

Motion-Compensated μ MRI of Tethered Insect Locomotion

Zur Erlangung des akademischen Grades eines
DOKTORS DER INGENIEURWISSENSCHAFTEN (Dr.-Ing.)

von der KIT-Fakultät für Maschinenbau des
Karlsruher Instituts für Technologie (KIT)
angenommene

DISSERTATION

von

M.Tech. Ajmal Chenakkara Jamal

Tag der mündlichen Prüfung:	12.08.2026
Erstgutachter:	Prof. Dr. Jan G. Korvink
Zweitgutachter:	Prof. Dr. Maxim Zaitsev
Weiterer Betreuerin:	Prof. Dr. Ulrike Wallrabe
Betreuender Mitarbeiter:	Dr. Mazin Jouda



This document is licensed under a Creative Commons
Attribution-ShareAlike 4.0 International License (CC BY-SA 4.0):
<https://creativecommons.org/licenses/by-sa/4.0/deed.en>

Karlsruher Institut für Technologie
Institut für Mikrostrukturtechnik
Hermann-von-Helmholtz Platz 1
76344 Eggenstein-Leopoldshafen

Selbstständigkeitserklärung

Ich versichere hiermit wahrheitsgemäß, die Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt, die wörtlich oder inhaltlich übernommenen Stellen als solche kenntlich gemacht und die Satzung des Karlsruher Instituts für Technologie (KIT) zur Sicherung guter wissenschaftlicher Praxis in der jeweils gültigen Fassung beachtet zu haben.

Eggenstein-Leopoldshafen, 12.08.2026

Ajmal Chenakkara Jamal

Abstract

Magnetic resonance imaging (MRI) provides a powerful, non-invasive means to probe biological structure and function across scales, offering unparalleled versatility and making it a promising modality for studying small, freely behaving model organisms. However, living systems are inherently dynamic, and motion presents a persistent challenge for MRI, mainly due to its relatively slow, sequential data acquisition. In clinical and preclinical imaging, well-established strategies exist to address the effects of physiological motion and small involuntary movements on image quality. These approaches, however, do not address large-scale behavioral motion. While such motion is rarely considered in clinical imaging, studies of model organisms often aim to observe biological processes *in vivo*, where natural movement is an integral part of the phenomena under investigation.

Micro-magnetic resonance imaging (μ MRI), the high-resolution adaptation of MRI for microscopic scales, offers a unique window into the internal morphology of small model organisms and represents an ideal modality for such investigations. Realizing this potential requires new imaging frameworks that treat motion as an intrinsic part of the MRI measurement, allowing natural behavior while modeling and correcting its effects during acquisition and reconstruction. Insects serve as excellent models for developing and validating these approaches. Towards this goal, this thesis explores the feasibility of MRI in small moving organisms and examines the limitations of implementing such a framework within a μ MRI system. Initial findings motivate the development of an experimental platform designed to study live, active insects under minimally constrained conditions using an *in situ* treadmill, combined with methods for motion-compensated imaging during tethered locomotion.

The MR-compatible treadmill system enables tethered insect locomotion within a 15.2 T ultra-high-field preclinical scanner. This configuration constrains motion largely within the coil field of view (FOV) while allowing near-natural walking behavior. In addition, an integrated MR-compatible computer vision system was developed, employing an optical flow algorithm for real-time behavioral motion tracking (mean tracking error of 2.04 ± 1.13 pixels), which provided a robust means to address motion-related challenges during imaging. A vision-based prospective gating approach was first introduced to compensate for the semi-periodic abdominal pitching during locomotion. This enabled snapshot imaging of a 2D sagittal section of the insect at $117\text{ }\mu\text{m}$ isotropic resolution, as well as an axial section through

the thorax at 39 μm resolution, clearly delineating internal morphology during behavior. While this enabled motion-compensated imaging, temporal resolution remained limited and performance depended on behavioral variability of the insect. Retrospective gating improved flexibility and enabled the possibility of dynamic imaging, but residual artifacts persisted due to the inherently non-rigid nature of insect body motion and measurement uncertainties.

To address the limitations of gating-based motion compensation, a deep learning motion correction step was applied after gating as part of a combined approach, leading to greater robustness against motion artifacts and improved image quality. Achieving true dynamic imaging, however, requires faster acquisition. To this end, acceleration was achieved via undersampled Cartesian acquisition, with a deep learning-based framework employed to mitigate the resulting reconstruction artifacts. In the absence of parallel imaging hardware, which limited spatial acceleration, integrating a view-sharing strategy with deep learning proved more effective, leveraging temporal redundancy to improve initialization and restore image quality in dynamic 2D imaging. While the resulting temporal resolution of 1.2 s/frame at an acceleration factor of $2.5\times$ remain modest, the proposed approach represents a promising direction for future non-invasive dynamic biomechanical studies of tethered insects.

These contributions provide a systematic framework for motion-inclusive μMRI , bridging the gap between high-resolution structural imaging and the study of small organisms in active, near-naturalistic conditions.

Zusammenfassung

Die Magnetresonanztomographie (MRT) bietet ein leistungsstarkes, nicht-invasives Verfahren zur Untersuchung biologischer Strukturen und Funktionen über verschiedene Skalen hinweg, bietet eine beispiellose Vielseitigkeit und macht sie zu einer vielversprechenden Methode zur Untersuchung kleiner, frei beweglicher Modellorganismen. Lebende Systeme sind jedoch von Natur aus dynamisch, und Bewegungen stellen für die MRT eine ständige Herausforderung dar, vor allem aufgrund der relativ langsamen, sequenziellen Datenerfassung. In der klinischen und präklinischen Bildgebung gibt es etablierte Strategien, um die Auswirkungen physiologischer Bewegungen und kleiner unwillkürlicher Bewegungen auf die Bildqualität zu bewältigen. Diese Ansätze berücksichtigen jedoch keine großräumigen Verhaltensbewegungen. Während eine solche Bewegung in der klinischen Bildgebung selten berücksichtigt wird, zielen Studien an Modellorganismen häufig darauf ab, biologische Prozesse *in vivo* zu beobachten, bei denen natürliche Bewegungen einen integralen Bestandteil der zu untersuchenden Phänomene bilden.

Mikro-magnetresonanztomographie (μ MRT), die hochauflösende Adaption von MRT für mikroskopische Maßstäbe, bietet einen einzigartigen Einblick in die innere Morphologie kleiner Modellorganismen und stellt eine ideale Methode für solche Untersuchungen dar. Um dieses Potenzial auszuschöpfen, sind neue Bildgebungskonzepte erforderlich, die Bewegung als wesentlichen Bestandteil der MRT-Messung betrachten, natürliches Verhalten zulassen und gleichzeitig dessen Auswirkungen während der Datenerfassung und Rekonstruktion modellieren und korrigieren. Insekten eignen sich hervorragend als Modelle für die Entwicklung und Validierung dieser Ansätze. Mit diesem Ziel vor Augen untersucht diese Arbeit die Machbarkeit der MRT bei kleinen, frei bewegenden Organismen und beleuchtet die Grenzen der Implementierung eines solchen Rahmens innerhalb eines μ MRT-Systems. Diese initialen Befunde motivieren die Entwicklung einer experimentellen Plattform, die darauf ausgelegt ist, lebende, aktive Insekten unter minimal eingeschränkten Bedingungen mithilfe eines *in situ*-Laufbands zu untersuchen, kombiniert mit Methoden zur bewegungskompensierten Bildgebung während der *tethered Fortbewegung*.

Das MR-kompatible Laufbandssystem ermöglicht die *tethered Fortbewegung* von Insekten innerhalb eines 15.2 T Ultrahochfeld-Präkliniskanners. Diese Konfiguration begrenzt die Bewegung weitgehend innerhalb des FOV der Spule, während gleichzeitig ein nahezu natür-

liches Laufverhalten. Darüber hinaus wurde ein integriertes MR-kompatibles Computer-Vision-System entwickelt, das einen Optical-Flow-Algorithmus zur Echtzeit-Erfassung von Verhaltensbewegungen einsetzt (mittlerer Tracking-Fehler von 2.04 ± 1.13 Pixeln) und damit eine robuste Methode zur Bewältigung bewegungsbedingter Herausforderungen während der Bildgebung bereitstellt. Ein visionsbasiertes prospektives Gating-Verfahren wurde initial eingeführt, um die semi-periodische Pitching-Bewegung des Abdomens während der Lokomotion zu kompensieren. Dadurch wurde eine Snapshot-Bildgebung eines 2D-sagittalen Schnitts des Insekts mit isotroper Auflösung von $117 \mu\text{m}$ sowie eines axialen Schnitts durch den Thorax mit $39 \mu\text{m}$ Auflösung ermöglicht, wodurch die innere Morphologie während des Verhaltens klar abgebildet werden konnte. Obwohl dadurch eine bewegungskompensierte Bildgebung erreicht wurde, blieb die zeitliche Auflösung begrenzt und die Leistung hing von der Verhaltensvariabilität des Insekts ab. Retrospektives Gating erhöhte die Flexibilität und eröffnete die Perspektive dynamischer Bildgebung, jedoch traten weiterhin Residualartefakte auf, bedingt durch die inhärent nicht-rigide Natur der Insektenkörperbewegung sowie Messunsicherheiten.

Um die Einschränkungen der gating-basierten Bewegungskompensation zu überwinden, wurde nach dem Gating ein Deep-Learning-basierter Korrekturschritt als Teil eines kombinierten Ansatzes angewandt, was zu einer höheren Robustheit gegenüber Bewegungsartefakten und einer verbesserten Bildqualität führte. Die Realisierung einer echten dynamischen Bildgebung erfordert jedoch eine schnellere Akquisition. Zu diesem Zweck wurde eine Beschleunigung durch unterabgetastete kartesische Akquisition erreicht, wobei ein Deep-Learning-gestütztes Framework zur Reduktion der daraus resultierenden Rekonstruktionsartefakte eingesetzt wurde. In Abwesenheit von Parallelbildgebungshardware, die die räumliche Beschleunigung einschränkte, erwies sich die Integration einer View-Sharing-Strategie mit Deep Learning als besonders effektiv, indem zeitliche Redundanz genutzt wurde, um die Initialisierung zu verbessern und die Bildqualität in der dynamischen 2D-Bildgebung wiederherzustellen. Obwohl die resultierende zeitliche Auflösung von 1.2 s/Bild bei einem Beschleunigungsfaktor von $2.5\times$ weiterhin moderat ist, stellt der vorgeschlagene Ansatz eine vielversprechende Richtung für zukünftige nicht-invasive dynamische biomechanische Untersuchungen an Insekten dar.

Diese Beiträge liefern einen systematischen Rahmen für bewegungsintegrierte μMRT und schließen die Lücke zwischen hochauflösender struktureller Bildgebung und der Untersuchung kleiner Organismen unter aktiven, nahezu naturgetreuen Bedingungen.

List of abbreviations

Abbreviation	Definition
MRI	Magnetic Resonance Imaging
MRT	Magnetresonanztomographie
μ MRT	Mikro-Magnetresonanztomographie
fMRI	Functional Magnetic Resonance Imaging
MRS	Magnetic Resonance Spectroscopy
CSI	Chemical Shift Imaging
MRSI	Magnetic Resonance Spectroscopic Imaging
PRESS	Point Resolved Spectroscopy
μ MRI	Micro-Magnetic Resonance Imaging
MRM	Magnetic Resonance Microscopy
NMR	Nuclear Magnetic Resonance
RF	Radiofrequency
SNR	Signal-to-Noise Ratio
CNR	Contrast-to-Noise Ratio
FOV	Field of View
FID	Free Induction Decay
FFT	Fast Fourier Transform
DFT	Discrete Fast Fourier Transform
PMC	Prospective Motion Correction
RMC	Retrospective Motion Correction
PNS	Peripheral Nerve Stimulation
VS	View Sharing
TR	Repetition time
TE	Echo time

Abbreviation	Definition
TE _{eff}	Effective echo time
GRE	Gradient Recalled Echo
SE	Spin Echo
EPI	Echo Planar Imaging
RARE	Rapid Acquisition with Relaxation Enhancement
TSE	Turbo Spin Echo
FSE	Fast Spin Echo
ETL	Echo Train Length
FLASH	Fast Low Angle Shot
CS	Compressed Sensing
ZF	Zero-Filling
PCB	Printed Circuit Board
TTL	Transistor–Transistor Logic
EMI	Electromagnetic Interference
PLA	Polylactic Acid
MLP	Multilayer Perceptron
CNN	Convolutional Neural Network
RNN	Recurrent Neural Network
AI	Artificial Intelligence
DL	Deep Learning
ReLU	Rectified Linear Unit
MSE	Mean Squared Error
RMSE	Root Mean Squared Error
MAE	Mean Absolute Error
SSIM	Structural Similarity Index Measure
PSNR	Peak Signal-to-Noise Ratio
GE	Gradient Entropy
YOLOv8	You Only Look Once
YOLOv8n	You Only Look Once Nano
COCO8	Common Objects in Context
OFCE	Optical Flow Constraint Equation

List of symbols and constants

Symbols/Constants	Definition
B_0	Static main magnetic field (T)
B_1	Transverse radio-frequency magnetic field
γ	Gyromagnetic ratio (42.577 MHz/T for ^1H)
M_0	Equilibrium magnetization
M_{xy}	Transverse magnetization
M_z	Longitudinal magnetization
T_1	Longitudinal (spin-lattice) relaxation time
T_2	Transverse (spin-spin) relaxation time
T_2^*	Effective transverse relaxation time
T_2'	Transverse relaxation time due to reversible field inhomogeneities
ΔB_0	Static magnetic field inhomogeneity
α	Flip angle or tip angle
k_x	Spatial frequency along the x -direction
k_y	Spatial frequency along the y -direction
ρ_0	Spin density
χ	Magnetic susceptibility
G_x	x -direction magnetic field gradient
G_y	y -direction magnetic field gradient
G_z	z -direction magnetic field gradient
k_B	Boltzmann constant ($1.381 \times 10^{-23} \text{ J K}^{-1}$)
$\hbar$	Reduced Planck constant ($1.054 \times 10^{-34} \text{ J s}$)

Contents

Abstract	i
Zusammenfassung	iii
List of abbreviations	v
1 Introduction	1
1.1 Motivation	1
1.2 Outline	4
1.3 Main results	6
2 Background	7
2.1 Nuclear magnetic resonance (NMR)	7
2.1.1 Nuclear spin and magnetic moment	8
2.1.2 Net magnetization	8
2.1.3 Radiofrequency (RF) excitation	9
2.1.4 Relaxation	10
2.1.5 Signal formation: FID vs. echoes	12
2.1.6 Chemical shift	14
2.2 Magnetic resonance imaging (MRI)	15
2.2.1 Spatial encoding	15
2.2.2 k -space formalism	17
2.2.3 Pulse sequences	19
2.2.4 Image contrast	21
2.3 Motion in MRI	23
2.3.1 Effects of subject motion	24
2.3.2 Motion estimation and correction	25
2.3.3 Motion tracking methods	27
2.3.4 Optical flow for motion tracking	28
2.4 Deep learning in MRI	30
2.4.1 Multi-layer perceptrons and convolutional networks	30
2.4.2 Transfer learning	34

3	Freely moving organism	37
3.1	The scale-of-motion problem	37
3.2	<i>Ex situ</i> platform: freely moving insect	39
3.2.1	<i>Ex situ</i> platform and ladybug beetle	40
3.2.2	Deep learning-based motion tracking	41
3.2.3	<i>Ex situ</i> tracking characterization	42
3.2.4	Kalman prediction	43
3.2.5	Limitations	47
3.3	Constrained locomotion: treadmill	49
3.3.1	The Sun beetle: <i>Pachnoda marginata</i>	50
3.3.2	Insect μ MRI challenges and considerations	53
3.4	Conclusions	55
4	<i>In situ</i> treadmill	57
4.1	Introduction	57
4.2	<i>In situ</i> treadmill: design and development	59
4.2.1	Treadmill design iterations	59
4.3	<i>In situ</i> MR-compatible optical imaging system	61
4.3.1	RF interference on the <i>in situ</i> optical imaging system	63
4.3.2	Influence of the <i>in situ</i> optical imaging system on MRI	63
4.4	Animal handling	65
4.5	Baseline vs. motion-affected MRI	66
4.6	Sun beetle tethered motion model	68
4.6.1	Multi-rigid approximation	68
4.7	Prospective gating	71
4.7.1	Computer vision-based prospective gating	72
4.7.2	Computer vision system performance characterization	74
4.8	MRI of the behaving insect	75
4.8.1	Axial MRI of the thorax	75
4.8.2	Sagittal MRI of the insect body	76
4.9	Discussion	79
4.10	Conclusion	80
5	Residual motion artifact correction	81
5.1	Retrospective gating	84
5.1.1	Retrospective gating for the initial correction	86
5.2	MRI of the dormant insect: internal organ motion	88
5.3	Deep learning for residual motion correction	90
5.3.1	Datasets	91
5.3.2	Motion artifact simulation	92
5.3.3	Deep learning model architecture and problem formulation	96

5.3.4	Model training and evaluation	97
5.3.5	Image quality metrics	98
5.3.6	Motion correction on simulated data	100
5.3.7	Motion correction on experimental data	100
5.4	Discussion and conclusion	102
6	Time resolved μMRI	105
6.1	Deep learning for MRI acceleration	106
6.1.1	MRI undersampling simulation	107
6.1.2	Deep learning model training	108
6.1.3	Model evaluation on simulated data	109
6.2	Retrospective view-sharing for dynamic MRI	111
6.2.1	k -space undersampling and view-sharing: implementation	112
6.2.2	Dynamic MRI: view-sharing with deep learning correction	113
6.2.3	Correlated dynamic μ MRI and optical imaging	115
6.3	Discussion and conclusion	116
7	<i>In vivo</i> MRS: towards functional characterization	119
7.1	<i>In vivo</i> localized ^1H MRS	119
7.1.1	Spectroscopic sequence and parameters	120
8	Conclusions and outlook	123
8.1	Conclusions	123
8.2	Outlook	125
	Bibliography	127
	Publications	141
	Appendix	143
A	Data availability	143
B	Imaging parameters	143
C	DL motion correction: simulated dataset distribution	144
	List of Figures	151
	List of Tables	153
	Acknowledgments	155

1 | Introduction

1.1 Motivation

Biological systems are inherently dynamic, and many processes of interest occur under conditions of natural, unrestricted motion. Imaging and spectroscopic techniques used to study such systems must therefore minimize perturbation and avoid invasive interventions that could alter physiological or behavioral states. Consequently, the ability to investigate freely behaving organisms is central to many areas of zoology, enabling the study of locomotion, biomechanics, neural activity, and other physiological processes within an ecologically relevant context.

