasets. Improvements are reported as median post–pre differences for MI, NCC, and SSIM, and as median pre–post differences for SAD. ΔMAD is defined as post–pre for all the metrics.

Metric	Median improvement A	Median improvement B	Median ΔMAD A	Median ΔMAD B
MI	0.1714	0.0748	0.0114	0.0047
NCC	0.0861	0.0553	-0.0078	-0.0001
SSIM	0.0891	0.0408	0.0120	0.0068
SAD	1.5057	0.5052	-0.2730	-0.1518

3.1.2 Similarity metrics on the entire dataset

As it was established that Pipeline A achieved the best performance in the 8 test cases, this pipeline was implemented to analyse the entire dataset.

All similarity metrics computed on the entire dataset showed an improvement after applying the selected pipeline configuration *A*. MI, NCC and SSIM exhibited higher values after MC, while SAD showed a reduction, indicating improved image consistency. Figure 3.2 presents the distribution of pre- and post-MC values for each metric. Boxplots are reported using the median and interquartile range, as these statistics are more robust to outliers compared to mean-based measures.

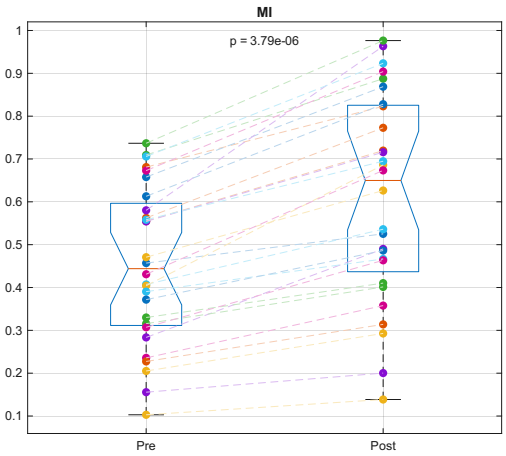

(a) Boxplot of MI.

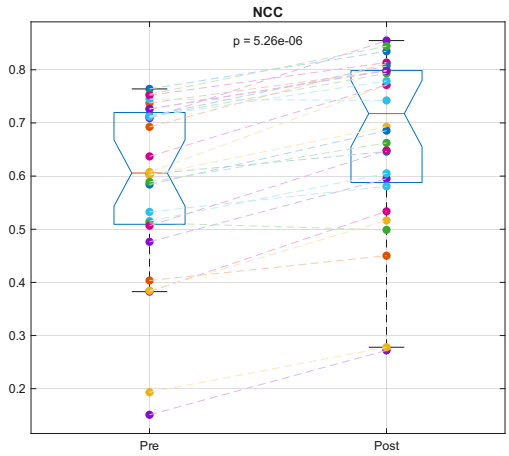

(b) Boxplot of NCC.

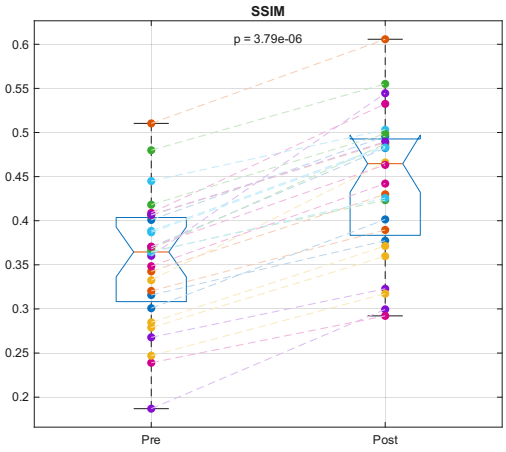

(c) Boxplot of SSIM.

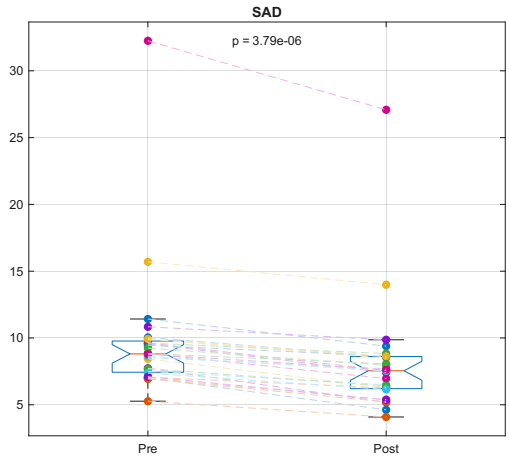

(d) Boxplot of SAD.

Figure 3.2: Comparison between pre- and post-motion correction values for all similarity metrics on the entire dataset. Boxplots illustrate the distribution of each metric, showing the overall improvement after MC.

A quantitative summary of the motion-corrected data on the entire dataset is reported in Table 3.2. The presented values correspond to the median improvement, ΔMAD and p-value from Wilcoxon signed-rank test to compare statistical significance.

Table 3.2: Quantitative effect of the selected motion-correction pipeline on the complete dataset (N = 28). The reported values to the median improvement and ΔMAD for all four metrics.

Metric	Improvement	ΔMAD	p-value
MI	0.1560	0.0483	3.79×10^{-6}
NCC	0.0873	-0.0150	5.26×10^{-6}
SSIM	0.0869	-0.0030	3.79×10^{-6}
SAD	1.3614	0.0831	3.79×10^{-6}

Figure 3.3 shows an example of a TIC extracted from pixel [68,56] in the first contrast-agent injection of the second-day acquisition of RASTRIC patient 47, before (on the left) and after (on the right) MC. The blue curve in the plot represents the raw pixel values, the red curve is the preprocessed TIC used for the analysis, the yellow curve represents the fitted LDRW model and the gray vertical lines represent the OOP frames selected within the MC. Explanation of the LDRW model and the extracted TIC can be found in Section 2.5.1 A visible reduction of signal fluctuations can be observed after MC, together with an increase in the overlap between the measured samples and the fitted curve.

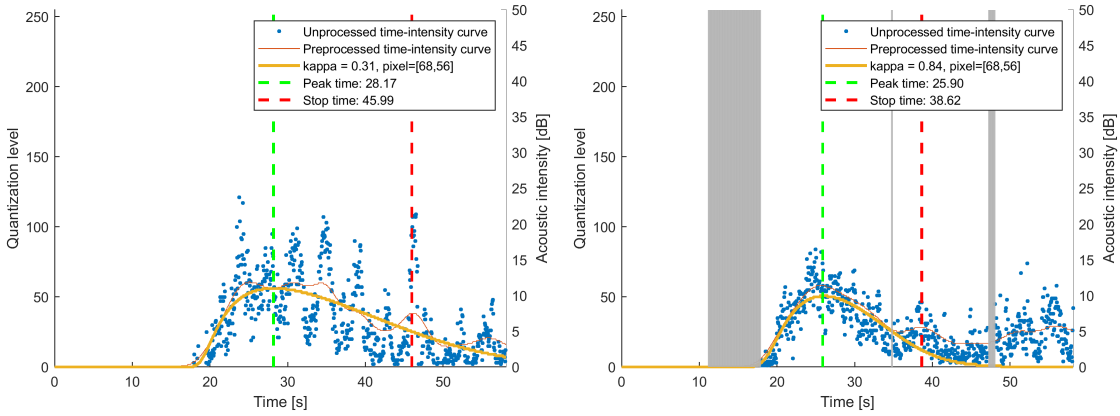

(a) Pre-motion correction. **(b)** Post-motion correction.

Figure 3.3: Example of time-intensity curves extracted from pixel [68,56] of RASTRIC patient 47 before and after motion correction.

3.2 Speckle size regularization

Figure 3.4 illustrates the results of speckle size as a function of depth before and after the regularization process for axial and lateral directions.

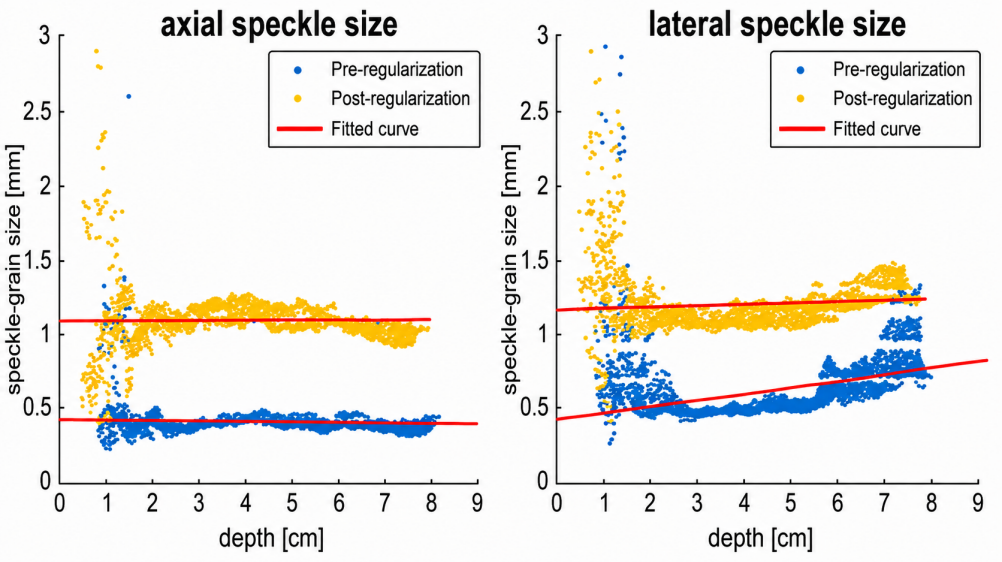

Figure 3.4: Fitted curve on axial (left) and lateral (right) speckle size.

Table 3.3 reports the coefficients of the fitted curves to the axial and lateral speckle size measurements before and after regularization.

Table 3.3: Linear regression coefficients obtained before and after speckle-size regularization for the axial and lateral directions. σ_{const} refers to the intercept, while σ_{var} is the slope of the fitted curve.

Direction	Regularization	σ_{const}	σ_{var}
Axial	Pre	-0.005	0.490
	Post	0.0007	1.128
Lateral	Pre	0.0728	0.503
	Post	0.0068	1.215

Before regularization, the axial speckle size remained approximately independent from depth, as indicated by the almost zero slope of the fitted curve, but the values were around 0.5 mm. After regularization, the axial dimension remains independent of depth but with dimensions around 1 mm. On the contrary, lateral speckle size was highly depth-dependent before regularization, presenting a positive slope, while after regularization the slope value reduces, as depth-depencency does.

Figure 3.5 shows representative 2D autocovariance maps before and after regularization. Prior to the regularization, the autocovariance has a strong lateral

elongation, indicating an anisotropic speckle pattern having a larger lateral dimension (6.04 px) than axial dimension (1.58 px). The regularization makes the autocovariance more symmetric and the speckle sizes more similar (lateral = 3.13 px, axial = 4.09 px) which suggests a more isotropic speckle pattern.

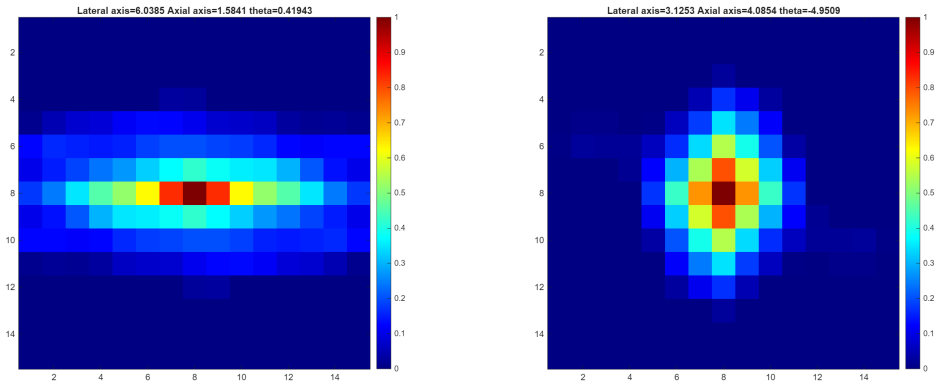

(a) Before regularization. **(b)** After regularization.

Figure 3.5: Examples of 2D autocovariance functions before and after speckle size regularization. After regularization, the autocovariance becomes more symmetric, reflecting a more isotropic speckle pattern.

Figure 3.6 shows the 2D autocovariance patterns obtained from two ROIs at different depths of imaging before and after speckle size. Before regularization, the deeper ROI shows a greater lateral elongation of the autocovariance, due to the depth dependency of the speckle pattern. After regularization, the autocovariance is closer together at different depths, which means the speckle size is isotropic and spatially uniform throughout the FOV.

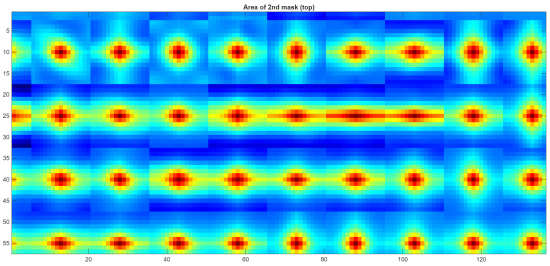

(a) Top ROI before regularization.

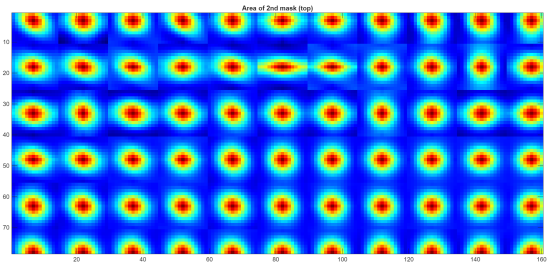

(b) Top ROI after regularization.

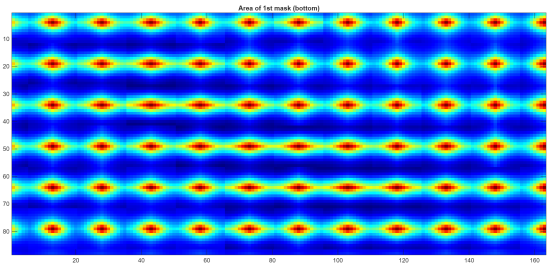

(c) Bottom ROI before regularization.

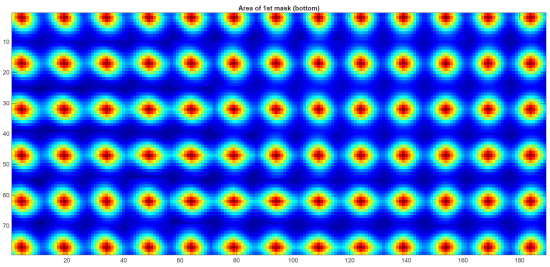

(d) Bottom ROI after regularization.

Figure 3.6: Representative 2D autocovariance patterns extracted from two ROIs located at different imaging depths before and after speckle size regularization.

3.3 Calibration experiment

From the fitting of the pairs (C, I_{lin}), an approximately linear relationship was obtained. In this fitting, the concentration of 10 mg/L was excluded, as the expected linear relationship between concentration and intensity is valid only at low UCA concentrations. The result is shown in Figure 3.7.

Consequently, the sensitivity of the scanner (a_1) and the baseline intensity (a_2) can be found. The estimated parameters are:

$$a_1 = 11.026 \tag{3.1}$$

$$a_2 = -0.1639 \tag{3.2}$$

with a coefficient of determination:

$$R^2 = 0.9569 \tag{3.3}$$

3.4 CUDI-based analysis

The following figures illustrates the spatial distribution and radial behavior of the parameters extracted for tumor analysis and characterization. A representative

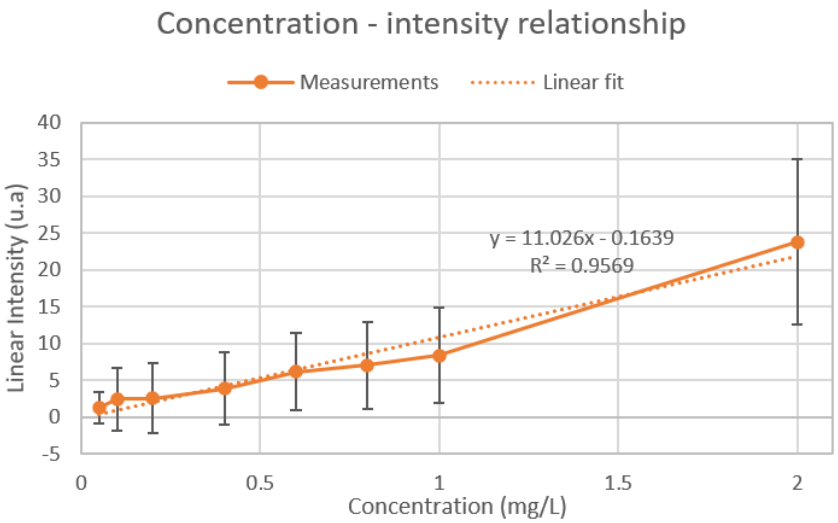

Figure 3.7: Relationship between concentration and intensity. The curve is based on the following concentrations in mg/L: 0.05, 0.1, 0.2, 0.4, 0.6, 0.8, 1, 2.

case is presented before extending the analysis to the entire dataset.