A wide range of experimental modalities and tools are available to address these goals, including X-ray, optical, and electrophysiological techniques. These approaches are based on distinct physical principles and provide different types of information, varying in spatial and temporal resolution as well as in their degree of invasiveness. Conventional experimental approaches often require sacrificing, physically restraining, or anesthetizing the organism, conditions that can significantly alter natural behavior. Despite these limitations, methods such as fluorescence imaging, optogenetics, and electrophysiology have developed into powerful tools for studying neural and physiological activity in behaving vertebrate and invertebrate model organisms, often under partially restrained conditions [1, 2, 3]. Similarly, high-speed X-ray cine-tomography has enabled four-dimensional imaging of internal dynamics in small living insects [4, 5]. Nevertheless, many of these approaches still impose physical or physiological constraints and remain invasive to varying degrees.

Magnetic resonance imaging (MRI) stands out among imaging techniques for its versatility and its ability to provide rich morphological, functional, chemical, and dynamic information about living matter non-invasively. It therefore represents a promising tool for studying freely moving biological systems. Although MRI techniques have been widely used to image dynamic aspects of internal physiology in both clinical studies and in preclinical small-animal research, their application has largely been restricted to otherwise static subjects. In particular, “behavioral motion” or large-scale body motion is typically avoided, and is generally treated as a limitation that must be minimized rather than an opportunity to be explored.

The fundamental limitation arises from the nature of MRI data acquisition itself. In general, raw MRI signals (k -space) are sampled sequentially and relatively slowly, with the imaging coordinate system defined by three orthogonal magnetic field gradients. Motion of the subject during acquisition disrupts the alignment between the imaging and subject coordinate systems, leading to artifacts in the reconstructed images and representing a long-standing challenge in MRI. Considerable efforts have been made to address various aspects of motion in clinical MRI, with a wide range of techniques spanning the entire MRI pipeline, from acquisition to image reconstruction [6, 7].

These constraints become even more pronounced in ultra-high-field systems, where longer acquisition times and increased sensitivity to field inhomogeneities amplify these artifacts. Although small animal models in preclinical MRI have advanced toward dynamic functional imaging of brain activity and internal physiology, even recent awake fMRI studies [8, 9, 10] typically require substantial physical restraint, and often brief pre-scan anesthesia, to mitigate motion artifacts. These challenges intensify in magnetic resonance microscopy (MRM), also known as micro-magnetic resonance imaging (μ MRI), a high-resolution variant of MRI optimized for small samples and typically achieving voxel resolutions below $100\text{ }\mu\text{m}^1$. μ MRI offers exceptional morphological detail in small organisms but is limited by high motion sensitivity and the spatial constraints of narrow-bore geometries, for the case of freely moving organisms. Consequently, μ MRI applications in invertebrate zoology have largely focused on static morphological imaging of immobilized, sedated, or postmortem specimens [11, 12, 13].

Flow-through MRI systems have demonstrated that high-resolution, non-invasive imaging of physiological processes is possible in live fish and marine invertebrates under controlled, simulated environmental conditions [14, 15, 16]. While these approaches relax some experimental constraints, they remain incompatible with unconstrained locomotion or spontaneous, self-directed behavior, including stimulus-driven responses under naturalistic conditions. Even in adult human MRI, the nominally optimal acquisition state depends on a combination of conscious cooperation, instructed immobility, physical stabilization, and often anxiety-driven self-restraint to minimize motion artifacts.

Collectively, these challenges highlight a fundamental limitation: the study of behavior, locomotion, and biomechanical processes using MRI remains largely restricted by current methodological constraints. Developing methods that allow high-resolution imaging of freely behaving subjects would represent a major advance, providing structural and functional insights under naturalistic conditions. To realize this potential, motion and dynamics must be treated as intrinsic components of the measurement rather than merely sources of artifacts.

¹ The terms μ MRI and MRI are used interchangeably in this thesis, depending on whether the context emphasizes the high-resolution experimental setup or the general MRI modality.

Fundamentally, this challenge consists of two main components: developing experimental frameworks that permit natural behavior within the imaging system, and designing strategies to mitigate motion artifacts that arise. Fully unconstrained movement is generally incompatible with MRI, while excessive restraint suppresses the behaviors of interest. These considerations motivate minimally constrained paradigms that stabilize the subject while preserving essential locomotor activity, using insects as the model system. Initial efforts focused on capturing single-frame snapshots during active locomotion states of the insect. This evolved toward dynamic imaging, enabling continuous observation of locomotor activity and the potential to quantify biomechanical processes.

In summary, the thesis investigates how μ MRI can be adapted to accommodate the behavioral motion of small organisms through novel experimental and methodological frameworks. The following sections will present the structure and main contributions of this thesis.

1.2 Outline

The theoretical background underlying the core principles and technologies addressed in this thesis is presented in Chapter 2. Initial exploratory investigations make Chapter 3. The principal scientific contributions of this work are presented in three manuscripts, which form the core Chapters 4, 5, and 6. Finally, potential directions for future research are discussed in Chapters 7 and 8. An overview of these is presented below:

Chapter 2 reviews the fundamentals of nuclear magnetic resonance (NMR) signal formation and describes how spatial information is encoded to generate images in MRI. It then proceeds with the effect of motion in MRI: motion during acquisition violates the assumption of a static subject and results in motion artifacts. Finally, the principles of deep learning are introduced within the scope of this work, highlighting their relevance to the applications developed throughout the thesis.

Chapter 3 introduces the paradigm of freely moving organisms in the context of MRI, extending the discussion of general motion-related challenges to the specific case of freely moving insects in μ MRI. The chapter highlights the associated technical and methodological challenges, and presents initial exploratory work that lays the foundation for the core contributions of the thesis. In particular, it introduces an *in situ* treadmill-based paradigm for imaging active, behaving organisms under minimal constraints.

Chapter 4 further develops the idea of an *in situ* treadmill as an effective platform for MRI of actively behaving insects. This chapter focuses on the design and development of the treadmill system, including the development and integration of an MR-compatible optical imaging setup and an optical flow-based vision algorithm for real-time motion tracking of active insects. The chapter concludes with the development of a prospective gating strategy for motion-compensated MRI of the sun beetle (*Pachnoda marginata*).

Chapter 5 presents a retrospective motion correction strategy for behaving insects. Prospective gating alone often yields poor temporal resolution due to the irregular locomotion patterns of tethered insects. To address this, a fully retrospective motion compensation approach was developed, combining retrospective gating with a deep learning-based residual motion artifact correction framework. This method enhances temporal resolution and advances the goal of dynamic MRI in active insects.

Chapter 6 focuses on the development of an accelerated MRI acquisition scheme for dynamic imaging. While Cartesian undersampling combined with supervised deep learning accelerates acquisition, excessive undersampling can lead to severe blurring. To mitigate this, a view-sharing preprocessing step followed by deep learning-based correction significantly improves image reconstruction quality, providing higher apparent temporal resolution to resolve the biomechanics of insect locomotion.

Chapter 7 explores the feasibility of *in vivo* localized magnetic resonance spectroscopy (MRS) in tethered insects by demonstrating stable acquisition of localized spectra, including water and lipid resonances, over extended durations.

Chapter 8 provides concluding remarks for the thesis and outlines potential directions for future research.

1.3 Main results

The research conducted within the scope of the doctoral project has led to three first-author publications, which are summarized below. A complete list of all publications, including co-authored works and conference contributions, is provided in Chapter 8.2.

1. **Motion compensated magnetic resonance imaging of an active sun beetle using an *in situ* treadmill [17]**

Journal: Scientific Reports (2025)

Authorship: First author

Content: This work introduces the *in situ* spherical treadmill platform for performing motion-compensated MRI of a tethered, walking sun beetle. The study focuses on the design and implementation of the treadmill system, and an MR-compatible computer vision pipeline for real-time motion estimation for a prospective gating scheme for motion compensated MRI.

Contribution: Conceptualization, methodology, software, investigation, data curation, visualization, writing.

2. **Residual motion artifact removal enables dynamic μ MRI of a behaving *Pachnoda marginata* [18]**

Journal: Journal of Magnetic Resonance (2025)

Authorship: First author

Content: This study focuses on the implementation of retrospective gating for motion-compensated μ MRI of a behaving sun beetle. Compared with prospective gating, this approach offers more relaxed constraints and increased flexibility. Residual motion artifacts introduced by the gating process are subsequently corrected using a deep learning-based framework.

Contribution: Conceptualization, methodology, software, investigation, data curation, visualization, writing.

3. ***In situ* time-resolved motion of a tethered *Pachnoda marginata*, AI-correlated using μ MRI and optical imaging [19]**

Journal: Journal of Magnetic Resonance (2026)

Authorship: First author

Content: Here, the limitations of gating in the presence of non-rigid, irregular motion of the sun beetle are addressed using an accelerated acquisition scheme. View-sharing of k -space improves undersampling efficiency, and a subsequent deep learning module refines the reconstruction. This ultimately paves the way for dynamic μ MRI of a behaving organism.

Contribution: Conceptualization, methodology, software, investigation, data curation, visualization, writing.

2 | Background

This chapter reviews the foundational concepts underlying magnetic resonance imaging (MRI), beginning with the physical principles that govern signal generation. It then describes how these signals are spatially encoded to produce images, and examines the impact of subject motion on image reconstruction. The chapter concludes with an overview of the theoretical foundations and methodologies of deep learning relevant to the applications developed in this thesis.

Parts of this chapter are based on the principles outlined in the textbook by Haacke et al. [20], with more detailed treatments available in the comprehensive references by Bernstein et al. [21], Levitt [22], and Callaghan [23].

2.1 Nuclear magnetic resonance (NMR)

The versatility of MRI arises from a fundamentally unique mechanism of image formation compared to conventional scattering or attenuation-based imaging modalities [23]. Instead of relying on the interaction of radiation with matter through absorption or scattering, MRI derives image contrast from the intrinsic magnetic properties of specific atomic nuclei (primarily ^1H in clinical imaging) possessing nonzero **nuclear spin**. The physical basis of MRI is the phenomenon of nuclear magnetic resonance (NMR), which exploits the intrinsic quantum mechanical property of nuclear spin and its associated magnetic moment. When placed in an external magnetic field and subjected to timed radiofrequency (RF) excitation, nuclear spins exhibit a characteristic relaxation response that depends on their local chemical and physical environment. This gives rise to the measurable MR signal and forms the basis for the extraction of spatially, temporally, and chemically specific information from the sample.

In this section, the fundamental principles of NMR are reviewed, beginning with the behavior of individual nuclear spins in a magnetic field and progressing to the formation of a measurable macroscopic magnetization from the ensemble of spins. This collective magnetization constitutes the primary source of the MR signal and provides the foundation for subsequent spatial encoding and image formation.

2.1.1 Nuclear spin and magnetic moment

Atomic nuclei with an odd number of protons and/or neutrons possess a quantum mechanical property known as nuclear spin angular momentum, denoted by I , with an associated nuclear magnetic dipole moment, μ , defined by the relationship:

$$\mu = \gamma I, \quad (2.1)$$

where γ is the gyromagnetic ratio, a nucleus-specific constant. For the ^1H nucleus, $\gamma/2\pi \approx 42.58 \text{ MHz/T}$.

When placed in a static external magnetic field B_0 , the magnetic moment experiences a torque $\tau = \mu \times B_0$. This interaction causes the magnetic moment to precess about the axis of the external field at a specific frequency known as the **Larmor frequency**, ω_0 :

$$\omega_0 = \gamma B_0. \quad (2.2)$$

This precession is an inherent property of the spin system and persists even at thermal equilibrium. The Larmor frequency establishes the fundamental resonance condition required for subsequent excitation and signal detection.

2.1.2 Net magnetization

At thermal equilibrium, nuclear spins in a sample preferentially align along the static magnetic field B_0 , giving rise to a net magnetization M_0 along the longitudinal (z) axis. For spin- $\frac{1}{2}$ nuclei such as protons (^1H), the spins occupy two possible energy states: the lower-energy α state (magnetic moment parallel to B_0) and the higher-energy β state (magnetic moment antiparallel to B_0). This is illustrated in Fig. 2.1(a–b).

The slight population difference between these two states is governed by the Boltzmann distribution. Under the high-temperature approximation (where the Zeeman energy is much smaller than the thermal energy, $\gamma\hbar B_0 \ll k_B T$), the fractional population difference is expressed as:

$$\frac{N_\alpha - N_\beta}{N} \approx \frac{\gamma\hbar B_0}{2k_B T} \ll 1, \quad (2.3)$$

where N_α and N_β are the number of spins in the α and β states, respectively, N is the total number of spins, γ is the gyromagnetic ratio, $\hbar$ is the reduced Planck constant, k_B is Boltzmann's constant, and T is the absolute temperature.

The resulting macroscopic magnetization is given by:

$$M_0 = \rho_0 \frac{\gamma^2 \hbar^2 B_0}{4k_B T} \mathbf{k}, \quad (2.4)$$

where ρ_0 is the equilibrium spin density and $\mathbf{k}$ is the unit vector along the z -axis. Although the population difference is extremely small (on the order of parts per million), it is sufficient to generate a measurable MR signal once the net magnetization is rotated into the transverse plane. The magnitude of M_0 scales linearly with the static field B_0 and inversely with temperature, providing the fundamental basis for MRI signal-to-noise ratio (SNR).




Figure 2.1: Nuclear spin magnetization and the energy diagram. **(a)** Nuclear spins at thermal equilibrium in the absence of an external magnetic field. **(b)** In the presence of a static magnetic field B_0 , a slight population imbalance results in a net longitudinal magnetization M_0 .

2.1.3 Radiofrequency (RF) excitation

At equilibrium, the net longitudinal magnetization M_0 is aligned with the static magnetic field B_0 . In this state, the magnetization is stationary and does not induce a detectable signal in the receiver coils. To generate a measurable MR signal, the magnetization must be tipped into the transverse (x - y) plane, as illustrated in Fig. 2.2(a-d).

The excitation process is commonly described using a semi-classical treatment. In this model, the collective quantum behavior of the spin ensemble is represented by a single macroscopic magnetization vector. Displacement of this vector is achieved by applying a RF magnetic field, $B_1(t)$, perpendicular to B_0 . To achieve resonance and effectively transfer energy to the spin system, the B_1 field must oscillate at the Larmor frequency ($\omega_0 = \gamma B_0$) in the laboratory frame. In a reference frame rotating at ω_0 , the RF field appears static, and the magnetization rotates about the B_1 axis. A pulse of appropriate amplitude and duration,

such as a $\pi/2$ pulse, rotates the longitudinal magnetization entirely into the transverse plane:

$$\alpha = \gamma B_1 \tau_{\text{rf}}, \quad (2.5)$$

where α is the flip angle and τ_{rf} is the RF pulse duration.

Following excitation, the transverse magnetization $\mathbf{M}_{xy}$ precesses at the Larmor frequency ω_0 , producing a time-varying magnetic flux through nearby receive coils. This changing flux induces voltage governed by **Faraday's law**,

$$\mathcal{E} = -\frac{d\Phi_B}{dt}, \quad (2.6)$$

yielding the **free induction decay (FID)**, the fundamental MR signal prior to spatial encoding. The resonance condition ($\omega_0 = \gamma B_0$) ensures efficient, coherent excitation.