Figures 3.8, 3.9, 3.10 and 3.11 represent the parameters derived from the IDC fitting analysis: respectively κ, peak time, wash-in time, mean transit time. Figures 3.12, 3.13, 3.14, 3.15 and 3.16 show the parameters derived from the spatiotemporal similarity analysis: coherence, correlation coefficient, orthogonal coherence, conditional entropy, mutual information.

For each parameter, a parametric colormap extracted from the CUDI framework is displayed with its corresponding color bar. In addition, two side-by-side plots are reported. For each plot, specific values are shown for each rim of the radial analysis under consideration and for the core. Negative rims correspond to the regions outside the manually delineated tumor contour, with -3 indicating the outermost rim and -1 representing the rim immediately outside the contour. Positive rims correspond to the regions inside the tumor, with 1 indicating the innermost rim closest to the contour, and 3 indicating the innermost rim, just preceding the core. The left plots report the radial profile of the parameter using the median value for each rim (blue continue line), and the 25th and 75th percentiles range (yellow and orange dashed lines). The right plots illustrate the full distribution of parameter values for each rim using boxplots and providing information about variability, distribution shape and outliers.

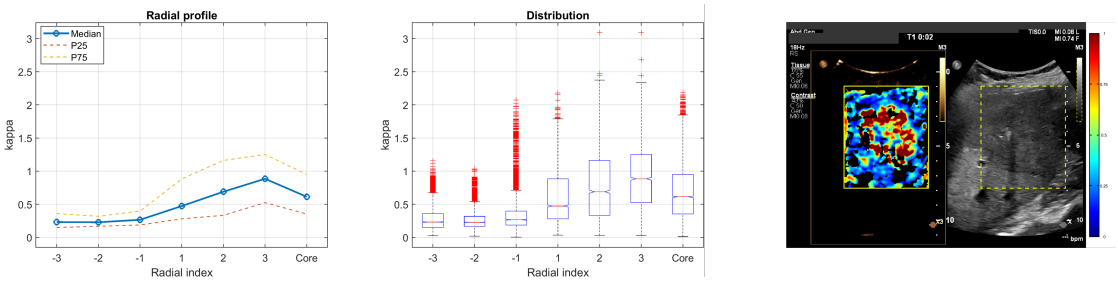

Figure 3.8: Parameter κ. On the left: radial profile. In the centre: distribution. On the right: colormap extracted from CUDI framework.

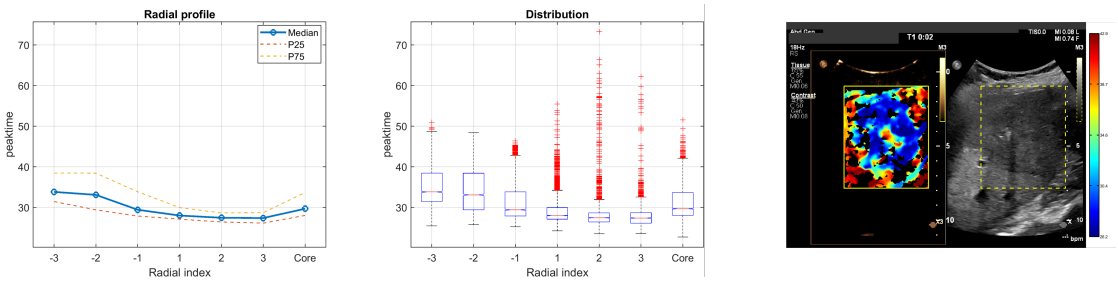

Figure 3.9: Parameter peak time. On the left: radial profile. In the center: distribution. On the right: colormap extracted from CUDI framework.

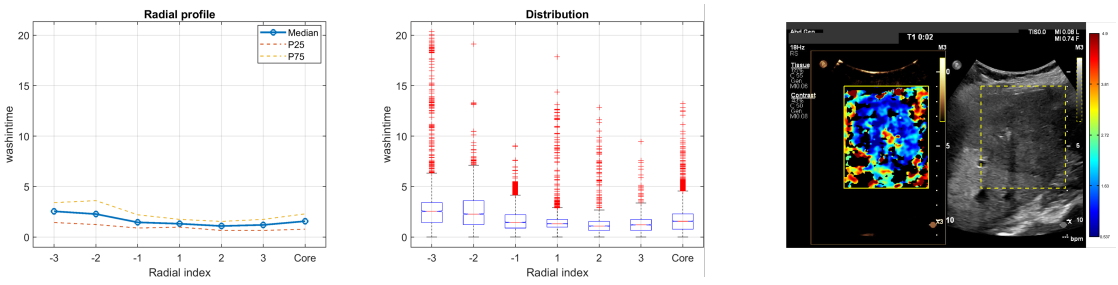

Figure 3.10: Parameter wash-in time. On the left: radial profile. In the center: distribution. On the right: colormap extracted from CUDI framework.

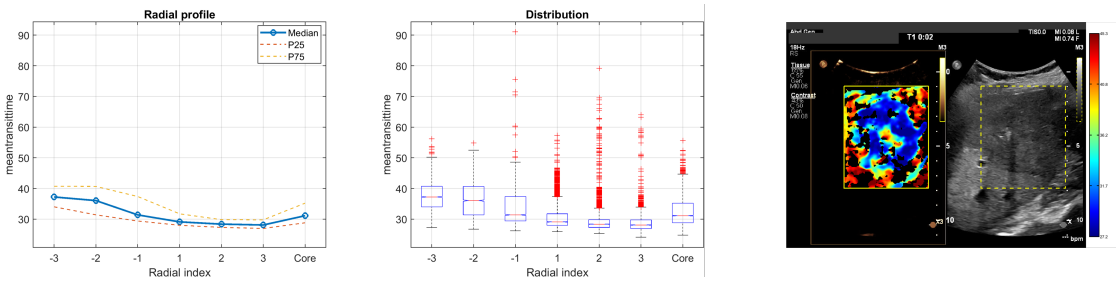

Figure 3.11: Parameter mean transit time. On the left: radial profile. In the center: distribution. On the right: colormap extracted from CUDI framework.

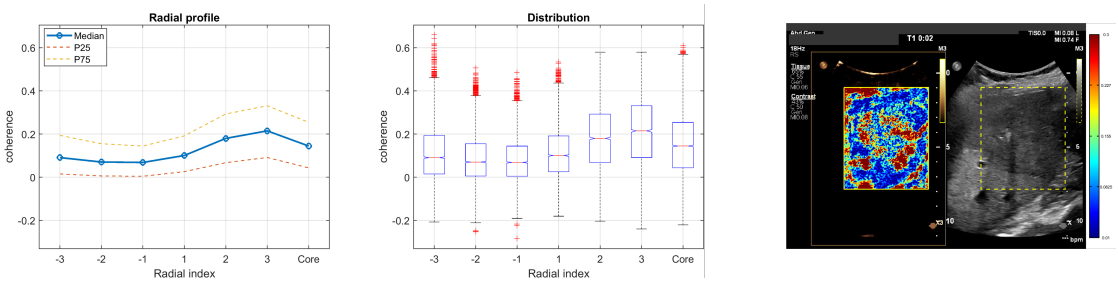

Figure 3.12: Parameter coherence. On the left: radial profile. In the center: distribution. On the right: colormap extracted from CUDI framework.

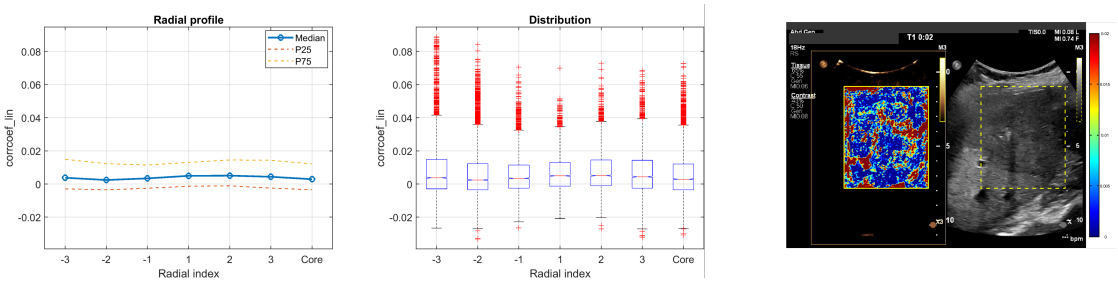

Figure 3.13: Parameter correlation coefficient. On the left: radial profile. In the center: distribution. On the right: colormap extracted from CUDI framework.