2.1.4 Relaxation

After the magnetization is tipped into the transverse plane, it does not remain constant. The transverse magnetization decays and the longitudinal magnetization recovers due to interactions between spins and their environment, a process known as *relaxation*. This evolution is described by the phenomenological *Bloch equation*:

$$\frac{d\mathbf{M}}{dt} = \gamma(\mathbf{M} \times \mathbf{B}) - \frac{M_x \mathbf{i} + M_y \mathbf{j}}{T_2} - \frac{(M_z - M_0) \mathbf{k}}{T_1} \quad (2.7)$$

where the first term, $\gamma(\mathbf{M} \times \mathbf{B})$, represents the precessional motion of the magnetization in the magnetic field, while the subsequent terms account for the two primary relaxation processes:

- **Longitudinal relaxation (T_1):** Also known as *spin-lattice relaxation*, this describes the recovery of the longitudinal magnetization M_z toward its equilibrium value M_0 . This process involves an exchange of energy between the spin system and the surrounding molecular environment (the “lattice”). As illustrated in Fig. 2.2(f), the recovery follows:

$$M_z(t) = M_z(0)e^{-t/T_1} + M_0 \left(1 - e^{-t/T_1}\right). \quad (2.8)$$

- **Transverse relaxation (T_2):** Often called *spin-spin relaxation*, this describes the decay of transverse magnetization M_{xy} due to the loss of phase coherence between neighboring spins. Unlike T_1 , this process involves no net energy change but rather a dispersal of the phase. As illustrated in Fig. 2.2(g), the decay is given by:

$$M_{xy}(t) = M_{xy}(0) e^{-t/T_2}. \quad (2.9)$$




Figure 2.2: RF excitation and relaxation of net nuclear magnetization in the rotating frame of reference. **(a)** A 90° RF excitation rotates the net longitudinal magnetization M_0 into the transverse (x – y) plane. **(b–d)** Magnetization relaxing back to equilibrium, with **(e)** showing the FID signal. **(f)** Recovery of longitudinal magnetization, characterized by the longitudinal relaxation time T_1 . **(g)** Decay of transverse magnetization, characterized by the transverse relaxation time T_2 .

These relaxation processes are tissue-specific; consequently, they govern the signal intensity, decay characteristics, and the resulting contrast between different anatomical structures. In practice, the observed transverse signal decays more rapidly than T_2 predicts. This effective relaxation time, denoted as T_2^* , accounts for both the intrinsic spin–spin interactions and the additional dephasing caused by spatial inhomogeneities in the static magnetic field (ΔB_0). This relationship is expressed as:

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} = \frac{1}{T_2} + \gamma \Delta B_0, \quad (2.10)$$

where T_2' accounts for the dephasing caused by these field variations. Consequently, T_2^* is always shorter than T_2 , as it reflects the combined influence of irreversible spin–spin interactions and the additional, often reversible, dephasing from field non-uniformities.

2.1.5 Signal formation: FID vs. echoes

Immediately following a 90° RF excitation pulse that tips longitudinal magnetization into the transverse plane, the precessing transverse magnetization induces a voltage in the receive coil, producing the FID signal, as mentioned in Section 2.1.3. The FID reflects the combined effects of intrinsic transverse relaxation (T_2) and dephasing due to magnetic field inhomogeneities (ΔB_0), therefore decaying with the effective time constant T_2^* while oscillating at the Larmor frequency.

Spin echo (SE): By applying a 180° RF refocusing pulse, static field inhomogeneity-induced dephasing can be reversed through phase reversal, forming an *echo signal* whose decay is now governed primarily by T_2 , as illustrated in Fig. 2.3. Under ideal refocusing conditions, the echo signal amplitude at echo time TE is given by,

$$S_{\text{echo}} \propto M_0 e^{-TE/T_2}. \quad (2.11)$$

Gradient recalled echo (GRE): Alternatively, echo signals may be formed using deliberately applied magnetic field gradients superimposed on B_0 . These gradients induce controlled, spatially dependent dephasing across the sample, which can be reversed by inverting the polarity to generate a *gradient echo*, which is illustrated in Fig. 2.4. Unlike spin echoes, gradient echoes do not refocus dephasing caused by static field inhomogeneities and therefore remain T_2^* -weighted. Gradient-echo formation relies on controlled gradient reversal rather than RF refocusing, enabling rapid generation of echo signals (important for imaging), while remaining sensitive to dephasing from static magnetic field inhomogeneities. Together, FID and echoes constitute the fundamental MR signals that form the basis for subsequent spatial encoding using magnetic field gradients and image formation.




Figure 2.3: Spin echo refocusing.




Figure 2.4: Gradient echo refocusing.

2.1.6 Chemical shift

The measured FID reflects the coherent precession of an ensemble of spin isochromats at their respective Larmor frequencies. Within this framework, **chemical shift** arises as a consequence of the microscopic magnetic environment of each nucleus. Electrons surrounding the nucleus partially shield it from the external field B_0 , leading to a slightly modified effective magnetic field and, therefore, a small shift in resonance frequency. Nuclei in different molecular environments thus precess at slightly different frequencies. This frequency offset relative to a reference compound is expressed in parts per million (ppm) as,

$$\delta = \frac{\omega - \omega_{\text{ref}}}{\omega_{\text{ref}}} \times 10^6. \quad (2.12)$$

In NMR spectroscopy, these frequency differences form the fundamental basis of molecular and structural characterization, as chemically distinct species produce separate spectral peaks. In MRI, by contrast, chemical shift does not primarily encode molecular structure but instead appears as a frequency-dependent spatial displacement in the readout direction, often observed as chemical shift artifacts [20]. Nevertheless, this effect can be deliberately exploited in techniques such as fat–water separation [24].

Beyond conventional imaging, chemical shift is also utilized in magnetic resonance spectroscopic imaging (MRSI), also known as chemical shift imaging (CSI). Here, the spatial encoding and spectral information are combined to map the distribution of specific metabolites *in vivo*, effectively performing localized spectroscopy across an entire imaging volume.

2.2 Magnetic resonance imaging (MRI)

The signal measured in an NMR experiment can be in the form of an FID or as an echo, and constitutes the collective summation of transverse magnetization from all excited voxels within the detection range of the RF coil. The defining feature of MRI lies in its ability to spatially encode this signal. To discriminate contributions from different spatial locations, RF pulses and magnetic field gradients must be applied with precise timing and coordination. This orchestrated pattern of RF and gradient manipulations is known as a **pulse sequence** and forms the fundamental framework for translating nuclear spin information from physical space into a spatially encoded signal.

2.2.1 Spatial encoding

The fundamental mechanism of spatial encoding is rooted in the Larmor relationship, $\omega = \gamma B$, which dictates that the precessional frequency of nuclear spins is linearly proportional to the local magnetic field strength. By superimposing three orthogonal, linear magnetic field gradients; G_x , G_y , and G_z , onto the static field B_0 , the Larmor frequency becomes a function of position:

$$\omega(\mathbf{r}) = \gamma(B_0 + \mathbf{G} \cdot \mathbf{r}), \quad (2.13)$$

where $\mathbf{r}$ is the position vector. To resolve a multidimensional image, this spatial dependence is utilized in two distinct manners depending on the timing of the gradient application relative to signal acquisition. Spatial information is first encoded through the application of gradients for finite durations prior to signal acquisition, which induce position-dependent phase shifts in the transverse magnetization (phase encoding). During signal readout, a gradient is subsequently applied so that spatial location is directly mapped onto precessional frequency (frequency encoding). Conventionally, MRI signal formation and encoding are described in a reference frame rotating at the Larmor frequency, with the z -axis aligned along the direction of the main magnetic field B_0 .

2.2.1.1 Slice selection

Spatial localization begins with the selective excitation of spins. By applying a linear magnetic field gradient G_z along the z -direction, the total magnetic field becomes position-dependent:

$$B(z) = B_0 + G_z z, \quad (2.14)$$

resulting in a spatially varying Larmor frequency $\omega(z) = \gamma B(z)$. When an RF pulse with center frequency ω_{RF} and bandwidth $\Delta\omega$ is applied simultaneously with the gradient, only spins within a specific spatial region are tipped into the transverse plane (see Fig. 2.5(a)). The slice position z_0 and thickness Δz are defined as:

$$z_0 = \frac{\omega_{\text{RF}}/\gamma - B_0}{G_z}, \quad \Delta z = \frac{\Delta\omega}{\gamma G_z}. \quad (2.15)$$

2.2.1.2 Phase encoding

Once spins are excited within the slice, spatial encoding within the xy -plane is required. Phase encoding is performed by applying a gradient pulse G_y for a duration τ_{pe} along the y -axis before signal acquisition. This gradient induces a position-dependent phase shift $\phi(y)$ that persists after the gradient is deactivated:

$$\phi(y) = \gamma \int_0^{\tau_{pe}} G_y(\tau) y d\tau = \gamma G_y y \tau_{pe}. \quad (2.16)$$

This phase imprint manifests as a position-dependent factor $e^{i\phi(y)} = e^{i\gamma G_y \tau_{pe} y}$ in the received signal. To resolve the y -dimension, the pulse sequence is repeated N_y times, each with a unique gradient amplitude G_y , thereby sampling different spatial phase offsets across the object.




Figure 2.5: Spatial encoding in MRI. **(a)** Slice selection. **(b)** Phase and frequency encoding for k -space sampling, and subsequent image reconstruction. (Brain MRI from [25]).

2.2.1.3 Frequency encoding

The final spatial dimension is encoded during signal acquisition (readout). A constant gradient G_x is applied along the x -axis, causing the Larmor frequency to be a function of x position:

$$\omega(x) = \gamma(B_0 + G_x x). \quad (2.17)$$

This induces a time-dependent phase accrual $\phi_x(x, t) = \int_0^t \omega(x) dt' = \gamma G_x t x$ (ignoring the constant B_0 term). The received signal $s(t)$ is the integral of the transverse magnetization across the excited slice, encoded with both the persistent phase from the y -gradient and the evolving phase from the x -gradient:

$$s(t) = \int \int \rho(x, y) e^{i(\gamma G_y \tau_{pe}) y} e^{i(\gamma G_x t) x} dx dy, \quad (2.18)$$

where $\rho(x, y)$ represents the spin density in the imaging plane.

2.2.2 k -space formalism

The spatial encoding process described above can be unified by introducing the concept of k -space, which essentially represents the spatial frequency domain of the image. We can define the spatial frequencies k_x and k_y as the time-integrals of the applied gradients:

$$k_x(t) = \frac{\gamma}{2\pi} \int_0^t G_x(\tau) d\tau = \frac{\gamma}{2\pi} G_x t, \quad (2.19)$$

$$k_y = \frac{\gamma}{2\pi} \int_0^{\tau_{pe}} G_y(\tau) d\tau = \frac{\gamma}{2\pi} G_y \tau_{pe}. \quad (2.20)$$

By substituting these k -variables into Equation 2.18, the signal equation takes the form of a 2D Fourier Transform. This can be written compactly using the k -space vector $\mathbf{k} = (k_x, k_y)$ and the position vector $\mathbf{r} = (x, y)$:

$$s(\mathbf{k}) = \iint \rho(\mathbf{r}) e^{-i2\pi(\mathbf{k} \cdot \mathbf{r})} d\mathbf{r}. \quad (2.21)$$

In this framework, the sampling trajectory through k -space is determined by the pulse sequence timing: each readout acquires one line along k_x , while phase-encoding gradients step through k_y positions. A schematic illustration of phase and frequency encoding used for k -space sampling and the subsequent **Fourier reconstruction** is shown in Fig. 2.5(b).

2.2.2.1 Image reconstruction

The Fourier reconstruction implies that the spatial spin density $\rho(\mathbf{r})$ can be recovered from the measured k -space signal $s(\mathbf{k})$ by applying the inverse Fourier transform:

$$\rho(\mathbf{r}) = \iint s(\mathbf{k}) e^{i2\pi(\mathbf{k} \cdot \mathbf{r})} d\mathbf{k}. \quad (2.22)$$

In practice, MRI data are acquired on a discrete Cartesian grid in k -space, where the measured samples are denoted as $s[m, n] = s(m\Delta k_x, n\Delta k_y)$. Image reconstruction is implemented numerically using a two-dimensional inverse discrete Fourier transform (2D-IDFT):

$$\rho[p, q] = \frac{1}{N_x N_y} \sum_{m=0}^{N_x-1} \sum_{n=0}^{N_y-1} s[m, n] e^{i2\pi\left(\frac{mp}{N_x} + \frac{nq}{N_y}\right)}. \quad (2.23)$$

Efficient computation is achieved using fast Fourier transform (FFT) algorithms. The reconstructed image represents a discretized estimate of the underlying spin density, while the image contrast is governed by the selected pulse sequence and timing parameters.

2.2.2.2 Field of View and resolution

The FOV and spatial resolution (Δ) are inversely related to the sampling density and the maximum extent reached in k -space, respectively:

- **Phase Encoding (y-direction):**

$$\text{FOV}_y = \frac{1}{\Delta k_y}, \quad \Delta y = \frac{1}{2k_{y,\max}}$$

- **Frequency Encoding (x-direction):**

$$\text{FOV}_x = \frac{1}{\Delta k_x}, \quad \Delta x = \frac{1}{2k_{x,\max}}$$

Thus, the FOV is determined by the spacing between adjacent k -space samples (Δk), whereas the spatial resolution depends on the total k -space coverage ($2k_{\max}$). Adequate sampling (satisfying the Nyquist criterion) is essential to avoid aliasing artifacts and to achieve the desired detail in the image.

2.2.2.3 SNR and k -space properties

The SNR is a primary metric of image quality, defined as the ratio of the mean signal amplitude to the standard deviation of background noise. The signal amplitude scales with voxel volume $\Delta x \Delta y \Delta z$, while noise increases with the receiver bandwidth and decreases with the square root of the total acquisition time. Combining these effects for conventional fully sampled Cartesian imaging yields the simplified scaling

$$\text{SNR} \propto \Delta x \Delta y \Delta z \cdot \sqrt{\frac{N_y \cdot N_{\text{avg}}}{\text{BW}_{\text{read}}}}, \quad (2.24)$$

where N_y is the number of phase-encoding steps, N_{avg} the number of signal averages, and BW_{read} the receiver bandwidth. This relation assumes fixed sequence parameters and fully sampled acquisition; in practice, other factors such as field strength, coil sensitivity, and relaxation times also affect SNR.

The k -space signal distribution, and its impact on image quality, depends on both the underlying object structure and the acquisition parameters such as voxel size, readout bandwidth, and phase-encoding coverage. Key k -space properties relevant to SNR and resolution include:

- **Energy distribution:** Signal energy is non-uniformly distributed; the k -space center (low spatial frequencies) contains the highest amplitudes and determines global image contrast and bulk SNR. The periphery (high spatial frequencies) encodes sharp edges and fine structural details but possesses much lower signal amplitudes.
- **Resolution trade-offs:** Increasing spatial resolution by extending k -space coverage ($k_{\max}$) incurs a significant SNR penalty. Smaller voxels contain fewer spins and therefore generate less signal per voxel, resulting in a reduced SNR.

- **Readout bandwidth:** The receiver bandwidth (BW_{read}) determines the sampling rate along k_x . While a higher bandwidth reduces the spatial displacement of chemical shift artifacts, it increases the integrated noise, thereby reducing the overall SNR.

2.2.3 Pulse sequences

MRI pulse sequences consist of a carefully timed series of RF excitations, magnetic field gradient switching that imparts spatially dependent frequency and phase shifts to the precessing spin ensembles, and signal acquisition periods for the k -space sampling. By varying the order, timing, and combination of these elements, different strategies can be employed to manipulate spin behavior, encode spatial information, and generate tissue contrast. These sequences can be broadly classified as spin-echo or gradient-echo methods, depending on how the transverse magnetization is refocused and the resulting signal is formed, each offering distinct imaging characteristics and contrast mechanisms.

2.2.3.1 Spin Echo (SE) imaging

Spin echo imaging builds on the basic refocusing mechanism described in Section 2.1.5. In an SE sequence, a 90° RF excitation pulse tips the net magnetization into the transverse plane, where spins begin to dephase due to spin-spin interactions and static magnetic field inhomogeneities. After a time interval of $TE/2$, a 180° RF refocusing pulse is applied, which inverts the phase of the spins and reverses dephasing caused by static field variations. This process causes the spins to rephase, producing a coherent signal known as the spin echo at the echo time (TE). Because the refocusing pulse compensates for static field inhomogeneities, the resulting signal decay reflects the true transverse relaxation time T_2 , making spin echo imaging highly robust and reliable. By repeating the sequence at a chosen repetition time (TR) and systematically varying spatial encoding gradients, spin echo imaging enables precise control of image contrast and forms the gold standard for T_2 -weighted MRI. Basic schematic representation of the sequence is shown in Fig. 2.6.

2.2.3.2 Rapid Acquisition with Relaxation Enhancement (RARE) imaging

The RARE [26] sequence (also referred to as Turbo/Fast Spin Echo, TSE/FSE) is a rapid imaging technique derived from the conventional spin echo sequence. Following a single 90° RF excitation pulse, a train of 180° RF refocusing pulses is applied, generating multiple spin echoes within one TR. Unlike the standard spin echo sequence, which acquires only a single line of k -space per TR, each echo produced in the RARE sequence is associated with a different phase-encoding gradient. In this way, multiple lines of k -space are sampled during a single TR, forming an echo train. The number of echoes acquired per TR is defined as the echo train length (ETL), also known as the RARE factor.

By collecting several k -space lines per TR, the RARE sequence substantially reduces the total acquisition time. However, because the echoes are recorded at different echo times,




Figure 2.6: Spin echo pulse sequence.

the resulting reconstructed image exhibits a combination of T_2 weighting across k -space. To control the final image contrast, an effective echo time (TE_{eff}) is defined, corresponding to the echo which predominantly determines the image contrast. Schematic representation of the sequence is shown in Fig. 2.7.




Figure 2.7: RARE pulse sequence.

2.2.3.3 Gradient Echo (GRE) imaging

The GRE pulse sequence is a fast imaging technique that differs fundamentally from the conventional spin echo sequence by using gradient reversals rather than RF refocusing pulses to form the echo, as discussed in Section 2.1.5. Following a single RF excitation pulse with a flip angle α (typically $\alpha < 90^\circ$), dephasing and rephasing gradient lobes are applied to generate a gradient echo at the TE. Schematic representation of the sequence is shown in Fig. 2.8.

Unlike spin echo and RARE-based sequences, GRE does not employ 180° RF refocusing pulses. As a result, dephasing caused by magnetic field inhomogeneities is not refocused,

and the signal decay is governed by T_2^* rather than T_2 . Phase-encoding gradients are applied between repetitions, allowing one line of k -space to be acquired per TR. Because GRE sequences can be implemented with short α , TE and TR values, they enable rapid data acquisition and flexible contrast manipulation. Image contrast is determined by the choice of α , TR, and TE, and may emphasize T_1 , T_2^* , or proton density effects. Additionally, the absence of refocusing pulses reduces RF power deposition, making GRE sequences particularly suitable for fast imaging and three-dimensional acquisitions.




Figure 2.8: Gradient echo pulse sequence.

2.2.4 Image contrast

One of the key features that makes MRI a highly versatile modality is its ability to generate a wide range of soft-tissue contrast through flexible control of acquisition parameters. The signal measured from tissues is governed by their intrinsic magnetic properties, primarily the equilibrium magnetization M_0 (proportional to proton density (PD)) and the longitudinal and transverse relaxation time constants T_1 and T_2 . By appropriately selecting the TR and TE, the relative contributions of these parameters to the detected signal can be selectively emphasized, thereby determining the effective image weighting.

Under the simplified assumptions of a conventional 90° – 180° spin-echo sequence, perfect refocusing, steady-state magnetization, and negligible effects from flip-angle imperfections, magnetic field inhomogeneities, and diffusion, the signal S arising from a given tissue can be approximated as,

$$S \propto M_0 \left(1 - e^{-\frac{TR}{T_1}}\right) e^{-\frac{TE}{T_2}}. \quad (2.25)$$

From Equation 2.25, maximum signal amplitude is obtained using a long TR in combination with a short TE, a regime in which contrast is dominated primarily by proton

density. Increasing the echo time enhances T_2 -weighting by allowing greater transverse relaxation-dependent signal decay, whereas reducing the repetition time limits longitudinal recovery, thereby emphasizing T_1 -dependent contrast. A combination of short TR and long TE results in substantial signal attenuation and a corresponding reduction in SNR.

While this signal model captures the primary contrast mechanisms in spin-echo MRI, it does not account for additional effects such as T_2^* decay, magnetic susceptibility, diffusion, or perfusion. These effects can be selectively probed using alternative pulse sequences and advanced imaging techniques, which extend MRI contrast beyond the basic relaxation-based framework considered here, and are beyond the scope of this thesis (see [20, 21] for a comprehensive treatment).

As discussed in Section 2.2.3.2, RARE sequences acquire multiple echoes at progressively longer TE_{eff} within a single TR, producing a mixture of T_2 contrast across k -space. The central k -space lines, corresponding to low spatial frequencies, are sampled early, when transverse relaxation has caused minimal signal decay, while the peripheral lines, corresponding to high spatial frequencies, are acquired later and experience greater T_2 decay. This progressive decay effectively modulates k -space, acting as a low-pass filter or modulation transfer function (MTF) that attenuates high-frequency components. The resulting effect is a characteristic blurring of T_2 -weighted images, with occasional ghosting artifacts. This loss of image sharpness represents a fundamental trade-off for the substantial reduction in acquisition time offered by RARE.

2.3 Motion in MRI

The demonstration of nuclear magnetic resonance in bulk matter [27, 28] enabled its rapid evolution into a powerful spectroscopy technique for chemical analysis and was extended to study molecular motion, diffusion, and flow [29, 30], well before clinical MRI emerged as an imaging modality. In such experiments, motion is encoded through the phase accumulated by spins in the presence of magnetic field gradients, providing sensitivity to diffusion and coherent flow. On the other hand, conventional MRI as introduced later uses gradient-induced phase shifts primarily to encode spatial position, assuming a static spin density during the acquisition.

So far, the fundamentals of magnetic resonance and spatial encoding in MRI have been discussed under the assumption of a static subject. In reality, biological systems are rarely stationary, and can violate assumptions inherent in the underlying MRI forward model. This section examines motion as a fundamental challenge in MRI, with emphasis on its impact on image reconstruction. Concepts established in clinical MRI serve as the foundation and are discussed, where appropriate, in the context of the μ MRI domain (or ultra-high-field imaging in general), in which motion-related effects tend to be more pronounced due to the microscopic spatial resolution and prolonged acquisition times [31, 7]. The treatment of motion in MRI in this section is not exhaustive; comprehensive treatments can be found in relevant textbooks [20, 7] and extensive reviews [6, 32].

NMR experiments are inherently sensitivity-limited due to the low characteristic energy associated with magnetic spin transitions, necessitating prolonged data acquisition to achieve adequate SNR. In MRI, acquisition times are further extended by the requirements of spatial encoding and the use of complex imaging pulse sequences. As a result, the temporal resolution of MRI is fundamentally constrained by an interplay between the physics of nuclear relaxation and practical hardware limitations.

In general, MRI relies on the sequential sampling of k -space, with total acquisition time determined primarily by the number of phase-encoding steps and the TR. The choice of TR is itself constrained by the need for sufficient longitudinal (T_1) recovery to maintain adequate signal intensity and tissue contrast. In addition, hardware limitations, most notably the maximum achievable gradient strength and slew rate, which restrict the speed at which k -space can be traversed. Safety considerations, particularly those related to peripheral nerve stimulation (PNS), further limit gradient switching rates. Consequently, the temporal scale of MRI acquisitions is often substantially longer than that of common physiological motions encountered in clinical settings, such as respiration, cardiac motion, and involuntary patient movement.

Spatial encoding is defined by three orthogonal gradient coils that establish a fixed imaging coordinate system. Motion of the subject during k -space acquisition introduces inconsistencies between the object's true spatial coordinates and this encoding frame, violating the

assumptions underlying image reconstruction and giving rise to motion-related artifacts. These artifacts manifest as image degradations such as blurring of edges and contrast, ghosting from moving structures, spin dephasing with associated signal loss, and, in some cases, spurious signal contributions. Beyond the timing of motion relative to k -space acquisition, the type and severity of these artifacts also depend on the characteristics of the motion itself and on the imaging sequence and its parameters, including the spatial encoding strategy, repetition scheme, and image resolution.

These considerations apply equally to μ MRI, with additional challenges arising from the use of ultra-high static magnetic fields and strong gradient systems. While these features enable exceptional spatial resolution and sensitivity, they also exacerbate susceptibility to motion-induced phase errors, signal loss, and geometric distortions [7, 31, 6]. Consequently, even small or subtle motions can lead to pronounced image artifacts in μ MRI, necessitating careful experimental design and motion-robust acquisition strategies. Historically, however, motion was not considered a limiting factor in μ MRI (or to some extent in preclinical small-animal MRI) as it is in clinical imaging, as the experimental protocols typically relied on fully constrained, anesthetized, or *ex vivo* specimens [33]. This could still cause residual motion effects, which had to be actively managed [7]. In general, motion artifacts were largely mitigated through experimental design rather than explicit correction strategies, and comparatively limited attention was devoted to the development of motion-compensation approaches specifically for μ MRI.

2.3.1 Effects of subject motion

The primary criterion for motion characterization in MRI is its spatiotemporal relationship to k -space acquisition, which ultimately determines its impact on image reconstruction. Rigid-body motion, in which the subject moves as a whole, can be considered as the simplest case, which can be modeled as a combination of translations and rotations. While most biological motion is inherently non-rigid, the rigid-body approximation is often sufficient for many practical scenarios and offers significantly reduced modeling complexity. When higher fidelity is required, non-rigid motion can be approximated as a superposition of multiple locally rigid motions, enabling tractable and useful modeling.

Another commonly adopted assumption is in-plane motion, where movement is confined to the selected two-dimensional imaging slice. Throughout this thesis, two-dimensional Cartesian k -space sampling with standard Fourier-based image reconstruction is assumed unless stated otherwise. In this context, inter-view motion; motion occurring between successive phase-encoding steps, plays a particularly important role, as it directly influences the consistency of k -space data and the resulting image quality. Consequently, the severity and type of motion-induced artifacts depend not only on the spatial and temporal characteristics of the motion but also on its timing relative to k -space acquisition.

Rigid-body **translational motion** in the spatial domain manifests as a phase modulation in k -space, with the phase offset proportional to the magnitude of the translation. This relationship follows from the Fourier shift theorem and can be expressed as,

$$S(k_x, k_y) = S_0(k_x, k_y) e^{-i2\pi(k_x\Delta x + k_y\Delta y)} \quad (2.26)$$

where,

- $S(k_x, k_y)$ is the k -space signal after translation.
- $S_0(k_x, k_y)$ is the original k -space signal before translation.
- k_x, k_y are the spatial frequencies in the x - and y -directions.
- $\Delta x, \Delta y$ are the translation distances in the x - and y -directions.

Rigid-body **rotational motion** in the spatial domain results in a corresponding rotation of the k -space coordinate system, consistent with the Fourier rotation theorem. For an object undergoing a rotation by an angle θ about the origin, the transformed k -space coordinates $\mathbf{k}'$ are given by:

$$\begin{bmatrix} k'_x \\ k'_y \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} k_x \\ k_y \end{bmatrix} \quad (2.27)$$

and the corresponding k -space signal transform is given as:

$$S(\mathbf{k}') = S_0(\mathbf{k}) \quad (2.28)$$

2.3.2 Motion estimation and correction

Motion correction strategies in MRI can be broadly classified into prospective and retrospective approaches, based on when the correction is applied. **Retrospective motion correction (RMC)** methods perform the correction after data acquisition, either using motion information acquired externally (e.g., optical tracking) or extracted directly from the MR data itself [6]. Traditionally, many motion correction techniques have relied on iterative reconstruction algorithms that incorporate motion parameters into the forward model, optimizing a cost function based on data consistency or image sharpness. However, these conventional approaches often suffer from high computational costs and the need for accurate motion estimates. To address these limitations, deep learning-based approaches have emerged, which learn to estimate and correct motion-induced artifacts directly from the acquired data, often without explicit motion parameter estimation [34, 35]. In general, the acquisition proceeds independently of any subject motion, and the correction is applied during post-processing.

In **prospective motion correction (PMC)**, the imaging volume is updated in real time to follow the subject's motion, thereby keeping the FOV fixed in the subject's reference

frame rather than the scanner frame [36, 6]. This is accomplished by dynamically adjusting the gradient encoding and RF parameters based on real-time motion estimates. Accurate and frequent pose measurements of the subject are required to ensure effective correction. Numerous approaches have been proposed to implement PMC in clinical research and practice, which primarily differ in the methods used for motion estimation and tracking. Early implementations relied on MR navigator-based techniques, which acquire additional MR data at multiple time points and compare them to estimate motion without the need for external hardware. Navigators can operate in both k -space and image space [37, 38, 6], and generally increase acquisition and processing time.

External tracking approaches employ specialized hardware to perform motion tracking and are generally sequence-agnostic. Numerous tracking modalities have been used [39, 6], with RF-based and optical systems are the most common. One RF-based approach uses small pickup coils attached to the subject, with their positions encoded in the received signal frequency in the presence of a gradient [40, 41]. Optical tracking systems represent an important class of techniques for prospective motion correction. These systems generally rely on some form of camera-based setup to extract the motion parameters of the subject, often using a marker attached to the object of interest [36]. Cameras may be placed either inside the scanner bore or externally [42, 43, 44, 45]. The tracked pose information is fed to the scanner in real time, enabling dynamic adjustment of gradient orientations and RF excitation. Typically, the tracking system reports the position and orientation of the object in its own coordinate frame, which must be transformed into the scanner coordinates through a process of cross-calibration. One specialized implementation is the moiré phase tracking (MPT) system, which employs markers that produce distinctive patterns encoding both translation and rotation of the attached object [46]. This approach provides more robust and accurate motion measurements compared to conventional markers. Recent developments have focused on markerless optical tracking, where surface features of the subject are captured as a point cloud, which is then matched across frames to estimate motion parameters [47]. Markerless tracking has been applied in prospective motion correction for diffusion-weighted echo planar imaging (EPI) [48] and high-resolution neuroanatomical MRI [49], typically using commercial 3D surface tracking systems. These advances demonstrate that both marker-based and markerless optical tracking approaches can be effectively integrated into PMC workflows to mitigate motion artifacts across a variety of MRI applications.

These developments can be broadly classified as strategies to address gross patient motion in clinical MRI, which are primarily irregular or unpredictable body movements (e.g, head motion during brain MRI). Some of these techniques can also be applied effectively to internal physiological motions, such as those encountered in cardiac or respiratory imaging. However, for periodic motions of this kind, the most effective commonly used approaches are **gating techniques**, either prospective or retrospective. Gating synchronizes data acquisition with a physiological signal, so that the imaging data are acquired at predefined

phases of that cycle. Gating remains the predominant motion compensation strategy employed in preclinical small animal imaging, as active prospective or retrospective motion correction is challenging at such small scales and rapid physiological rates [7].

An overview of the various methods and considerations for addressing subject motion has been presented in the context of both clinical and preclinical MRI. In principle, these developments provide a foundation for extension to the μ MRI domain, where they are directly relevant to the challenge of imaging behaving organisms. Within this context, accurate motion estimation and tracking are critical, and the following sections provide the necessary theoretical and methodological background for the motion estimation strategies employed in this thesis.

2.3.3 Motion tracking methods

In this thesis, optical imaging-based motion tracking is employed to estimate subject motion, with a detailed description of the experimental setup provided in Section 4.3. Real-time video acquired from the imaging system serves as the primary data source for motion estimation. Motion information is extracted using computer vision techniques, which encompass algorithms for automated analysis and interpretation of visual data. Two classes of computer vision methods are used throughout this work:

1. **Deep learning**-based methods, which learn motion-related features directly from data using neural networks. This approach is used as the motion estimation strategy in the exploratory Chapter 3, which proposes a freely moving organism paradigm and examines the associated challenges. Owing to their ability to track complex motion patterns, deep learning techniques provide robust motion parameter estimation, particularly in scenarios involving unrestricted or highly variable motion.
2. **Optical flow**-based methods, which estimate motion by tracking pixel intensity variations between successive video frames [50]. These approaches offer a relatively simple implementation and enable near real-time motion tracking. They are particularly well suited to the constrained locomotion paradigm of the *in situ* treadmill, where movement occurs in a relatively controlled manner. Owing to their computational efficiency and robustness in structured environments, optical flow-based techniques have been used as the primary motion estimation strategy throughout this thesis. Further details on optical flow tracking within the scope of the thesis are provided in the following section.

2.3.4 Optical flow for motion tracking

A sparse optical flow-based algorithm was developed to track the motion of a predefined keypoint on anatomical landmarks of the specimen. Optical flow describes the apparent velocity field of brightness patterns within a sequence of images. The method relies on the brightness constancy assumption [51], which states that the intensity of a specific point in the image remains constant over time, even as its spatial position changes (see Fig. 2.9). This assumption can be mathematically formulated as:

$$I(x, y, t) = I(x + \delta x, y + \delta y, t + \delta t). \quad (2.29)$$

Here $I(x, y, t)$ is the image pixel intensity at position (x, y) and time t . Expanding this using a Taylor series and ignoring higher-order terms gives:

$$I(x + \delta x, y + \delta y, t + \delta t) \approx I(x, y, t) + \frac{\partial I}{\partial x} \delta x + \frac{\partial I}{\partial y} \delta y + \frac{\partial I}{\partial t} \delta t. \quad (2.30)$$

Since the intensity is assumed to be constant:

$$\frac{\partial I}{\partial x} u + \frac{\partial I}{\partial y} v + \frac{\partial I}{\partial t} = 0. \quad (2.31)$$

Here, $u = \delta x / \delta t$ and $v = \delta y / \delta t$ are the components of the optical flow vector, equation (2.31) is known as the optical flow constraint equation (OFCE). One common method to solve the OFCE is the Lucas-Kanade method [50]. It assumes that the flow is essentially constant in a local neighborhood of the pixel under consideration. The optical flow is then determined by solving the following system of linear equations for each pixel in a window:




Figure 2.9: Sparse optical flow tracking: keypoint moves from position 1 to position 2 under brightness constancy assumption.

$$A\mathbf{v} = \mathbf{b}, \quad (2.32)$$

$$A = \begin{bmatrix} \frac{\partial I}{\partial x}(x_1, y_1) & \frac{\partial I}{\partial y}(x_1, y_1) \\ \frac{\partial I}{\partial x}(x_2, y_2) & \frac{\partial I}{\partial y}(x_2, y_2) \\ \vdots & \vdots \\ \frac{\partial I}{\partial x}(x_n, y_n) & \frac{\partial I}{\partial y}(x_n, y_n) \end{bmatrix}, \quad \mathbf{v} = \begin{bmatrix} u \\ v \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} -\frac{\partial I}{\partial t}(x_1, y_1) \\ -\frac{\partial I}{\partial t}(x_2, y_2) \\ \vdots \\ -\frac{\partial I}{\partial t}(x_n, y_n) \end{bmatrix}. \quad (2.33)$$

Here $\mathbf{A}$ is the matrix of image gradients, $\mathbf{b}$ represents the vector of negative temporal gradients of the image intensity at different points, and $\mathbf{v} = [u \ v]^T$ is the optical flow vector. A solution can be obtained using the Moore-Penrose inverse [52, 53] $\mathbf{A}^+$ of $\mathbf{A}$:

$$\mathbf{v} = \mathbf{A}^+ \mathbf{b} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}. \quad (2.34)$$

The detailed algorithm, including keypoint selection, gradient computation, and sparse flow estimation, is presented in Algorithm 1.