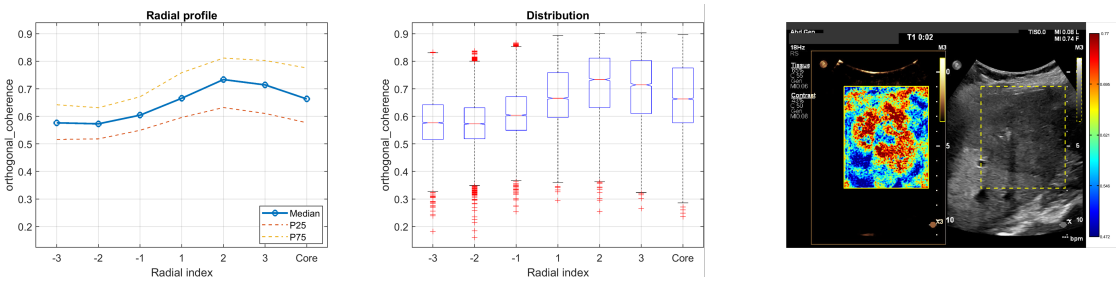

Figure 3.14: Parameter orthogonal coherence. On the left: radial profile. In the center: distribution. On the right: colormap extracted from CUDI framework.

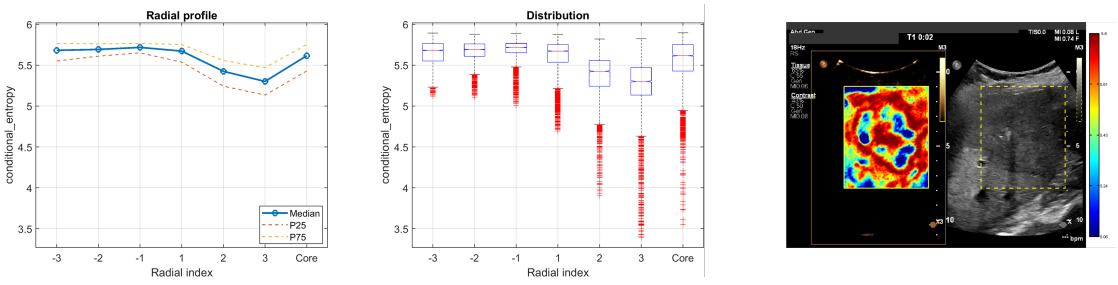

Figure 3.15: Parameter conditional entropy. On the left: radial profile. In the center: distribution. On the right: colormap extracted from CUDI framework.

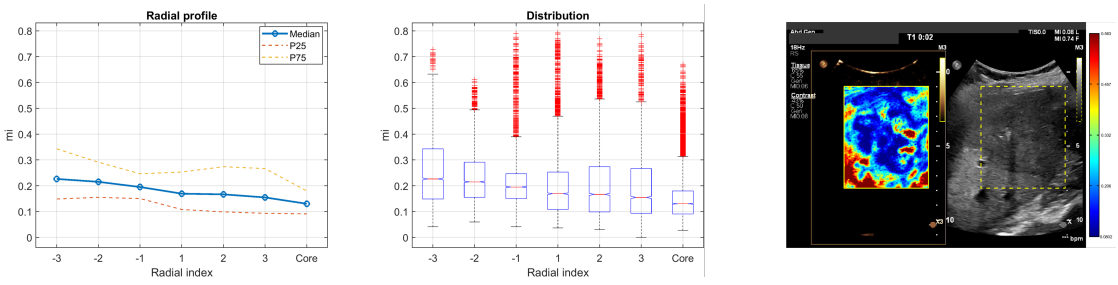

Figure 3.16: Parameter mutual information. On the left: radial profile. In the center: distribution. On the right: colormap extracted from CUDI framework.

Table 3.4 reports the p-values obtained from Wilcoxon signed-rank tests comparing outer, inner and core regions for each parameter. The test was computed considering only datasets satisfying the minimum valid-pixel threshold, as described in Section 2.5.3.

Table 3.4: Statistical comparison between outer, inner, and core regions for all analyzed parameters. P-values were obtained using paired Wilcoxon signed-rank tests. Significant differences ($p < 0.05$) are highlighted in bold and marked with a *. N indicates the number of datasets considered for the analysis.

Parameter	Outer vs Inner	Inner vs Core	Outer vs Core	N
Coherence	0.228	0.849	0.218	26
Conditional entropy	0.713	0.395	0.517	26
Correlation coefficient	0.568	0.585	0.585	26
κ	**0.005***	0.408	**0.003***	22
Mean transit time	**<0.001***	**0.050***	**<0.001***	22
Mutual information	**0.025***	0.144	**0.006***	26
Orthogonal coherence	**<0.001***	**0.049***	**<0.001***	26
Peak time	**0.003***	0.062	**<0.001***	22
Variance	0.077	0.355	**0.016***	22
Wash-in time	0.463	0.845	0.327	18

Among the analysed parameters, parameters such as coherence, correlation coefficient, conditional entropy and wash-in time do not present significant differences between the analysed regions. In contrast, five of them emerge as statistically significant differences in at least two of the three regional comparisons: κ, peak time, mean transit time, orthogonal coherence, mutual information. In particular, mean transit time was the only parameter showing statistically significant differences across all the three regional comparisons.

The distribution of these five parameters is illustrated in Figures 3.17 and 3.18. For each parameter, the boxplots illustrate the distribution of the values for outer, inner and core regions across all the sequences. Colored markers indicate the value associated with each single dataset, while dashed lines connect the values from the same lesion across the three regions. This visualization is helpful to appreciate both population-wide distributions and lesion-specific trends during the analysis.

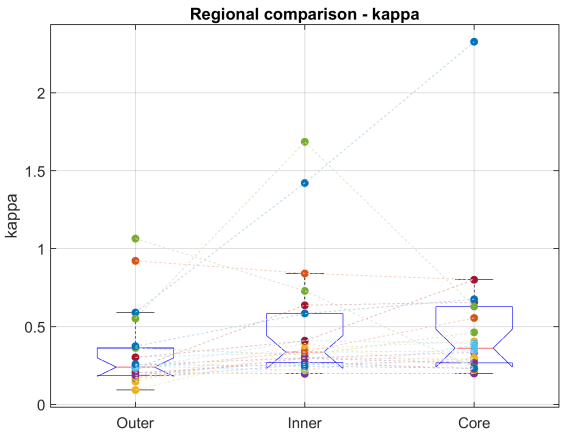

(a) Kappa.

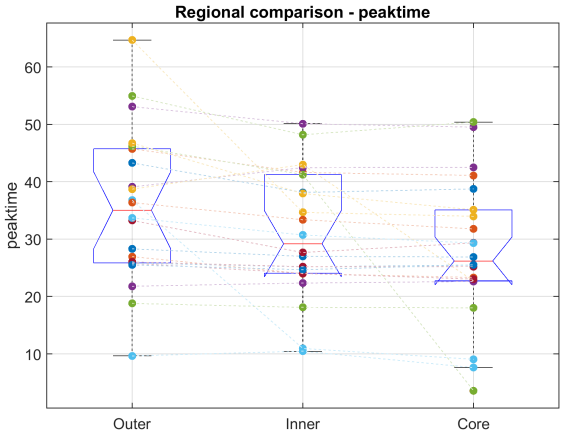

(b) Peak time.

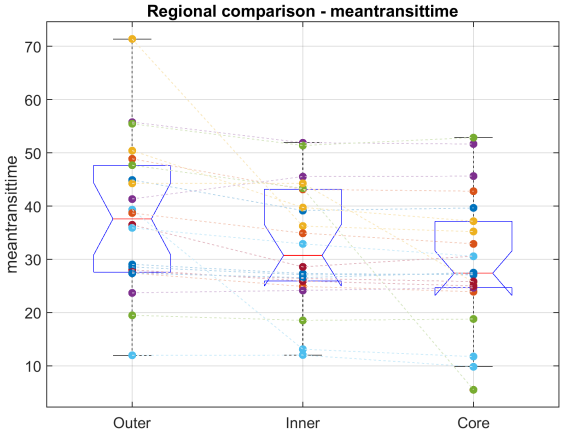

(c) Mean transit time.