Algorithm 1: Lucas–Kanade sparse optical flow algorithm for motion tracking [17].

Input: First frame I_1 , selected point $p = (x, y)$, window size $W = 20 \times 20$

Output: Tracked point position in subsequent frames

/* Initialization */

Capture the first frame I_1 from the live video stream;

Select a point $p = (x, y)$ in I_1 to track;

Define the window size $W = 20 \times 20$;

/* Pre-compute Gradients */

Compute the partial derivatives $\frac{\partial I}{\partial x}$ and $\frac{\partial I}{\partial y}$ for I_1 using Sobel operators;

Compute the temporal partial derivative $\frac{\partial I}{\partial t}$ between I_1 and the next frame I_2 ;

/* Track the point through subsequent frames */

for each subsequent frame I_{k+1} **do**

 Extract the W window centered at $p_k = (x_k, y_k)$ from I_k ;

 Compute the partial derivatives $\frac{\partial I}{\partial x}$, $\frac{\partial I}{\partial y}$ and temporal partial derivative $\frac{\partial I}{\partial t}$ within the window;

 /* Formulate the optical flow equations */

$$\begin{bmatrix} \frac{\partial I}{\partial x}(p_1) & \frac{\partial I}{\partial y}(p_1) \\ \frac{\partial I}{\partial x}(p_2) & \frac{\partial I}{\partial y}(p_2) \\ \vdots & \vdots \\ \frac{\partial I}{\partial x}(p_n) & \frac{\partial I}{\partial y}(p_n) \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = - \begin{bmatrix} \frac{\partial I}{\partial t}(p_1) \\ \frac{\partial I}{\partial t}(p_2) \\ \vdots \\ \frac{\partial I}{\partial t}(p_n) \end{bmatrix}$$

 ;

 /* Solve using least squares */

$$\begin{bmatrix} u \\ v \end{bmatrix} = - \left(\begin{bmatrix} \frac{\partial I}{\partial x} & \frac{\partial I}{\partial y} \end{bmatrix}^T \begin{bmatrix} \frac{\partial I}{\partial x} & \frac{\partial I}{\partial y} \end{bmatrix} \right)^{-1} \begin{bmatrix} \frac{\partial I}{\partial x} & \frac{\partial I}{\partial y} \end{bmatrix}^T \begin{bmatrix} \frac{\partial I}{\partial t} \end{bmatrix}$$

 ;

 Update the point position $p_{k+1} = (x_k + u, y_k + v)$ for the next frame I_{k+1} ;

end

2.4 Deep learning in MRI

The physical principles of MR image formation and the impact of motion on image reconstruction have been discussed in the previous chapters. Beyond traditional model-based correction methods, data-driven approaches have proven highly effective in mitigating motion and related artifacts in MRI. Among these, deep learning techniques have demonstrated strong performance due to their ability to learn complex patterns directly from data, offering versatile solutions for various MRI artifact correction problems [34, 54, 55].

This thesis employs deep learning methods to address two main challenges: (1) motion tracking using an *ex situ* computer vision system in Chapter 3, and (2) MR image artifact correction, which encompasses motion artifact correction in Chapter 5 and undersampling artifact correction in Chapter 6. To provide context for these applications, the following sections introduce the fundamental principles of learning representations from data, progressing from basic neural architectures to the multi-scale models employed in this work.

Deep learning is a subfield of artificial intelligence (AI) that focuses on learning representations and functions directly from data. Unlike rule-based systems, learning-based approaches model complex relationships by approximating a function that maps inputs to outputs using examples. In this sense, a deep learning model can be viewed as a parameterized function approximator that captures patterns, structures, and abstractions from large datasets. These methods are loosely inspired by the structure of biological neural systems, particularly interconnected neurons. Artificial neural networks consist of layers of computational units (“neurons”) that transform input signals through weighted connections and nonlinear activation functions. While far simpler than biological neurons, these artificial units enable hierarchical feature learning, where lower layers extract basic patterns and deeper layers progressively learn more abstract representations. This paradigm has driven major breakthroughs across diverse domains, from convolutional neural networks (CNNs) in computer vision to recurrent neural networks (RNNs) and transformers in sequential tasks like natural language processing [56, 57, 58].

This section builds in part on established deep learning literature. For a comprehensive and in-depth treatment, the reader is referred to Goodfellow et al. [56], LeCun et al. [57], and Bishop [59, 60].

2.4.1 Multi-layer perceptrons and convolutional networks

A fundamental building block of artificial neural networks is the **perceptron** [61], one of the earliest models of a computational neuron, which is illustrated in Fig. 2.10(a). Given an input vector $\mathbf{x} \in \mathbb{R}^d$, the perceptron computes a weighted sum followed by a nonlinear activation:

$$y = \phi(\mathbf{w}^\top \mathbf{x} + b), \quad (2.35)$$

where $\mathbf{w} \in \mathbb{R}^d$ denotes the weight vector, $b \in \mathbb{R}$ is the bias term, and $\phi(\cdot)$ is an activation function such as the step function, sigmoid, or rectified linear unit (ReLU). The introduction of this nonlinear activation is critical; without it, any composition of multiple layers would remain a simple linear transformation, incapable of modeling the complex, non-linear mappings required for many practical problems. Geometrically, the perceptron implements a linear decision boundary in the input space. While simple and computationally efficient, a single perceptron is limited to learning only linearly separable functions.




Figure 2.10: Artificial neural networks. **(a)** Single-layer perceptron model. **(b)** Multilayer perceptron with two hidden layers.

To model complex, nonlinear relationships, multiple perceptrons are stacked to form multilayer feedforward neural networks, also known as multilayer perceptron (MLP) (Fig. 2.10(b)). For an L -layer network, the transformation at layer ℓ is,

$$\mathbf{h}^{(\ell)} = \phi^{(\ell)} \left(\mathbf{W}^{(\ell)} \mathbf{h}^{(\ell-1)} + \mathbf{b}^{(\ell)} \right), \quad (2.36)$$

where $\mathbf{h}^{(0)} = \mathbf{x}$ is the input, $\mathbf{W}^{(\ell)}$ and $\mathbf{b}^{(\ell)}$ are the weights and biases of layer ℓ , and $\phi^{(\ell)}$ is an element-wise nonlinear activation. The network output is $\mathbf{h}^{(L)}$. Stacking layers enables the model to learn hierarchical feature representations, where deeper layers capture increasingly abstract patterns. Under mild conditions, such networks are universal function approximators, meaning they can approximate a wide class of functions to arbitrary accuracy [56, 62].

Training these networks involves learning the parameters $\theta = \{\mathbf{W}^{(\ell)}, \mathbf{b}^{(\ell)}\}_{\ell=1}^L$ by minimizing a loss function $\mathcal{L}(\theta)$. For a dataset $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$, this loss is defined as

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^N \ell(f(\mathbf{x}_i; \theta), \mathbf{y}_i), \quad (2.37)$$

where $f(\mathbf{x}_i; \theta)$ denotes the network output and $\ell(\cdot, \cdot)$ is a task-dependent loss function (e.g., mean squared error (MSE) or mean absolute error (MAE) for image-to-image regression, cross-entropy for classification). Gradients of this loss with respect to parameters are computed efficiently using the **backpropagation algorithm**, which applies the chain rule of calculus through the computational graph [56, 63].




Figure 2.11: CNN processes input MRI through convolutional layers with activation function to extract hierarchical features, interleaved pooling to reduce resolution, and final fully connected layers for task specific output.

While fully connected layers, as in MLPs, treat all input features equally, many data modalities, especially images, exhibit strong local structure. **CNNs** incorporate this prior knowledge through three key ideas: local connectivity, parameter sharing, and spatial invariance [64, 56]. Instead of a dense weight matrix, CNNs use convolutional kernels. For an input feature map X and kernel K , a 2D convolution operation is:

$$Y(i, j) = \sum_m \sum_n X(i + m, j + n) K(m, n), \quad (2.38)$$

where (i, j) indexes spatial positions. Each kernel is applied across the entire input, producing a feature map that detects specific patterns (e.g., edges or textures). Parameter sharing

drastically reduces the number of learnable weights compared to fully connected layers, improving statistical efficiency and generalization.

A typical convolutional layer can be written as,

$$\mathbf{H}^{(\ell)} = \phi(\mathbf{K}^{(\ell)} * \mathbf{H}^{(\ell-1)} + \mathbf{b}^{(\ell)}), \quad (2.39)$$

where $*$ denotes convolution and $\mathbf{H}^{(\ell)}$ is a stack of feature maps. CNN architectures often include pooling operations, such as max-pooling,

$$Y(i, j) = \max_{(m, n) \in \mathcal{R}_{ij}} X(m, n), \quad (2.40)$$

which reduce spatial resolution and lower computational cost, while retaining most salient features, introducing a degree of translation invariance. By stacking convolutional, activation, and pooling layers, CNNs learn hierarchical spatial representations particularly well suited to visual and audio data [57, 56]. An example CNN processing a brain MRI input is shown in Fig. 2.11. The network begins with convolutional layers and associated activations, as discussed above, which extract hierarchical spatial features from the image. These are then flattened (vectorized) and passed through fully connected layers to produce task-specific outputs, such as binary tumor classification (e.g., presence or absence of a tumor) or multi-class predictions.

The progression from perceptrons to deep convolutional architectures reflects a shift from simple linear classifiers to highly structured, hierarchical models that exploit nonlinear representation learning and domain-specific inductive biases. While standard CNNs excel at extracting hierarchical features, they can struggle with tasks requiring precise localization and dense pixel-level predictions, such as image-to-image regression, semantic segmentation, or image analysis. This can be particularly pronounced in μ MRI, where the images often contain fine anatomical structures at micrometer resolution, subtle intensity variations, and complex morphology. Pooling layers, although useful for computational efficiency and robustness to small translations, inherently reduce spatial resolution and can discard these fine-grained details.

To address such challenges, specialized architectures have been developed that preserve and recover spatial information while leveraging the hierarchical feature learning of CNNs. Among these, the **U-Net** architecture [65] and variants are widely adopted for image-to-image mapping tasks. At a high level, U-Net follows an encoder-decoder design: the contracting (encoder) path captures context through successive convolution and pooling operations, while the expansive (decoder) path progressively upsamples feature maps and integrates high-resolution features from the encoder via skip connections. This structure allows the network to retain fine-grained spatial information while utilizing deep hierarchical representations, making it highly effective for dense prediction problems. Further details on this will be provided in the context of the specific application in Section 5.3.3.

2.4.2 Transfer learning

All deep learning methods used in this thesis operate within the framework of supervised training, where labeled data are used to optimize the network parameters $\theta = \{\mathbf{W}^{(\ell)}, \mathbf{b}^{(\ell)}\}_{\ell=1}^L$, as described in Section 2.4.1. In the context of insect μ MRI and optical imaging, the limited availability of labeled data poses a challenge for training networks from scratch. To address this, a transfer learning strategy is employed: a model pretrained on a large source dataset (e.g., human brain MRI or natural image datasets) is used to initialize the network for the target task, enabling the reuse of previously learned features, and thereby improving both training efficiency and generalization on smaller datasets, as illustrated in Fig. 2.12.




Figure 2.12: Transfer learning pipeline.

Formally, let $\mathcal{D}_S = \{(\mathbf{x}_i^S, \mathbf{y}_i^S)\}_{i=1}^{N_S}$ denote the source dataset, and $\mathcal{D}_T = \{(\mathbf{x}_j^T, \mathbf{y}_j^T)\}_{j=1}^{N_T}$ denote the target dataset, with $N_T \ll N_S$. The network f_θ is first trained on $\mathcal{D}_S$ to learn the mapping $f_\theta : \mathbf{x}^S \mapsto \mathbf{y}^S$. For transfer learning, the parameters θ learned from the source dataset are used to initialize the network, which is then fine-tuned on $\mathcal{D}_T$ to adapt the learned representations to the target domain [66, 67]. Early layers of the network capture generalizable spatial features, while later layers are adjusted to the specific characteristics of the insect data [68].

In summary, transfer learning is used for two applications in this thesis:

1. **Motion tracking:** In the early exploratory Chapter 3, the focus is on accurately estimating both the motion trajectory and the pose of a freely moving insect in an *ex situ* platform. A model pretrained on a large-scale natural image dataset has been used here, which provides effective and computationally efficient feature initialization for transfer to specialized domains. This model is then fine-tuned using small

dataset of custom optical images of insects, thereby allowing adaptation to the specific characteristics of the target insect data. Further details on this setup are provided in Section 3.2.1.

2. **μ MRI artifact correction:** In the context of insect μ MRI, two major challenges are addressed: (1) motion-induced artifacts in Chapter 5 and (2) artifacts arising from undersampled reconstruction in Chapter 6. Although these problems are treated in separate chapters, they share a common methodological framework within the broader narrative of this work. Both artifact types are addressed as image-to-image regression problems using a modified U-Net–based architecture [65]. The model was pretrained on human brain MRI datasets to leverage general spatial features, and subsequently fine-tuned on the insect μ MRI data to adapt to the smaller, domain-specific dataset. The early convolutional layers learn generalizable patterns transferable across datasets, while the customized architecture enables more accurate correction of motion and undersampling artifacts. The complete details of the problem formulation, training, fine-tuning, and implementation are presented in the respective chapters.

3 | Freely moving organism

Behavioral arenas provide a controlled experimental interface through which an organism's physiological and behavioral responses can be investigated under laboratory conditions. These environments are typically designed to minimize external constraints, thereby enabling spontaneous movement, interactions, and internally driven behaviors while maintaining sufficient experimental control. Such arenas are standard across experimental modalities, and exist along a continuum of invasiveness: fully freely moving animals at one end, versus complete immobilization, allowing only subtle responses (e.g., eye movements) at the other. MRI traditionally occupies the highly constrained extreme, with subjects fully immobilized and anesthetized.

The central objective of this thesis, namely “MRI of behaving organisms”, therefore requires moving beyond this traditionally constrained regime toward experimental setups that accommodate movement during imaging. The chapter begins by exploring the feasibility of a freely moving organism platform within the context of μ MRI. The primary considerations include the design of a suitable platform that permits a sufficient degree of freedom to enable natural movements, integrated with a computer vision-based real-time motion tracking system to accurately capture and compensate for movement during imaging. All of this must operate within the constraints of an ultra-high-field preclinical MRI scanner. The practical challenges and limitations inherent to such a system are analyzed, leading to the development of a minimally constrained *in situ* treadmill paradigm and the selection of an appropriate organism for testing the underlying hypothesis. Together, these design choices transform the initial conceptual framework into a practical and experimentally viable paradigm, in which motion and dynamics are treated as intrinsic components of the measurement, thereby establishing the foundation for the remainder of the thesis.

3.1 The scale-of-motion problem

While many of the assumptions and motion models discussed could, in principle, be extended to μ MRI studies of active or freely behaving organisms, the case of a truly freely moving animal introduces a fundamental paradigm shift, with few direct analogues in established clinical or preclinical MRI. In this regime, the magnitude of the animal's displacement can

often be comparable to, or exceed, the characteristic length scales of the coil geometry and the associated spatial field distributions, including RF, gradient, and static magnetic fields.

The primary technical challenge lies in optimizing the coil geometry and filling factor, with motion as an additional constraint. A surface coil–integrated platform can offer the necessary degrees of freedom to accommodate model organisms like *C. elegans*, whose motion is largely restricted to a two-dimensional plane. Such experimental configurations are standard in modalities or platforms based on optical microscopy or optogenetics [69, 70]. In MRI, if the organism’s trajectory can be estimated and tracked with sufficient accuracy, prospective motion correction (PMC) techniques can, in principle, be used to dynamically update the imaging coordinate system and continuously follow a chosen orthogonal two-dimensional slice. This concept was demonstrated *in silico* by Reischl et al. [71], as illustrated in Fig. 3.1, and serves as the initial motivation and starting point for this thesis project.




Figure 3.1: Prospective motion correction in μ MRI: The imaging volume dynamically tracks the subject’s motion, enabling MR imaging of freely moving specimens in an *in silico* study. (Adopted from Reischl et al. [71]).

However, realizing such a platform in practice presents several challenges. Even under static conditions, MRI of *C. elegans* or similar sized models remains beyond the capabilities of current MRI systems (as of 2026), primarily due to severe SNR limitations corresponding to the extremely small voxel volumes. The chosen animal model must be sufficiently large to yield adequate SNR within a selected 2D slice or 3D volume with reasonable signal averaging. At the same time, the platform or the behavioral arena must allow the organism

to move freely. A single coil large enough to encompass such an arena would suffer from poor filling factor, resulting in negligible signal sensitivity.

Furthermore, moving beyond immobilized subjects introduces complex field interactions. Although this is also true for clinical MRI, motion relative to encoding gradients is the primary source of k -space inconsistencies and the resulting image artifacts. However, in a freely moving paradigm, the subject also moves through spatial variations in the RF excitation (B_1) and static (B_0) magnetic fields. Central to these B_0 variations is magnetic susceptibility χ . At the microscopic scale, the interface between the organism's body and its environment (e.g., agar, water, or air) creates local magnetic field gradients. As the specimen moves, these susceptibility-induced field distortions shift dynamically. This not only causes spatial warping and signal voids but also results in a moving B_0 inhomogeneity that can outpace the capabilities of traditional static shimming. These effects represent a significant hurdle for micro-scale, high-field imaging where even sub-millimeter displacements can lead to total signal dephasing.

Additional challenges lie in achieving accurate motion estimation and tracking of the organism within an *in situ* imaging platform, which is a critical prerequisite for enabling motion-resolved MR acquisition. MR-based approaches, such as navigator techniques, exhibit inherent limitations when estimating motion with multiple degrees of freedom at the spatial accuracy and temporal resolution required for behavioral tracking. In contrast, computer vision-based systems currently provide state-of-the-art performance in motion tracking and therefore represent a natural and well-suited solution for this task, as well as for the broader objective of behavioral phenotyping, provided that an MR-compatible system can be developed and integrated for high-field MRI systems.

3.2 *Ex situ* platform: freely moving insect

The spatial scale of any behavioral arena is fundamentally limited by the physics of the measurement modality. Instrumentation imposes physical constraints that restrict an organism's degrees of freedom and, by extension, its accessible behavioral repertoire. A central challenge in experimental design is therefore the extent to which a laboratory configuration can faithfully approximate the organism's native ecological context.

The progression of the thesis is grounded in the central challenge of enabling MRI measurements in the presence of organism motion. At this stage, the overall engineering challenge can be divided into two main components: (1) the MR measurement problem, which was discussed previously in the context of a prior *in silico* study on *C. elegans*, and (2) the motion estimation problem.

The MR measurement problem is primarily defined by the constraints imposed by the microscopic spatial scales of μ MRI and the resulting SNR limitations (discussed further in Section 3.3.2). These challenges necessitate the selection of a biological specimen or

model system with sufficient tissue volume to generate detectable MR signals. While larger animals naturally provide higher signal due to their greater size, insects offer a practical and accessible alternative for validating the proposed paradigm, particularly in settings where vertebrate facilities are unavailable. Their combination of manageable size, measurable tissue volume, and complex behavioral repertoire makes them an ideal first model. MRI has already emerged as a powerful non-invasive tool for the structural characterization of insects and other small invertebrates, with demonstrated applications in selected functional studies [12, 72]; however, prior work has largely been limited to static or immobilized specimens, as is typical in MRI.

Robust motion estimation is a central component of effective motion compensation strategies, regardless of the specific organism or the correction approach. This consideration motivated addressing the motion estimation as the primary step, prior to tackling the MR imaging challenges associated with the prospective animal model. Importantly, this does not require an *in situ* MRI setup at this stage. Instead, the motion tracking methodology can be developed and validated using a computer vision-based system implemented on a simple *ex situ* experimental platform as follows.

3.2.1 *Ex situ* platform and ladybug beetle

Coccinella septempunctata, commonly known as the seven-spot ladybird beetle or ladybug, was selected as the specimen for developing and validating the vision-based motion tracking system. Figure 3.2 presents a labeled illustration highlighting the external morphology of the ladybug. This species was chosen primarily because its locomotion appears predominantly rigid when viewed from a top-down perspective, which simplifies motion estimation and tracking and closely aligns with the rigid-body motion models discussed earlier (Section 3.1) in the context of prospective motion correction (PMC). The primary motivation for this choice was to explore whether such motion characteristics could ultimately enable the application of PMC-based approaches for MRI of an active insect. When confined to a low-height enclosure, the beetle's movement is effectively restricted to a quasi-two-dimensional plane, allowing planar tracking methods to capture its behavior while maintaining biologically relevant locomotion. Specimens were collected from a local garden and maintained in a custom enclosure at room temperature (approximately 22–25 °C), with a diet of fruits and water.

The *ex situ* insect platform was fabricated using a 3D printer (Ultimaker 3, polylactic acid (PLA)) with dimensions of $34 \times 34 \times 6 \text{ mm}^3$ and enclosed by a transparent plexiglas lid. A schematic of the setup is shown in Fig. 3.3(a). Behavioral motion data were recorded using a commercial webcam (Microsoft HD-3000) positioned above the platform, as shown. This controlled configuration enables precise acquisition of locomotor behavior while constraining the beetle to a planar motion regime. The following sections detail the motion tracking and validation workflow, focusing on a deep learning (DL)-based approach for estimating and analyzing insect locomotion within the *ex situ* platform.




Figure 3.2: External morphology of a ladybug. (Modified from [73]).

3.2.2 Deep learning-based motion tracking

Deep learning methods currently represent the state of the art for various computer vision problems, particularly for tasks such as object detection, tracking, and pose estimation. Compared to traditional image-processing techniques, deep neural networks exhibit substantially improved robustness to noise, variations in illumination, and complex backgrounds, resulting in markedly higher accuracy. Consequently, vision-based methods have found widespread applicability across diverse areas of science and technology. In particular, significant progress has been made in applying deep learning-based computer vision algorithms to behavioral studies of model organisms [74].

The problem addressed here involves real-time estimation of the two-dimensional position (x, y) and the in-plane pose (orientation) θ of a freely moving insect. While conventional image-processing approaches, such as thresholding, centroid-based detection or even optical flow-based techniques can provide reliable position estimates under controlled conditions, accurate pose estimation, especially in a markerless setting, remains highly challenging. To overcome these limitations, a deep CNN-based approach capable of jointly inferring position and orientation from image data in real time has been used.

Specifically, the you only look once (YOLOv8) model [75], a deep learning model, pretrained on the Microsoft common objects in context (COCO8)-Pose dataset [76], is utilized. This approach leverages the paradigm of transfer learning, in which a CNN trained on a large, problem-agnostic dataset can be fine-tuned using a smaller, task-specific labeled dataset in a supervised manner, as discussed in Section 2.4.2. Such a strategy enables efficient learning while significantly reducing the amount of training data required, often by orders of magnitude, compared to training a deep learning model from scratch.

3.2.2.1 YOLOv8 model architecture

YOLOv8 is a family of deep learning-based computer vision models originally proposed by Redmon *et al.* [77] for real-time object detection. The you only look once nano (YOLOv8n) variant used here is a lightweight and optimized implementation developed by Ultralytics [75], capable of object detection, pose estimation, classification, and segmentation. Here, YOLOv8n was employed for keypoint detection, specifically to identify two anatomical landmarks (head and abdominal apex) of the insect. These keypoints provide an initial estimate of the organism's two-dimensional position and orientation. Using pretrained weights significantly reduces the amount of task-specific training data required.

At a high level, YOLOv8 architecture consists of four main components: input, backbone, neck, and head. The input module standardizes images through resizing, normalization, and data augmentation. The backbone extracts hierarchical feature representations from the image, while the neck combines multi-scale features to enhance detection performance. The head predicts object classes, bounding boxes, and keypoint locations. The architecture is designed to jointly learn detection and pose estimation, sharing features across tasks, which improves efficiency and accuracy. Comprehensive architectural details and implementation procedures are provided in the official Ultralytics documentation [75]. While YOLOv8 was a practical and effective choice at the time of the experiments (2023), it is one of several pretrained models suitable for these tasks, and more recent architectures and solutions have since been developed that may offer improved performance.

3.2.3 *Ex situ* tracking characterization

Three videos of insect behavior were recorded and used to generate the training, validation, and testing datasets. Individual frames were extracted using a custom Python script. The dataset was split into 150 training images, 50 validation images, and 100 test images. The overall DL pipeline based on the pretrained YOLOv8 model is illustrated in Fig. 3.3(b). All images were manually annotated with head and tail (abdominal apex) keypoints using a web-based annotation tool (`makesense.ai`).

The objective of the tracking system is to estimate the real-time two-dimensional position (x, y) and orientation θ of the insect. The model is trained to detect the head and tail as two keypoints in each frame, with the line joining these points defining the rigid-body pose of the insect in the image plane. Although this work focuses on two-dimensional pose estimation, the approach can be readily extended to three-dimensional motion tracking or augmented with additional keypoints for more complex motion models.

The model was initialized with weights pretrained on the COCO8-Pose dataset [76] and subsequently fine-tuned on the task-specific dataset. Fine-tuning was performed using 150 labeled images with an initial learning rate of 3×10^{-4} . Both training and validation losses converged after 100 epochs, at which point the learning rate had decayed to 5.96×10^{-5} ,




Figure 3.3: Deep learning-based motion tracking. **(a)** Schematic illustration of the *ex situ* insect tracking setup. **(b)** Schematic of the deep learning-based tracking pipeline, showing fine-tuning of a pre-trained model using custom insect data.

with a total training time of approximately one hour. The trained model was evaluated on a test dataset comprising 100 non-contiguous image frames, sampled from the test video at intervals of $\Delta t = 0.27$ s.

The inference time for a single frame was approximately 13 ms on the workstation (hardware specifications are provided in Section 5.3.4). Quantitative evaluation was performed by computing the Euclidean distance between predicted and manually annotated keypoints. For the test dataset, the mean localization error was approximately 7 ± 3 pixels (0.242 ± 0.104 mm), as shown in Fig. 3.4, where Fig. 3.4(a) shows the model-estimated (x, y) trajectory of the insect, Fig. 3.4(b) presents the histogram of Euclidean distance errors, Fig. 3.4(c) shows the pose estimation (θ) results, and Fig. 3.4(d) provides representative test images highlighting keypoint detection with both small and large localization errors. Figure 3.5 illustrates deep learning-based pose tracking on a separate *ex situ* ladybug test video, with sampled frames and timestamps highlighting the temporal progression of behavioral motion.

3.2.4 Kalman prediction

Motion estimation accuracy is a critical factor in the design of motion compensation strategies for MRI. Regardless of the tracking modality, the estimation process is affected by inherent system latency and measurement noise. In computer vision-based tracking systems, this latency arises from image acquisition at the camera sensor and propagates through subsequent processing stages, including keypoint detection by the algorithm. As a result, a mismatch occurs between the measured motion and the object's true motion at the same time point. In prospective or adaptive motion correction systems in MRI, where latency is significant, this mismatch may necessitate predicting the subject's future position to compensate, enabling more proactive motion correction [78, 79, 71].




Figure 3.4: *Ex situ* tracking validation on test dataset ($n = 100$). **(a)** Model-estimated (x, y) position trajectory of the insect. **(b)** Histogram of the Euclidean distance error between model-predicted keypoints and manual annotations, showing a mean error of 7 ± 3 pixels. **(c)** Pose (θ) estimation results obtained from the model. **(d)** Representative test images illustrating keypoint detection with both small and large localization errors (green crosses: manual annotations; red crosses: model detections).




Figure 3.5: Deep learning–based tracking of ladybug body pose in the *ex situ* setup. Sampled video frames with timestamps show the motion trajectory.

To investigate the feasibility of short-term motion prediction and assess the temporal robustness of the tracking framework, a Kalman filter–based algorithm [80] was implemented for recursive state estimation under Gaussian noise assumptions. A linear Kalman filter with a constant-velocity motion model was implemented in Python using the OpenCV library [81]. The system is parameterized by a state transition matrix A , a measurement matrix H , and a process noise covariance matrix Q . The state vector consists of four variables $(x, y, \dot{x}, \dot{y})$, representing two-dimensional position and velocity, while the measurement vector contains the observed positions (x, y) . The transition matrix propagates the state forward in time under a linear motion assumption, while model uncertainty is captured through the process noise covariance matrix Q .

Let the true system state at time t be denoted by $\mathbf{x}_t$, its estimate by $\hat{\mathbf{x}}_t$, and the noisy measurement by $\mathbf{z}_t = H\mathbf{x}_t + \mathbf{v}_t$. The predicted state estimate is given by

$$\hat{\mathbf{x}}_t^- = A\hat{\mathbf{x}}_{t-1}, \quad (3.1)$$

and the corrected (updated) estimate is computed as,

$$\hat{\mathbf{x}}_t = \hat{\mathbf{x}}_t^- + K_t (\mathbf{z}_t - H\hat{\mathbf{x}}_t^-), \quad (3.2)$$

where K_t denotes the Kalman gain.

The algorithm was evaluated using a subset of the test dataset comprising 1000 image frames (approximately 27 s of recorded motion). The deep learning–based tracking model provided midpoint position estimates, defined as the mean location between the detected

head and tail keypoints. These served as noisy measurements for the Kalman filter. To examine the effect of increasing temporal intervals on estimation accuracy, the trajectory was artificially undersampled, and the resulting estimation error was quantified using the mean absolute error (MAE). This analysis provides an indication of the temporal horizon over which short-term predictive modeling remains reliable.




Figure 3.6: Kalman filter-based state estimation under temporal subsampling. **(a)** Comparison of the insect’s X-position trajectory obtained from the deep learning-based tracker and the corresponding Kalman filter state estimates for an effective temporal interval of 0.27 s. **(b)** MAE between the detected positions and Kalman estimates as a function of the temporal interval. Error bars indicate the standard error of the mean.

A representative case with a sampling factor of 10, corresponding to an effective prediction interval of 0.27 s, is shown in Fig. 3.6(a). The plot compares the deep learning-based position estimates with the corresponding Kalman filter state estimates. For this interval, the MAE between the detected positions and Kalman estimates was approximately 0.438 mm. The MAE and standard error of the mean for different temporal intervals are summarized in Fig. 3.6(b), providing an indication of the effective temporal horizon over which short-term predictive modeling remains reliable.

The Kalman filter yields more accurate state estimates under approximately linear motion conditions. Deviations from linear motion, such as abrupt directional changes, introduce velocity updates that can manifest as transient overshoot in the estimated trajectory. These effects become more pronounced as the temporal interval increases. Kalman filter-based approaches for smoothing motion trajectories and enabling short-term forward state propagation have been widely used in real-time animal tracking [82], as well as in prospective and retrospective motion correction within clinical MRI environments [83, 84, 85, 86]. While data-driven deep learning approaches could in principle also be applied for motion forecasting, classical Kalman filtering provides a comparatively simple, transparent, and computationally efficient solution that requires minimal training data and is well suited for real-time implementation.

Compared to the often abrupt and unpredictable rigid-body motion encountered in clinical MRI, the two-dimensional locomotion of small insects is typically piecewise smoother and more gradual within the central region of the experimental platform, although rapid directional changes may occur near boundaries or during specific behaviors. These motion characteristics suggest that small organisms are, in principle, well-suited for PMC strategies, in which MRI acquisition can be continuously adapted based on real-time tracking, as demonstrated in the *in silico* study by Reischl et al. [71].

From a motion-tracking perspective, current optical tracking technologies are sufficiently advanced to support the envisioned PMC-based implementation of MRI in freely behaving organisms. The primary challenges arise from the interaction between the constraints of the MR measurement problem and the practical requirements of precise *in situ* motion tracking within a narrow-bore ultra-high-field scanner, as discussed in the following sections.

3.2.5 Limitations

The ladybug (*Coccinellidae septempunctata*) chosen here primarily serves as a test system for validating the computer vision-based motion tracking and prediction. Although no explicit MR imaging experiments were conducted on this particular specimen, the challenges associated with freely moving organisms, described earlier in this chapter, remain relevant. While the tissue volume of the insect is larger than that of organisms such as *C. elegans*, limitations related to coil filling factor are expected to reduce the effective SNR in *in vivo* imaging.

In addition, designing an *in situ* behavioral platform that allows sufficient freedom of movement introduces further complexity. For the motion estimation problem, the integrated optical imaging system must simultaneously provide a large FOV and adequate spatial resolution, which complicates the overall system design. In the case of the horizontal preclinical MRI system, the optical axis requires additional relay optics [87], and accommodating these components within the magnet bore, limited to a diameter of 60 mm, impose significant practical constraints on the design and layout of both the imaging and behavioral platform components.

These challenges underscore the practical limitations of a freely moving platform for MRI experiments. To balance biologically relevant behavior with experimental feasibility, a minimally constrained ***in situ* spherical treadmill**-based platform offers a compromise for imaging live, active specimens. In this setup, the insect is tethered on its dorsal surface, effectively fixing its body relative to the imaging frame while allowing natural locomotion relative to its own body on the treadmill. This arrangement relaxes the design requirements for both the behavioral platform and the computer vision-based tracking system. Such treadmill-based platforms have been successfully implemented in other imaging modalities, demonstrating their effectiveness for studying active behavior under controlled

conditions. The following section introduces the proposed platform in detail, and the remainder of the thesis focuses on its development along with the associated methods for motion-compensated MRI of active insects.