Figure 3.17: Regional comparison of IDC fitting parameters. Boxplots summarize the distribution across all datasets, while dashed lines connect values belonging to the same lesion.

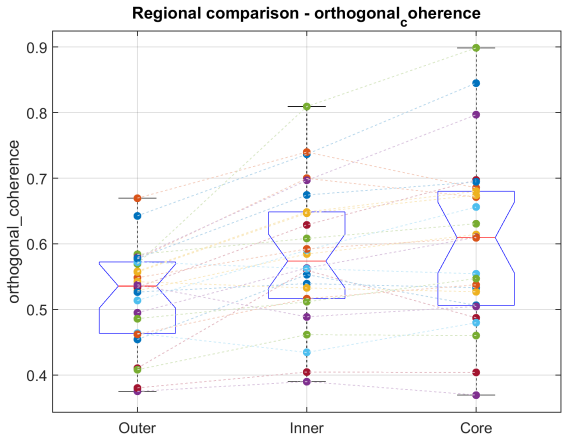

(a) Orthogonal coherence.

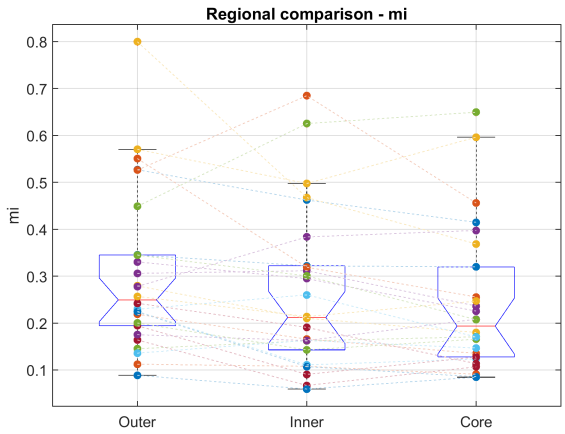

(b) Mutual information.

Figure 3.18: Regional comparison of the most significant spatiotemporal similarity parameters. Boxplots summarize the distribution across all datasets, while dashed lines connect values belonging to the same lesion.

Chapter 4

Discussion

4.1 Motion Correction

4.1.1 Choice of best pipeline

As shown in Section 3.1.1, both pipelines improved similarity metrics in most cases. However, the magnitude and consistency of these improvements differed between the two versions. From the results, it can be observed that the pipeline A gives better improvement than pipeline B with respect to NCC, SSIM and MI. Similarly, a larger decrease for A was seen in the SAD (in which more positive improvement was seen as less SAD), suggesting a more effective reduction of the intensity inconsistencies at the frame-to-frame level. The improvements in NCC and SSIM are larger for the larger, indicating that image similarity and local structural information were better preserved following MC. The greatest differences were seen in MI, indicating a better preservation of the statistical relationships between frames and signal characteristics, which is important for good extraction and analysis of TIC. The larger reduction in SAD further supports a more effective compensation of frame-to-frame intensity inconsistencies.

All metrics had the same p-value < 0.01 when using the paired Wilcoxon signed-rank test. This is because of the small sample size (n = 8) and the fact that the same pipeline consistently performed better than the other pipeline in all of the datasets. Thereby, all the paired differences were positive and the two-sided p-value for eight paired differences was the smallest possible value obtainable with the Wilcoxon signed-rank test.

Changes in temporal variability quantified through ΔMAD, are dependent from metrics. Pipeline A shows a larger reduction in variability for NCC and SAD, while MI and SSIM present small increases in ΔMAD for both pipelines. In addition, some failure cases are observed for B, where improvements are minimal or even negative.

These findings suggests that the superior performance of pipeline A stands in the stage at which the two-tier strategy is applied. Rigid registration is primarily intended to compensate for large-scale global displacements, whereas deformable registration is more suitable for correcting residual local deformations. In pipeline A, frames are first aligned to motion-consistent SRFs, which which reduces inter-cycle variability and gives a more stable initial alignment for the next deformable registration. On the other hand, if rigid registration is done directly towards the MRF, as is the case in pipeline B, it may lead to larger registration errors, which the deformable registration stage has to compensate for in the following pipeline step.

Furthermore, rather than using a standard spatial registration, deformable registration is used directly between each frame and the MRF in pipeline A to ensure that they share a common spatial reference throughout the sequence. Pipeline B propagates deformable transformations through intermediate SRFs, which leads to higher chances of local deformations being accumulated and causing progressive spatial distortions.

This interpretation is in line with the initial concept of the two-tier approach proposed by Ta et al., which incorporated intermediate respiratory references, to better ensure the robustness of the registration when long CEUS acquisitions and breathing motion are present. In previous CEUS motion correction studies similar observations have been made, with intermediate reference frames providing more stable registrations than a direct alignment to a single reference image, taking into account the respiratory cycles.

The evaluation was performed on a limited dataset and ROI configuration. Therefore, even if these results indicate A as the better pipeline, more validation on a larger and more diverse set of data would be needed to establish generalizability of the findings. Based on these considerations, A was chosen for further analysis, because it not only shows the highest improvements for all similarity measures but also a similar or slightly better temporal stability.

4.1.2 Similarity metrics on the entire dataset

The boxplots in Section 3.1.2 reveal that in all of the evaluated metrics there is a general improvement of the image sequences after applying MC, demonstrating the efficienty if the pipeline for improving temporal and spatial consistency of the image sequence.

In the post MC sequence, MI had higher shared information between frames and a greater temporal coherence of the signal, meaning that the corresponding anatomical structures aligned better over time. Similarly, the higher NCC indicates more stability in pixel intensities across frames, therefore more reliability in spatial correspondence of anatomical features throughout the sequence. These results

are reinforced by the improvement in SSIM. Differently from NCC, SSIM also preserves local structural charcateristics, such as contrast and texture. Thus, its rise indicated that in addition to the intensity patterns also the structure of images is better preserved across frames. This demonstrates higher spatial coherence and lower motion-induced distortion. The reduction in SAD confirms a decrease in spatial misalignment and increase in stability after MC in the sequence.

These results are further strengthened by the use of median-based statistics, where the observed improvements are not due to a few frames that are affected by extreme motion, but are instead observed across the entire sequence of frames.

While MC generally leads to an overall improvement in similarity metrics, the behavior of the MAD is slightly different. As shown in Table 3.2, MI and SAD present positive ΔMAD values, indicating an increase in metric dispersion after MC. On the other hand, NCC and SSIM report small negative values, corresponding to a reduction in variability metric. These small changes need some caution in the interpretation. In particular, the higher ΔMAD value of the MI may be attributed to the fact that MC restores spatial consistency and reduces temporal averaging effects. By reducing this motion-induced blurring, local differences between pixels become more distinct and well-defined, leading to an increased dispersion of the metric values. Consequently, a rise in MAD should not be considered as a sign of bad MC, but rather as an indicator of improved spatial consistency and reduced motion-induced smoothing caused by tissue displacement. In contrast, the reduction in variability for NCC and SSIM suggests a more homogeneous improvement across the analysed dataset.

Overall, regardless the limited changes in ΔMAS, these results demonstrate that the implemented MC pipeline successfully improves image stability by enhancing alignment, intensity-based and structural consistency. This is particular relevant for subsequent CUDI-based analysis, since greater temporal coherence of pixel-wise signals allows TICs to be extracted and fitted more reliably, ultimately enabling a more robust assessment of spatiotemporal similarity and tissue heterogeneity. This improvement in TIC fitting can be seen in Figure 3.3. Before correction, the raw intensity pixel values exhibit clear oscillations. These variations recall the breathing movement and tissue displacement, which cause the tracked pixel to take place into different positions over time. After MC, the intensity pixel values appear more compact and better aligned with the TIC shape, The oscillation movement is reduced, sign of the effectiveness of the implemented MC pipeline. This create a smoother and more coherent temporal evolution of the signal. Consequently, the fitted LDRW model is closer to the measured data. In this way, a more reliable estimation of perfusion-related parameters can be carried out.

This result supports the quantitative results obtained from the similarity metrics and underlines the importance for MC to enable proceeding with subsequently analysis in a more reliable and robust manner. In this way, temporal consistency

of pixel-wise signals and characterization of tissue perfusion and heterogeneity are enhanced.

4.2 Speckle Size Regularization

As expected, before regularization the speckle pattern presented a clear anisotropy, characterized by lateral elongation depth-dependent and axial dimension independent of depth, as demonstrated by the near-zero slope of the fitted model in Figure 3.4 and Table 3.3. However, the measured dimensions were below the desired value of 1 mm. The observed behavior is consistent with the physical principle that CEUS signals originate from the collective nonlinear backscatter of microbubble populations rather than isolated reflectors. The depth-dependent elongation of speckle patterns can be explained by the divergence of ultrasound beams emitted by the convex probe, leading to an increase in lateral resolution with depth and therefore to intrinsic anisotropy of the speckle field [20].

After regularization, the axial size remained independent of depth while approaching the target dimension. In contrast, lateral speckle size exhibited a reduction in the slope of the fitted curve, indicating a more homogeneous lateral speckle size throughout the FOV.

Additional proof is illustrated in Figure 3.5. Before regularization, the estimated speckle shape is wrongly anisotropic, showing a larger lateral size than axial size. After regularization, the dimensions of the two axis become similar, resulting in a closer to isotropic speckle pattern. This observation is consistent with Figure 3.6, where visual examples extracted from near and far field ROIs show more uniform speckle size after regularization.

From CUDI perspective, these results are very important. Depth-dependent variations in speckle size may introduce artificial differences in local image texture and affect the subsequent analysis of similarity-based parameters. By reducing lateral depth-dependency and generating a more homogeneous speckle pattern, the regularization procedure improves the comparability of measurements obtained at different image depths. Moreover, it was proved that the reliability of analysis of similarity in space and time is greatly influenced by the regularity of the speckle pattern [19] [20].

Overall, the results show that speckle size regularization framework is able to decrease the depth-dependent anisotropy of the speckle field while maintaining spatial consistency. This achieves a more uniform representation of local image texture, and creates more reliable basis for the subsequent extraction of dispersion-related parameters as part of the CUDI framework.

4.3 Calibration experiment

The results indicate that, even if the exact form of the scanner's compression function is difficult to identify, using simplified modeling assumptions the linear relationship between concentration and linearized intensity is well preserved. The high coefficient of determination $R^2 \approx 0.96$ confirms that the selected model provides a good approximation of the measured data across the entire analysed concentration range. A marginal increase in variability at higher concentration can be observed, as visible with larger error bars. This behavior is aligned with the documented nonlinear response of UCAs: at higher concentrations, physical phenomena such as bubble-to-bubble interactions and acoustic attenuation begin to interfere with signal linearity and affect it.

These findings support the validity of adopting a simplified linear model to describe the relationship between microbubble concentration and acoustic intensity within the low-concentration regime, which is an interesting area for clinical CEUS applications. Therefore, within this regime, the relationship between UCA concentration and linearized intensity can be reasonably approximated as linear. This justifies the use of a linear model in subsequent CUDI analysis and related quantitative approaches.

4.4 CUDI-based analysis

The representative case illustrated from Figure 3.8 to Figure 3.16 highlights the effectiveness of the radial analysis to reveal spatial variations in both perfusion-related and spatiotemporal similarity parameters. Several parameters exhibited non-uniform radial profiles, sign of spatial heterogeneity between the outer, inner and core of the tumor. Even if the radial behavior observed in this case cannot be generalized to the entire dataset, this example demonstrates how the proposed radial analysis captures spatial patterns that would not be visible when averaging parameter values over the entire lesion.

The global analysis on all the datasets identified several parameters exhibiting statistically significant differences between the compared regions. In particular, κ, mean transit time, peak time, mutual information and orthogonal coherence exhibit significant differences in at least two of the three regional comparisons, suggesting the presence of reproducible spatial patterns across lesions.

The dispersion-related parameter κ shows significant differences between the outer and inner regions and between the outer region and the tumor core. Since κ is related to the dispersion characteristics of the local TIC, these findings may suggest that UCA transport dynamics differ between the tumor and the surrounding tissue. The absence of significant differences between the inner region and the

core indicates that this transition may occur primarily at the tumor contour. Peak time and mean transit time show some of the highest differences. Mean transit time is the only IDC fitting parameter that reaches p-values < 0.05 in all three regional comparisons. These parameters indicate a progressive increase from the surrounding tissue towards the tumor core, suggesting that the temporal characteristics of contrast enhancement vary clearly across the different spatial regions and reflecting heterogeneous perfusion dynamics within and around the lesion.

Regarding the spatiotemporal similarity metrics, the largest regional differences in measures are observed in the orthogonal coherence measure with significant differences (p-value < 0.05) in all comparisons. In particular, differences in the outer region indicate that the temporal behaviour of neighboring TICs is significantly different from the outer tissue and the tumour.This result supports the hypothesis that orthogonal coherence is able to capture spatial heterogeneity in local contrast enhancement patterns, giving complementary information to conventional perfusion parameters. Mutual information also presents significant differences between the outer region and both the inner and the core zones of the lesion. Since mutual information quantifies the statistical depencency and the amount of shared information between neighboring TICs, these findings suggest that the temporal organization of contrast enhancement differs when moving from outer to tumoral tissue.

In contrast, coherence, correlation coefficient, conditional entropy and wash-in time do not exhibit statistically significant regional variations. This absence of significant differences may imply that these parameters are either less sensitive to localized changes within the lesion or influenced by inter-patient variability.

One possible biological explanation for these regional differences may be the heterogeneous vascular structure that is often encountered in the liver metastases. At the tumor boundary, the activation of angiogenesis has been linked to tumor growth, because the tumor cells come into contact with the nearby liver parenchyma that stimulates the formation of new microvessels. On the other hand, the central regions of metastases may present lower vascular density, hypoxia or necrotic regions. This results in an outer, inner and core zone of the lesion that can have very different contrast agent transport and perfusion characteristics. These differences in vascular architecture may explain the observed differences in κ, peak time and mean transit time, which are influenced by the local contrast transport and perfusion characteristics. In the same way, variations in orthogonal coherence and mutual information between regions could be an indicator of variations in spatial organization and connectivity of the underlying microvascular network.

The observations on the extracted parameters are aligned with previous CUDI studies, which demonstrated that dispersion-related parameters and local spatiotemporal similarity metrics are sensitive to angiogenesis-induced alterations

in microvascular architecture, including changes in vessel density, tortuosity and branching patterns.

Among the dispersion parameters, κ is always reported as one of the most informative parameters. Studies on prostate cancer reported significant differences in κ between benign and malignant tissue, which indicates that the parameter is sensitive to the vessel tortuosity, the microvascular density and the organization of the vessels that are linked to angiogenesis [19]. The results agree with the idea that κ is not only a measure of the intensity of perfusion but also a measure of the spatial variability of the contrast transport dynamics. Similarly, Schmit et al. found that the parameter κ is one of the most discriminative parameters for the characterization of uterine disorders [31]. Their findings showed that the dispersion-related parameters offer additional information to the conventional perfusion parameters derived from TIC and that structural vascular differences, which are not fully accounted for using perfusion kinetics, can be partly explained using dispersion-related measures.

Other parameters based on coherence have also been found to be highly vascular architecture dependent. The prostate cancer datasets showed that the coherence-related metrics also had better discriminatory power than several traditional perfusion parameters, suggesting that temporal similarity between adjacent TICs holds significant information about the underlying microvascular network [19]. These findings agree with the significant differences found in the present study for orthogonal coherence and suggest that coherence-based approaches are sensitive to the vascular heterogeneity.

Different conclusions for conditional entropy have been reported. In studies of breast lesions, the researchers found that conditional entropy was one of the most discriminatory similarity parameters describing the spatiotemporal behavior of tumors, with the temporal contrast behavior of malignant vascular networks being less predictable than that of benign tissue [30]. In contrast, in the present study no significant regional differences were found in conditional entropy. This may indicate that the parameter is more suitable for differentiating the various types of lesions rather than for distinguishing heterogeneity in a particular lesion.

CUDI was also applied to the case of adenomyosis, where both κ and spatiotemporal similarity metrics presented heterogeneously high values when compared to healthy myometrial tissue. These findings were interpreted as evidence of angiogenesis-related alterations in microvascular architecture and local dispersion within the tissue, showing that CUDI can capture spatially heterogeneous microvascular patterns in diseased tissue [32].

Taken together, these findings support and extend the observations previously reported in the CUDI literature. Significant regional differences were observed for several dispersion- and similarity-related parameters, indicating that CUDI can be used not only to characterize vascular differences between tissues but also to

investigate regional variability within individual liver metastases. In particular, the significant differences involving the outer region suggest that important vascular changes may occur at the tumor boundary.

By separately evaluating outer, inner, and core tumor regions, the proposed radial framework directly addresses the spatial analysis concept previously suggested in the CUDI literature. The observed differences indicate that relevant information may be lost when parameters are averaged over the entire lesion.