3.3 Constrained locomotion: treadmill

Unrestricted movement within a naturalistic environment remains challenging for conventional MRI, as discussed in the previous sections. Treadmill-based paradigms, which impose only minimal constraints on natural movement, provide a practical compromise between fully unrestricted behavior and complete immobilization. Such setups have been widely adopted across disciplines including biomechanics, exercise physiology, metabolic research, and neuroethology [88, 89, 90, 91]. Representative examples of such experimental setups are illustrated in Fig. 3.7. These approaches share a common objective: to allow voluntary movement within a controlled, minimally stressful environment while maintaining spatial consistency that facilitates reliable access to experimental instrumentation and protocols. Although not directly related to behavioral imaging, clinical MRI studies have employed MR-compatible treadmill systems within the imaging workflow for pre-scan exercise or stress induction outside the magnet bore. This approach enables controlled cardiovascular and exercise stress studies, with imaging performed immediately afterward [92, 93].




Figure 3.7: Treadmill-based experimental setups in the literature. **(a)** Head-fixed *Drosophila* walking on a spherical treadmill during whole-cell patch-clamp recording with optogenetic stimulation [94]. **(b)** An air-cushioned spherical treadmill used for navigation studies in *Cataglyphis* desert ants [95].

Although Section 3.2.5 highlighted the challenges of implementing a fully freely moving paradigm in MRI, the insect remains a highly relevant model for the thesis hypothesis. The fundamental considerations regarding SNR and tissue volume outlined earlier remain valid. Building on these insights, the following sections introduce a new insect model, which serves as the core biological system enabling the methods developed to test the central hypotheses. Further challenges and methodological considerations associated with this model will be discussed in the context of the *in situ* treadmill-based experimental framework as presented in the following sections.

3.3.1 The Sun beetle: *Pachnoda marginata*

The discussion of challenges in a freely moving MRI paradigm motivated the design and development of the proposed *in situ* treadmill platform. Unlike conventional biology-first approaches, where an established animal model (e.g., rat or mouse) drives the development of motion correction or imaging instrumentation, this thesis adopts a deliberate instrumentation or methods-first strategy. Here, the treadmill design preceded biological model selection, as a means for exploring viable pathways to validate the “ μ MRI of the behaving organism” objective. Several specimens were initially considered within this context; insects ultimately emerged as the optimal proof-of-concept. Demonstrating the viability of this approach in insects establishes a foundation that may be translated to more complex and behaviorally relevant model systems, within the fundamental constraints of MRI discussed throughout this work.

While the ladybug (*Coccinellidae septempunctata*) served as a useful initial model for validating vision-based tracking (Section 3.2.1), its limited tissue volume is insufficient for robust MRI. Initially, the relatively small size of the ladybug was a prerequisite for a proposed freely moving experimental setup. However, the transition to a tethered treadmill system has relaxed these size constraints, permitting the use of larger insects. Consequently, the sun beetle (*Pachnoda marginata*) was selected as the primary tethered organism for this thesis due to its favorable size and biomechanical characteristics. This species is readily available, easy to maintain, and provides a relatively better filling factor for the commercial volume RF coil used, which is essential for achieving the SNR required to resolve internal anatomical features. Furthermore, its distinct segmented morphology, featuring a clear head, thorax, and abdomen, facilitates identification of stable landmarks for precise motion tracking and estimation.




Figure 3.8: External morphology of a sun beetle.

In addition, *P. marginata* exhibits relatively controlled locomotion and physiological resilience to the high-field MRI environment, making it an ideal surrogate for validating

motion compensation strategies. The beetle's complex tissue composition, including chitin, air sacs, and hemolymph, introduces substantial magnetic susceptibility (χ) variation, which presents significant challenges in imaging but provides a rigorous testbed to ensure the robustness of the imaging frameworks before applying them to widely studied but smaller, more signal-limited models like *Drosophila melanogaster* in the future.

Figure 3.8 presents a dorsal view of the adult *Pachnoda marginata*, with the head, thorax, and abdomen clearly delineated. Adult specimens typically measure approximately 20 mm to 25 mm in length, 10 mm to 12 mm in width, and 8 mm to 10 mm in height, providing a favorable size for high-resolution μ MRI.

3.3.1.1 Sun beetle on the treadmill

Figure 3.9 presents multiple perspective views of the tethered sun beetle mounted on the treadmill in the experimental configuration employed in this thesis. The tethering arrangement enables sustained locomotion while maintaining compatibility with the mechanical and electromagnetic constraints of an ultra-high-field 15.2 T preclinical scanner equipped with a commercial RF volume coil. The system is designed to preserve behavioral activity within the spatial limitations imposed by the imaging environment. Further details on the design, development, and characterization of the experimental setup are provided in Chapter 4.



Figure 3.9: Tethered sun beetle on a spherical treadmill.

3.3.1.2 Anatomical reference MRI

Throughout this thesis, 2D MR imaging of the beetle is presented, focusing on sagittal slices aligned with the longitudinal body axis and axial slices through the thorax (position references in Fig. 3.10(c)). The sagittal orientation captures the major internal compartments of the beetle, enabling visualization of the head capsule, the foregut and anterior midgut, and the posterior digestive tract extending into the abdomen. These structures are surrounded by air-filled tracheal networks and hemolymph-filled cavities, which may not be visible at this resolution, but introduces strong magnetic susceptibility contrasts if appropriate imaging sequences are not selected.

High-resolution anatomical reference images of the inactive insect are shown in Figure 3.10(a) and (b), with the corresponding acquisition parameters detailed in Section 4.5. Axial thoracic images emphasize the dense musculature responsible for locomotion, including the dorsal longitudinal muscles, dorsoventral muscles, and basalar muscles. These muscle groups occupy a substantial fraction of the thoracic volume and are mechanically coupled to the rigid exoskeleton.




Figure 3.10: Anatomical reference image of *Pachnoda marginata*. (a) Sagittal MRI showing the abdomen. (b) Axial MRI showing the thorax. (c) MRI slice position. (Modified from [17].)

In this thesis, biological specimens are primarily used in an engineering sense, serving as functional models (or “biological phantoms”) for testing and benchmarking the imaging frameworks. While biological context is provided to guide interpretation, the emphasis remains on methodological development rather than detailed biological or entomological analysis. Magnetic resonance methods have been used in entomological and broader zoological research, and the detailed discussions of these applications and their biological implications can be found in the articles by Ziegler et al. [12] and Hart et al. [72].

Throughout this thesis, the terms insect and beetle may be used interchangeably. The phrase “behaving model organism” or the related variations has been used to emphasize that the insect is alive, active, and capable of performing normal behaviors, rather than to imply its

use as a model in the conventional biological or ethological sense. No conclusions regarding the organism's biological or behavioral significance are drawn from this terminology.

3.3.2 Insect μ MRI challenges and considerations

The methodological advances presented in the thesis deal with the case of motion in the μ MRI regime. A key underlying assumption is that baseline static imaging provides ideal reference measurements with respect to reconstructed image quality. In practice, however, this assumption may not always hold, as static acquisitions can themselves be affected by various nonidealities, and necessary choices have to be made in terms of the experimental parameters, particularly the imaging sequence selection.

μ MRI is fundamentally constrained by SNR limitations arising from the scaling of voxel volume with spatial resolution, as briefly discussed in Section 3.2.5. For isotropic voxels, the detectable MR signal is directly proportional to the voxel volume V . The methods developed in this thesis employed a commercial ^1H RF volume coil with a birdcage configuration (inner diameter 35 mm; Bruker Biospin, Ettlingen, Germany) at 15.2 T. For conductive biological samples, the receive chain operates in the **sample-noise dominated regime** [96]. Here, the image SNR per voxel scales approximately linearly with the voxel volume for fixed bandwidth and sequence parameters [97],

$$\text{SNR} \propto V \propto (\Delta x)^3, \quad (3.3)$$

where Δx denotes the isotropic voxel dimension. In μ MRI applications, Δx typically ranges from 10 μm to 100 μm , resulting in a massive SNR penalty compared to conventional clinical MRI with voxel sizes on the order of 1 mm to 2 mm. For example, reducing the voxel edge length from 1 mm to 10 μm corresponds to a 10^6 -fold decrease in voxel volume and hence SNR per voxel, necessitating long acquisition times and ultra-high field strengths to achieve viable image quality.

Operation at ultra-high magnetic fields (e.g., 15.2 T) provides an SNR gain; in the sample-noise limit, SNR can exhibit super-linear dependence on B_0 (up to approximately $B_0^{7/4}$ [96] in some experimental regimes). However, these gains only partially compensate for the severe signal loss imposed by voxel miniaturization. Furthermore, at high fields, the SNR advantage is increasingly constrained by physiological noise and shortened T_2^* relaxation times, diminishing the marginal benefit of further increases in B_0 . As a result, efficient sequence design and signal utilization become critical for achieving better μ MRI image quality.

3.3.2.1 Susceptibility-induced inhomogeneities

In addition to SNR limitations, μ MRI is particularly susceptible to magnetic field inhomogeneities arising from local susceptibility variations ($\Delta\chi$). At 15.2 T, even minor differences

in magnetic susceptibility between adjacent tissues or materials are magnified. In insect imaging, prominent air–tissue interfaces, such as tracheal networks and the chitinous exoskeleton, introduce strong local susceptibility gradients (G_χ) [13, 72].

While the absolute field shift ΔB_0 is largely independent of voxel size, the resulting intravoxel dephasing is determined by the internal field gradient, $G_\chi = \nabla(\Delta B_0)$, across the voxel. The characteristic phase spread ϕ within a single voxel may be approximated by

$$\phi \approx \gamma (G_\chi \cdot \Delta x) \text{ TE}, \quad (3.4)$$

where γ is the gyromagnetic ratio and TE is the echo time. These internal field variations lead to intravoxel dephasing, causing signal attenuation and distortion of echo amplitudes. In severe cases, phase accrual can exceed 2π within a single voxel over typical echo times, resulting in essentially irreversible signal loss. Such effects are particularly detrimental for gradient recalled echo (GRE)-based imaging, where susceptibility-induced intravoxel dephasing directly leads to T_2^* -related signal loss. In multi-echo acquisitions, this effect can accumulate across the echo train, corrupting T_2 - and T_2^* -weighted contrast and compromising image fidelity, particularly at ultra-high fields.




Figure 3.11: Sagittal MRI of sun beetle using (a) FLASH, and (b) RARE (anatomical reference).

These combined effects are evident in the sagittal MRI acquired with the GRE-based fast low angle shot (FLASH) sequence shown in Fig. 3.11(a), particularly in comparison with the rapid acquisition with relaxation enhancement (RARE) acquisition in Fig. 3.11(b) of the same slice of a dormant insect. It should be noted that this serves only as a qualitative illustration of the image characteristics typically obtained with a GRE-based sequence in insect imaging, using the RARE image as an anatomical reference. It does not constitute a direct comparison between the two sequences, as their acquisition parameters differed and were not individually optimized for maximal image quality. Imaging parameters can be found in Appendix B

3.3.2.2 Implications for sequence selection

These severe SNR penalties and susceptibility sensitivity place stringent constraints on the selection of imaging sequences for insect μ MRI. To mitigate susceptibility-induced dephasing, echo times must be kept short, while SNR efficiency must be maximized to offset voxel-scale signal loss.

These requirements motivate the use of RARE sequences for imaging throughout this work. By employing a train of 180° refocusing pulses, RARE sequences refocus static field inhomogeneity effects and suppress T_2^* -related signal loss, thereby recovering coherence lost to reversible dephasing, as described in Section 2.2.3.2. However, while RARE compensates for susceptibility-induced dephasing, it does not eliminate intrinsic spin–spin relaxation. The transverse magnetization still decays exponentially according to T_2 , imposing a fundamental limit on echo train duration.

In principle, increasing the RARE factor (or ETL) enables substantial acceleration by acquiring multiple phase-encoding lines per excitation, thereby reducing total scan time. However, both intrinsic T_2 decay and residual dephasing across the echo train impose a practical upper bound on the usable echo train length. As the echo train extends, later echoes experience reduced signal amplitude due to T_2 relaxation and accumulated phase dispersion, leading to image blurring and contrast modulation across k -space.

Consequently, the parameters selected for the RARE acquisitions in the following chapters reflect a compromise between acceleration, SNR efficiency, and susceptibility robustness. Echo spacing, effective echo time, and echo train length were optimized primarily through empirical experimentation, guided by the physical constraints outlined above. These settings define the general acquisition strategy, while further methodological details specific to each chapter are presented in the corresponding sections.

3.4 Conclusions

This chapter developed the core hypothesis of the thesis by translating the exploratory findings into a concrete experimental framework. It began with the conceptual foundations of imaging freely moving organisms in the context of μ MRI, followed by the development of an *ex situ* platform for validating a deep learning–based motion tracking framework. Analysis of the motion tracking system’s performance and limitations, together with the spatial constraints and system complexity of the proposed platform, revealed that fully unconstrained motion in MRI remains challenging. These considerations motivated a transition toward a more controlled experimental paradigm of an *in situ* treadmill, culminating in the selection of a suitable insect model and a detailed examination of the associated μ MRI challenges.

Together, these elements define the scope of the thesis and the design space of the proposed platform, highlighting the interplay between MR measurement constraints, motion tracking

requirements, and biological considerations. The following chapters build on this foundation and develop the core methodological and experimental contributions of the thesis.

4 | *In situ* treadmill

The spherical treadmill provides a minimally constrained behavioral platform that enables position-consistent magnetic resonance measurements of an active insect. This chapter focuses on the design and implementation of the platform within the narrow bore of an ultra-high-field magnet. While the platform effectively addresses the spatial constraints of the magnet bore and enables MRI of tethered insect locomotion, conventional high-resolution μ MRI acquisitions still require scan times that exceed the physiological and behavioral timescales of the insect's motion. As a result, active motion compensation is necessary. To address this, an integrated computer vision system was developed to perform real-time motion estimation and tracking of the tethered insect. The resulting motion information is used in a prospective gating framework to synchronize data acquisition with the insect's movement, enabling spatially consistent k -space acquisition. This approach substantially reduces image artifacts arising from gross abdominal pitching motion during tethered locomotion.

The contents of this chapter are based on the author's contributions to the article "Motion compensated magnetic resonance imaging of an active sun beetle using an *in situ* treadmill," published in *Scientific Reports* (2025) [17].

4.1 Introduction

The geometrical design of the *in situ* treadmill is fundamentally constrained by the physical dimensions of the MRI scanner, which in turn limits the range of animal models that can be accommodated for behavior or locomotion-based imaging experiments. In practice, the primary limiting factor is the magnet bore diameter, with additional constraints imposed by the geometry of the gradient and RF coils. While, in principle, gradient and RF coil designs offer some flexibility through custom fabrication, the bore diameter constitutes a fixed upper bound on the allowable experimental volume. In contrast to other modalities such as optical imaging, electrophysiology, or X-ray-based treadmill experiments, where experimental instrumentation can often be developed or arranged around the treadmill, accessibility in μ MRI is limited to the openings of the scanner bore. More flexibility can be achieved by positioning some instrumentation *ex situ* and using relay mechanisms, like in the case of a fiber-optic-based excitation systems for optogenetic stimulation and readout

in the context of opto-fMRI [98]. Beyond spatial constraints, a fundamental consideration is the MR-compatibility of any objects placed inside the bore. A wide variety of materials and approaches have been developed for constructing MR-compatible systems, which have been applied across numerous preclinical and clinical MRI platforms and applications [99, 100]. In this work, PLA-based additive manufacturing was employed to rapidly prototype and build the spherical treadmill and its associated setup, providing a lightweight, customizable, and an MR-safe platform suitable for integration within the commercial RF volume coil in the ultra-high-field magnet.

The primary purpose of the treadmill system is to confine the insect's movements within the RF coil and MRI field of view, rather than to eliminate motion entirely. As discussed in Section 2.3, motion can lead to image artifacts, signal loss, and spatial misregistration, ultimately limiting achievable spatial resolution and quantitative accuracy. Accurate estimation and characterization of subject motion is a fundamental requirement for effective motion correction in MRI. As a result, substantial effort has been devoted to the development of motion estimation and correction techniques in both clinical and preclinical MRI, some of which have been discussed in Section 2.3.2. Motion tracking using external sensors, particularly optical imaging, offers a practical means of capturing tethered locomotor dynamics and thereby facilitates the mitigation of subsequent motion-related challenges.

Providing optical access to the treadmill-based experimental setup enables continuous, real-time observation of subject motion and behavior during MRI acquisition. Optical imaging system was developed and specifically optimized for operation within the ultra-high-field magnet environment, together with a dedicated computer vision algorithm for motion estimation. Conventional optical tracking approaches typically rely on rigid markers attached to the subject's body [36, 42, 101]. While effective in many human and animal imaging applications, marker-based tracking is invasive and impractical for small, freely moving organisms such as insects. To overcome these limitations, a markerless motion estimation strategy based on sparse optical flow was adopted. This approach enables robust tracking of user-defined anatomical features of the insect body across successive video frames, thereby facilitating accurate estimation of motion fields without requiring any physical modification or further constraints of the subject.

Comprehensive characterization of the developed optical tracking framework was conducted, encompassing detection performance, tracking accuracy, and end-to-end system latency. Using the extracted two-dimensional motion information, a prospective motion-gating strategy was implemented to enable selective acquisition of k -space data during periods of consistent or minimal motion. By synchronizing MR data acquisition with motion states, this approach effectively reduces motion-induced artifacts and improves the reconstructed image quality. Such gating strategies are well established in both clinical and preclinical MRI, where they are commonly used to synchronize image acquisition with physiological processes such as respiration or cardiac cycles [102, 103, 104]. Extending these principles

to behavior-driven motion gating provides a flexible and non-invasive pathway toward high-quality MRI snapshots of tethered insect locomotion.