It is important to consider that this work was relatively limited in size. Furthermore, the number of valid datasets varied between parameters due to the minimum pixel threshold implemented to ensure statistical robustness. Consequently, the observed trends should be considered exploratory and require further validation in larger datasets.

Ultimately, the proposed radial analysis framework successfully reveals spatially heterogeneous patterns difficult to appreciate with traditional approaches. The significant differences observed for kappa, mean transit time, peak time, orthogonal coherence, and mutual information suggest that both perfusion-related and spatiotemporal similarity parameters contain complementary relevant information for liver metastases regional characterization.

Overall, while previous CUDI studies mainly demonstrated the ability of dispersion-based parameters to distinguish tissues with different vascular architectures, the present work suggests that the same framework can also reveal spatial heterogeneity within a single lesion, extending the application of CUDI to liver metastases. The proposed radial framework allows the separate evaluation of outer, inner, and core tumor regions, highlighting the potential of spatially resolved vascular characterization and extending the application of CUDI beyond conventional lesion-averaged analysis.

Chapter 5

Conclusions

This thesis aimed to investigate the use of contrast enhanced ultrasound (CEUS) and contrast ultrasound dispersion imaging (CUDI) as new quantitative tools to characterize the microvascular heterogeneity of liver metastases.

For this purpose, a full processing pipeline was developed. Proposed framework consisted of motion compensation, CEUS signal calibration, speckle size regularization, extraction of the perfusion related parameter and the spatiotemporal similarity parameter using CUDI, and assessment of these parameters by a radial analysis of the lesions.

Particular attention was devoted to the challenges associated with clinical CEUS acquisitions, which are often affected by respiratory and probe-induced motion. A number of motion compensation methods were tested and compared. The final framework was able to incorporate both rigid and deformable registration and led to better alignment of CEUS sequences, allowing accurate extraction of time–intensity curves and calculations of subsequent quantitative parameters. The results demonstrate that, for quantitative CEUS analysis, liver acquisitions in clinics will require strong motion compensation.

The analyses performed were shown to be feasible for clinical CEUS examinations of liver metastases using CUDI. The reliable extraction of time–intensity curves was possible thanks to the motion correction, and the calibration and speckle size regularization ensured consistency with the assumptions of the CUDI framework.

The radial analysis showed that there are substantial regional variations for some parameters such as the dispersion-related parameter κ, the mean transit time, the peak time, the orthogonal coherence, and the mutual information. Based on these findings, it seems that both the perfusion-related and the spatiotemporal similarity parameters have useful and complementary information regarding the intratumoral vascular heterogeneity. Specifically, the differences between outer, inner and core regions indicate that information of interest may be lost when the parameters are averaged across the entire lesion.

The results are aligned with previous CUDI studies, where dispersion and similarity parameters have been shown to be sensitive to microvascular changes related to angiogenesis. Most previous studies have examined the difference in tissue vascular architecture between different tissues but in the present work the same framework was used to explore regional variation within individual liver metastases.

In conclusion, this work showed the importance of reliable motion compensation for quantitative analysis of CEUS, that CEUS radial CUDI analysis can give relevant information about intratumoral vascular heterogeneity and that both the spatiotemporal similarity parameter and the perfusion-related similarity parameter offer complementary information in the characterisation of liver metastases.

However, several limitations should be noted with these results. There were relatively few datasets available, and the number of valid observations varied for each parameter, due to quality requirements of the analysis pipeline. The present results, therefore, must be regarded as preliminary and need confirmation in a larger patient population.

The methodological and clinical implications are recommended for further studies. In terms of methodology, the motion correction framework could be enhanced by minimizing user dependency and maximizing automation. The statistical threshold method employed for quality control and out-of-plane motion detection could be improved. Moreover, the clustering algorithm applied to identify similar respiratory states requires fairly regular breathing cycles and might not be optimal when breathing cycles are asymmetric or irregular. Respiratory variability could also be accounted for in more advanced approaches to increase the robustness of the registration and consistency in the TIC.

Larger patient cohorts are needed from a clinical perspective to examine the reproducibility and generalizability of the identified perfusion related and spatiotemporal similarity parameters. The clinical relevance of these quantitative descriptors becomes apparent when considering how they may be correlated with the delivery of therapeutic drugs, therapeutic responses, and patient outcomes. It is therefore important to explore the correlation between CEUS-based measures of vascular heterogeneity and intratumoral drug concentration and clinical outcome to therapy in future studies. In this way, future development of quantitative CEUS biomarkers for liver metastases characterization and the development of personalized treatment strategies could be assessed.

Bibliography

[1] Kevin Martin. «Introduction to B-mode imaging». In: *Diagnostic Ultrasound: Physics and Equipment*. Cambridge University Press. 2010, pp. 1–3 (cit. on pp. 2, 3).

[2] Kevin Martin and Kumar Ramnarine. «Physics». In: *Diagnostic Ultrasound: Physics and Equipment*. Cambridge University Press. 2010, pp. 4–22 (cit. on p. 2).

[3] Adham N. Abou Ali, Alessandra Fittipaldi, Joao Rocha-Neves, Barbara Ruaro, Filippo Benedetto, Zeina Al Ghadban, Giulia Simon, Sandro Lepidi, and Mario D'Oria. «Clinical applications of contrast-enhanced ultrasound in vascular surgery: State-of-the-art narrative and pictorial review». In: *JVS-Vascular Insights* 3 (2025), p. 100254. ISSN: 2949-9127. DOI: https://doi.org/10.1016/j.jvsvi.2025.100254. URL: https://www.sciencedirect.com/science/article/pii/S2949912725000716 (cit. on p. 5).

[4] {Sherif R.Z.} Abdel-Misih and Mark Bloomston. «Liver Anatomy». English. In: *Surgical Clinics of North America* 90.4 (Aug. 2010), pp. 643–653. ISSN: 0039-6109. DOI: 10.1016/j.suc.2010.04.017 (cit. on p. 6).

[5] Wenchao Xu, Jia Xu, Jianzhou Liu, Nanzhou Wang, Li Zhou, and Junchao Guo. «Liver Metastasis in Cancer: Molecular Mechanisms and Management». In: *MedComm* 6 (Feb. 2025). DOI: 10.1002/mco2.70119 (cit. on p. 6).

[6] Elijah Trefts, Maureen Gannon, and David Wasserman. «The liver». In: *Current Biology* 27 (Nov. 2017), R1147–R1151. DOI: 10.1016/j.cub.2017.09.019 (cit. on p. 6).

[7] Zavos Chris. *Liver: Anatomy, Function, Common Diseases, and Foods*. Accessed: 30 May 2026. 2025. URL: https://peptiko.gr/en/liver-anatomy-function-common-diseases-and-foods/ (cit. on pp. 6, 7).

[8] Wenchao Xu, Jia Xu, Jianzhou Liu, Nanzhou Wang, Li Zhou, and Junchao Guo. «Liver Metastasis in Cancer: Molecular Mechanisms and Management». In: *MedComm* 6 (Feb. 2025). DOI: 10.1002/mco2.70119 (cit. on pp. 7, 8).

[9] Mirela Danila, Alina Popescu, Roxana Șirli, Ioan Sporea, Alina Martie, and Mădălina Sendroiu. «Contrast enhanced ultrasound (CEUS) in the evaluation of liver metastases.» In: *Medical ultrasonography* 12 3 (2010), pp. 233–7. URL: https://api.semanticscholar.org/CorpusID:19793148 (cit. on pp. 7–9).

[10] WebMD Editorial Contributors. *What Is a Hypoechoic Mass?* Medically reviewed by Laura J. Martin, MD. Accessed: April 30, 2026. 2025. URL: https://www.webmd.com/cancer/what-is-hypoechoic-mass (cit. on p. 8).

[11] Pavel Taimr, Vivian Jongerius, Chulja Pek, Nanda Krak, Bettina Hansen, Harry Janssen, Herold Metselaar, and Casper Eijck. «Liver Contrast-Enhanced Ultrasound Improves Detection of Liver Metastases in Patients with Pancreatic or Periampullary Cancer». In: *Ultrasound in medicine & biology* 41 (Oct. 2015). DOI: 10.1016/j.ultrasmedbio.2015.06.019 (cit. on p. 8).

[12] Ehsan Safai Zadeh, Helmut Prosch, Ahmed Ba'ssalamah, Hajo Findeisen, Amjad Alhyari, Nils Raab, and Christian Görg. «Contrast-enhanced ultrasound of the liver: basics and interpretation of common focal lesions». In: *RöFo - Fortschritte auf dem Gebiet der Röntgenstrahlen und der bildgebenden Verfahren* 196 (Jan. 2024). DOI: 10.1055/a-2219-4726 (cit. on pp. 8, 9).

[13] Lars Peter Larsen. «Role of contrast enhanced ultrasonography in the assessment of hepatic metastases: A review». In: *World journal of hepatology* 2 (Jan. 2010), pp. 8–15. DOI: 10.4254/wjh.v2.i1.8 (cit. on p. 9).

[14] B. Müllhaupt et al. «Assessment of focal liver lesions in non-cirrhotic liver – expert opinion statement by the Swiss Association for the Study of the Liver and the Swiss Society of Gastroenterology». In: *Swiss Medical Weekly* 153.9 (2023). Accessed: 2026-04-30, p. 40099. DOI: 10.57187/smw.2023.40099. URL: https://smw.ch (cit. on pp. 9, 10).

[15] Thodsawit Tiyarattanachai, Simona Turco, John R. Eisenbrey, Corinne E. Wessner, Alexandra Medellin-Kowalewski, Stephanie Wilson, Andrej Lyshchik, Aya Kamaya, and Ahmed El Kaffas. «A Comprehensive Motion Compensation Method for In-Plane and Out-of-Plane Motion in Dynamic Contrast-Enhanced Ultrasound of Focal Liver Lesions». In: *Ultrasound in Medicine & Biology* 48.11 (2022), pp. 2217–2228. ISSN: 0301-5629. DOI: https://doi.org/10.1016/j.ultrasmedbio.2022.06.007. URL: https://www.sciencedirect.com/science/article/pii/S0301562922004409 (cit. on pp. 11, 12).

[16] Jyotirmoy Banerjee, Yuanyuan Sun, Camiel Klink, Renske Gahrmann, Wiro J. Niessen, Adriaan Moelker, and Theo van Walsum. «Multiple-correlation similarity for block-matching based fast CT to ultrasound registration in liver interventions». In: *Medical Image Analysis* 53 (2019), pp. 132–141. ISSN: 1361-8415. DOI: https://doi.org/10.1016/j.media.2019.02.003. URL: https:

//www.sciencedirect.com/science/article/pii/S1361841518305334 (cit. on p. 11).

[17] Ji Zhang, Yanrong Zhang, Juan Chen, Gonghao Ling, Xiangyu Wang, and Haibo Xu. «Respiratory motion correction for liver contrast-enhanced ultrasound by automatic selection of a reference image». In: *Medical Physics* 46.11 (2019), pp. 4992–5001. DOI: https://doi.org/10.1002/mp.13776. eprint: https://aapm.onlinelibrary.wiley.com/doi/pdf/10.1002/mp.13776. URL: https://aapm.onlinelibrary.wiley.com/doi/abs/10.1002/mp.13776 (cit. on pp. 11, 22, 23).

[18] Casey N. Ta, Mohammad Eghtedari, Robert F. Mattrey, Yuko Kono, and Andrew C. Kummel. «2-Tier In-Plane Motion Correction and Out-of-Plane Motion Filtering for Contrast-Enhanced Ultrasound». In: *Investigative Radiology* 49 (2014), pp. 707–719. URL: https://api.semanticscholar.org/CorpusID:18899059 (cit. on pp. 12, 19, 22, 23, 27).

[19] M.P.J. Kuenen. «Contrast-ultrasound dispersion imaging for prostate cancer localization». English. Phd Thesis 1 (Research TU/e / Graduation TU/e). Electrical Engineering, 2014. ISBN: 978-90-386-3568-2. DOI: 10.6100/IR769554 (cit. on pp. 12, 13, 38, 40, 41, 43, 45, 67, 70).

[20] Ferenc Kandi, Catarina {Dinis Fernandes}, Simona Turco, Eva {de Bock}, {Lynda J.M.} Juffermans, {Judith A.F.} Huirne, and Massimo Mischi. «2D Contrast-Ultrasound Dispersion Imaging of Angiogenesis in Adenomyosis: First Experimental Measurements». English. In: *IEEE International Workshop on Medical Measurement and Applications, MEMEA 2024*. 2024 IEEE International Symposium on Medical Measurements and Applications, MeMeA 2024, MeMeA 2024 ; Conference date: 26-06-2024 Through 28-06-2024. United States: Institute of Electrical and Electronics Engineers, July 2024. DOI: 10.1109/MeMeA60663.2024.10596726. URL: https://memea2024.ieee-ims.org/ (cit. on pp. 13, 30, 31, 34, 67).

[21] The MathWorks, Inc. *MATLAB*. Version R2025b. The MathWorks, Inc. Natick, Massachusetts, USA, 2025. URL: https://mathworks.com (cit. on p. 14).

[22] The MathWorks, Inc. *MATLAB*. Version R2022a. The MathWorks, Inc. Natick, Massachusetts, USA, 2022. URL: https://mathworks.com (cit. on p. 14).

[23] Ian T. Jolliffe and Jorge Cadima. «Principal component analysis: a review and recent developments». In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374.2065 (Apr. 2016), p. 20150202. ISSN: 1364-503X. DOI: 10.1098/rsta.2015.0202. eprint: https://royalsocietypublishing.org/rsta/article-pdf/doi/10.1098/rsta.2015.0202/1381479/rsta.2015.0202.pdf. URL: https://doi.org/10.1098/rsta.2015.0202 (cit. on p. 25).

[24] S. Klein, M. Staring, K. Murphy, M.A. Viergever, and J.P.W. Pluim. «Elastix : a toolbox for intensity-based medical image registration». English. In: *IEEE Transactions on Medical Imaging* 29.1 (2010), pp. 196–205. ISSN: 0278-0062. DOI: 10.1109/TMI.2009.2035616 (cit. on p. 26).

[25] Global Ultrasound Institute. *Mechanical index (MI)*. https://globalultrasoundinstitute.com/pocuspedia/mechanical-index-mi/. Consulted on 20th June 2026. n.d. (Cit. on p. 31).

[26] C. U. Lee, G. K. Hesley, T. A. Pierson, R. L. Higgins, and M. W. Urban. «Breast ultrasound knobology and the knobology of twinkling for marker detection». In: *Translational Breast Cancer Research* 5 (Oct. 2024). PMID: 39534581; PMCID: PMC11557156, p. 28. DOI: 10.21037/tbcr-24-30 (cit. on p. 32).

[27] De Guo, Hong Yang, Cong Zhang, and Dong Liu. «Optimization of Ultrasonic Imaging Using Persistence and Filter Technology». In: *Advanced Materials Research* 716 (July 2013), pp. 542–547. DOI: 10.4028/www.scientific.net/AMR.716.542 (cit. on p. 32).

[28] Jeff Powers and Frederick Kremkau. «Medical ultrasound systems». In: *Interface focus* 1 (May 2011), pp. 477–89. DOI: 10.1098/rsfs.2011.0027 (cit. on p. 32).

[29] Massimo Mischi, Erik Korsten, and Ton Kalker. «Videodensitometric Methods for Cardiac Output Measurements». In: *EURASIP Journal on Advances in Signal Processing* 2003 (Apr. 2003). DOI: 10.1155/S1687617203211185 (cit. on pp. 34, 37).

[30] C. Chen, S. Turco, P. Kapetas, et al. «Spatiotemporal analysis of contrast-enhanced ultrasound for differentiating between malignant and benign breast lesions». In: *European Radiology* 34.7 (2024), pp. 4764–4773. DOI: 10.1007/s00330-023-10500-x. URL: https://doi.org/10.1007/s00330-023-10500-x (cit. on pp. 45, 70).

[31] C. M. Schmit, B. Stoelinga, P. de Jong, P. M. van de Ven, H. A. M. Brölmann, and J. A. F. Huirne. «Contrast-Enhanced Ultrasound Imaging of Uterine Disorders: A Systematic Review». In: *Journal of Ultrasound in Medicine* 40.12 (2021), pp. 2567–2581. DOI: 10.1002/jum.15682. URL: https://doi.org (cit. on p. 70).

[32] Ferenc Igor Kandi, Catarina Dinis Fernandes, Simona Turco, Eva de Bock, Lynda Juffermans, Judith Huirne, and Massimo Mischi. «2D Contrast-Ultrasound Dispersion Imaging of Angiogenesis in Adenomyosis: First Experimental Measurements». In: *2024 IEEE International Symposium on Medical Measurements and Applications (MeMeA)*. 2024, pp. 1–6. DOI: 10.1109/MeMeA60663.2024.10596726 (cit. on p. 70).