The following sections provide a detailed account of the concepts, methods, and design choices that define the experimental platform outlined above, which serves as the foundation for the subsequent work in this thesis.

4.2 *In situ* treadmill: design and development

The overall design criteria of the spherical treadmill were governed by three primary constraints: mechanical integration within the available commercial radiofrequency (RF) coil geometry, maintenance of unobstructed optical access for motion tracking, and MR compatibility. The experimental platform for the rigid tethering of the active insect was realized as a precision 3D-printed assembly, fabricated via a fused deposition modeling (FDM) process using a Prusa MK4 system and Prusament polylactic acid (PLA) filament (Prusa Research a.s., Prague, Czech Republic), as shown in Fig. 4.2(a). The overall assembly dimensions (64 mm × 32 mm × 34 mm) were selected to match the form factor of a commercial 35 mm inner-diameter quadrature ^1H RF volume coil in a birdcage configuration, tuned and matched to 650 MHz for the 15.2 T ultra-high-field magnet (Bruker BioSpec; Bruker BioSpin GmbH, Ettlingen, Germany). This hardware configuration served as the standardized imaging environment for all experiments described in this thesis unless stated otherwise.

A 3D-printed hollow ball with a wall thickness of 1 mm and an outer diameter of 17 mm serves as the insect walking platform (see Fig. 4.2(a–c)). The primary structure features a hemispherical slot of 18 mm diameter, with 6 air vents of 2 mm diameter, connected to a laboratory dry air supply delivered through PVC tubing of 4 mm inner diameter, regulated by a pneumatic shut-off valve, as represented in Fig. 4.5 and Fig. 4.1(a–b). The ball is suspended within this slot. The insect is tethered to a vertically adjustable support structure using a UV-cured adhesive (Delo Industrial Adhesives GmbH & Co. KG, Windach, Germany) and positioned above the ball as shown in Fig. 4.1(b) and Fig. 4.2(e). The air-cushioned ball provides a near-frictionless platform, allowing the insect to walk freely in any direction relative to its fixed body.

4.2.1 Treadmill design iterations

The primary geometrical constraint for the treadmill design was the 35 mm inner diameter RF volume coil, while a second key consideration regarding the treadmill operation itself was optimizing the stability of the ball during the pneumatic actuation. Initial prototyping revealed that configurations with multiple air vents provided greater mechanical stability than a single-vent design, as illustrated in Fig. 4.3(a). Several multi-vent variants were evaluated empirically. While four-vent designs were fabricated and tested, a six-vent configuration consistently exhibited superior stability during operation. The final treadmill




Figure 4.1: Design and development of treadmill setup. (a) Integrated treadmill with optical imaging system inside 15.2 T magnet. (b) Spherical treadmill setup inside the ^1H RF coil. (Modified from [17].)




Figure 4.2: Design and development of treadmill setup. (a–d) CAD rendering of the treadmill assembly. Treadmill with active beetle (e) *in situ* and (f) *ex situ* ^1H RF volume coil. (Adapted from [17].)

design incorporating six air vents is shown in Fig. 4.3(b), with the corresponding wireframe model presented in Fig. 4.3(c). The airflow characteristics were not explicitly modeled; instead, the final configuration was determined through iterative rapid prototyping and experimental evaluation.




Figure 4.3: CAD renderings of treadmill design iterations. (a) Initial design with a single air vent. (b) Revised design with six air vents for improved airflow stability. (c) Wireframe model of the six-air-vent design highlighting internal geometry.

The treadmill assembly was designed in a modular configuration to facilitate placement within the RF coil and to enable precise positioning of the insect at the center of the coil. The PLA-based 3D-printed structure is fully compatible with high magnetic fields and produced no detectable susceptibility artifacts during MRI acquisition. Collectively, this design provides the insect with an effectively infinite walking plane while maintaining a fixed in-plane orientation relative to the imaging FOV. In addition, the tethering configuration largely preserves both B_0 and B_1 field homogeneity during imaging by design, as opposed to a potential freely moving animal paradigm discussed in Chapter 3. The present assembly was configured to satisfy the geometric constraints of the commercial cylindrical RF volume coil, thereby minimizing design complexity for the proof-of-concept stage, as described above. A reverse design approach, in which the RF coil geometry is optimized around the treadmill platform, provides a clear direction for future gains in sensitivity. In particular, preclinical small animal MRI commonly employs separate transmit–receive coil configurations, and integrating bespoke surface coils for localized imaging could improve the filling factor and SNR relative to the current volume resonator setup.

4.3 *In situ* MR-compatible optical imaging system

The real-time motion and behavioral responses of the insect inside the ultra-high-field magnet were captured using an integrated optical imaging setup. The design was primarily constrained by the magnet bore geometry, and by the requirement for MR-compatibility. The treadmill positioned within the RF volume coil provided direct horizontal optical access, eliminating the need for complex optical relays. The optical imaging system was developed from a commercial webcam (Microsoft LifeCam HD-3000, Microsoft Corporation, Redmond, WA, USA) with a maximum frame rate of 30 Hz, and image resolution of 1280×720 pixels, as shown in Fig. 4.4(c). The webcam was modified for MR-compatibility by replacing its metallic components with non-magnetic materials. All original metallic screws connecting

the camera body to its printed circuit board (PCB) were removed and replaced with a hot glue bond to ensure structural stability. The integrated microphone was also removed. The modified camera was then housed within a custom 3D-printed mount, as depicted in Fig. 4.4(a–d), and secured with polymer screws. The MR-compatibility of the setup was studied using a water phantom sample and is provided in Section 4.3.2.




Figure 4.4: Design and development of the *in situ* MR-compatible optical imaging system. (a) Perspective view. (b) Side view. (c) Webcam and housing. (d) Exploded view. (Adapted from [17].)

For illumination, a single surface-mount white LED (2 mm, 5 V rigid strip type) was glued to the camera mount (see Fig. 4.4(c)), and was driven using the camera’s power supply. The camera and LED were powered from the USB connection to the computer (Lenovo Thinkpad E14, Lenovo Group Limited, Beijing, China), which was used for grabbing the live video feed and image processing for the gating trigger signal. A 20 m USB active repeater cable connected the *in situ* camera to the computer in the instrument console room (see Fig. 4.5). The camera’s manual focus ring was used for precise focusing, and the system was aligned horizontally with the RF coil at a distance of 45 mm from the coil center inside the magnet bore. For improved accessibility, the imaging setup was extended using a 1 m carbon fiber tube, as shown in Fig. 4.1(a). The primary reason for choosing a low-cost commercial CMOS webcam with modest spatial and temporal resolution was to minimize system complexity and accelerate prototyping of MR-compatible modifications. This configuration met the requirements of the experiments described in this thesis and provided a proof-of-concept for optical integration within an ultra-high-field magnet.

The overall schematic of the *in situ* treadmill with the integrated optical imaging system is shown in Fig. 4.5. The experimental components and operations are distributed across physically separate rooms. The fully electromagnetic interference (EMI)-shielded magnet room contains the MRI hardware and the treadmill setup, the adjacent instrumentation room houses the spectrometer console, laboratory air supply valve, and computer vision system, while MRI acquisition and monitoring are performed from a dedicated control room.

All systems can be controlled remotely, either from the control room or via the computer vision workstation in the control room using remote desktop access.




Figure 4.5: Schematic diagram of the *in situ* treadmill with an integrated computer vision system, a platform for performing MRI of behaving insects. (Adopted from [17]).

4.3.1 RF interference on the *in situ* optical imaging system

Although the camera was modified to be MR-compatible, there remained a possibility that radiofrequency pulses and gradient switching could interfere with its operation. However, empirical observations showed no significant RF interference or data corruption from the camera during MRI acquisition or prolonged use inside the magnet. The camera employs digital USB transmission with differential signalling and shielding, which reduces susceptibility to potential electromagnetic interference. The RF excitation pulse had an estimated amplitude of $B_1(t) \approx 3.74 \mu\text{T}$, based on a sinc-modulated waveform with a duration of 1.575 ms, a bandwidth of 2 kHz, and a flip angle of 90° . These observations suggest that the combination of camera placement outside the coil (45 mm from centre of the coil), low RF power, digital transmission, and shielding was sufficient to prevent interference. More aggressive RF or gradient regimes, or the use of unshielded cabling, may require additional shielding or filtering strategies, as demonstrated in previous MR-compatible device designs [87].

4.3.2 Influence of the *in situ* optical imaging system on MRI

Integrating the optical imaging system within the ultra-high-field magnet may introduce distortions to the B_0 field, potentially affecting the accuracy of MR measurements. This

effect was characterized using B_0 field mapping experiments with a uniform water phantom as the imaging sample. The camera system was positioned 45 mm from the center of the RF volume coil, as in all tethered insect experiments described in this thesis (Fig. 4.6). A test tube with a diameter of 17 mm and a length of 120 mm, uniformly filled with distilled water, was placed inside the 35 mm ^1H volume coil. B_0 field maps were acquired both with and without the optical imaging system inside the magnet. An additional B_0 map was obtained with the optical system present and active shimming enabled to evaluate the effectiveness of correcting the static B_0 field distortions.




Figure 4.6: ΔB_0 distribution with and without the camera system, and with camera + active shim for a water sample. **(a)** Central axial slice. **(b)** Coronal slice [z-axis]. (Adapted from [17].)

The standard *FieldMap* sequence available in Paravision 360 v3.3 (Bruker BioSpin GmbH, Ettlingen, Germany) was used for all experiments. Acquisitions were performed with an isotropic FOV of $(40\text{ mm})^3$ and a matrix size of $128 \times 128 \times 64$, resulting in an in-plane resolution of $312.5\text{ }\mu\text{m} \times 312.5\text{ }\mu\text{m}$. The FOV was centered on the sagittal z-axis of the 76 mm long volume coil to capture the region of best B_0 homogeneity. Mapping parameters included an TE of 1.529 ms and a TR of 35 ms, with no signal averaging (NEX = 1), and the total scan time was 276 s.

The B_0 field distribution of the central axial slice for the three aforementioned cases is shown in Fig. 4.6(a). The mean and standard deviation of the B_0 field variation for each slice, $\Delta B_{0,z}$,

were calculated and plotted to characterize field inhomogeneity along the z -direction, as shown in Fig. 4.6(b). The presence of the active camera introduced an approximately 450 Hz lateral shift in the B_0 field distribution, as evident from the measured B_0 map in Fig. 4.6(b). This shift appeared to be static and was effectively corrected using active shimming with the *Study Shim* routine in Paravision 360 v3.3, as shown in Fig. 4.6(a) and (b). These results demonstrate that the optical imaging system, when appropriately positioned and combined with standard shimming procedures, does not significantly compromise the homogeneity of the static magnetic field in the region of interest.

4.4 Animal handling

Following the development and characterization of the spherical treadmill platform and the MR-compatible optical imaging system in preceding sections, the transition to *in vivo* experimentation required a robust tethering interface and protocol that secured the beetle within the coil and magnet isocenter. Live adult sun beetle (*Pachnoda marginata*) specimens were used throughout the study. The specimens were acquired from an online pet store (thepetfactory.de), housed in a custom-made enclosure at room temperature, and provided with fruit and water as their diet, as specified by the vendor.




Figure 4.7: Animal handling and tethering procedure. **(a)** The beetle is placed in a small ventilated plastic container and immobilized via cold-anesthesia. **(b)** Spherical treadmill setup used for tethered locomotion. **(c)** The removable PLA-based tether is prepared, and a single drop of UV-curable adhesive is applied at the designated attachment site. **(d)** The cold-anesthetized beetle is tethered following adhesive curing. **(e)** The tethered beetle is transferred onto the spherical treadmill. **(f)** The treadmill with the tethered beetle is slid into the RF volume coil. **(g)** Final positioning of the setup for MRI experiments.

For tethering the insect, a UV-curable adhesive (Delo Industrial Adhesives GmbH & Co. KG, Windach, Germany) was used. UV-curable resins are frequently used for insect tethering, as they provide on-demand, fast curing and strong mechanical support while allowing

precise positioning prior to polymerization. To facilitate accurate alignment and minimize physiological stress, beetles were immobilized via cold-anesthesia at 2 °C for 15 min. A single drop of adhesive was applied to the 3D-printed PLA-based tether, which was then positioned on top of the dorsal region of the beetle's body, approximately near the scutellum. The adhesive was cured under a UV source for 10 s; subsequent visual inspection confirmed that this brief exposure was non-damaging and did not impede the transition from dormancy to active locomotory states. Following curing, the tethered beetle was gently transferred onto the spherical treadmill setup, and carefully positioned inside the RF volume coil for the MRI experiments. Upon completion of the experiments, the adhesive tether was easily removed using a cotton swab soaked in warm water. The adhesive did not leave any residue on the beetle's body. The live beetle was subsequently checked for any signs of damage and returned to the enclosure. A detailed illustration of animal handling is provided in Fig. 4.7.

To account for anatomical variability among individuals of the same species, a single-subject design was adopted for each experimental cohort. Unless otherwise specified, all comparative data presented within a specific study or results section were acquired from the same individual beetle within a single continuous session. This approach minimizes inter-subject variability and ensures that observed changes in signal intensity or contrast are attributable to experimental variables rather than morphological inconsistencies. Following the conclusion of the experiments, the subjects were euthanized.

4.5 Baseline vs. motion-affected MRI

Initial imaging experiments were performed to establish a high-resolution anatomical baseline. By utilizing the treadmill assembly without pneumatic activation, the tethered insect remained in a dormant, stationary state, allowing for extended acquisition times and high SNR. These static references, illustrated in Figures 3.10(a) and (b), serve as a ground truth for identifying key internal anatomical structures, including the brain region, digestive system components, and longitudinal thoracic muscle groups, prior to the introduction of tethered locomotion. Unless otherwise stated, all MR imaging presented in this thesis was performed using two-dimensional (2D) acquisitions.

For imaging the thorax region of the insect (Fig. 3.10(b)), a RARE sequence (Paravision 360 v3.3, Bruker BioSpin GmbH, Ettlingen, Germany) was used with TE of 5.78 ms, TR of 3000 ms, slice thickness of 1 mm, and an image matrix of 1024×1024 pixels over a $20 \text{ mm} \times 20 \text{ mm}$ FOV, resulting in an in-plane resolution of $20 \mu\text{m} \times 20 \mu\text{m}$. Five signal averages were acquired with a RARE factor of 3, yielding a total acquisition time of 1 h 25 min 15 s per image.

For the sagittal section (Fig. 3.10(a)), the RARE sequence was applied with TE = 5.24 ms, TR = 1000 ms, slice thickness = 1 mm, and an image matrix of 256×256 pixels over a $30 \text{ mm} \times 30 \text{ mm}$ FOV, corresponding to an in-plane resolution of $117 \mu\text{m} \times 117 \mu\text{m}$. No signal averaging, with a RARE factor of 1, resulting in a total acquisition time of 256 s per image.

The absence of pneumatic activation does not always guarantee complete dormancy. Pre-scan immobilization via cold anesthesia at 2 °C, or a fully restraint configuration can also be used to minimize movement for the baseline imaging. In general, extended periods of inactivity occur due to the diurnal behavior of sun beetles. However, internal organ motion persists even in the absence of gross body movement, particularly affecting the sagittal abdominal slice more than the thorax. Consequently, relatively shorter acquisition times were chosen for abdominal imaging, while acquisition parameters for the thorax were optimized to maximize SNR and improve the visibility of fine muscle structures. These baseline protocols provided reference images that were used to evaluate subsequent behavioral imaging strategies.




Figure 4.8: Effect of locomotion-induced motion artifacts on MRI reconstruction during tethered experiments. **(a)** Reference image acquired while the insect was alive but in a quiescent state, resulting in minimal motion. **(b)** Motion-corrupted image showing artifacts predominantly arising from abdominal and head movements.

During the tethered locomotion, the dominant source of motion arises from abdominal pitching during walking. In comparison, magnitude of internal organ motion is relatively small and can be neglected as of now. Gross abdominal motion during MRI acquisition produces a combination of blurring and ghosting artifacts along the phase-encoding direction (vertical in the image), with severity depending on the extent of motion. This effect is illustrated in Fig. 4.8(b), which shows a motion-corrupted image that can be directly compared with the corresponding reference image acquired under minimal-motion conditions while the insect was in a quiescent state (Fig. 4.8(a)). The primary sources of these artifacts are abdominal and head movements. The following section addresses these motion characteristics within

the framework of a multi-rigid-body model of tethered insect locomotion and subsequently develops the corresponding motion compensation strategy.

4.6 Sun beetle tethered motion model

The spatial and temporal characteristics of insect motion relative to the k -space sampling dictates the severity of the image artifacts, and play a central role in devising effective motion-compensation strategies. It has been demonstrated that abdominal pitching during tethered locomotion is the dominant source of image artifacts. The overall motion of a walking insect on a treadmill consists of several components, including rapid leg movements, gross body displacement driven by semi-periodic abdominal pitching, irregular head motion, and internal organ dynamics. These motions differ substantially in amplitude, frequency, and spatial extent, and therefore contribute differently to MRI phase errors or the k -space corruption.

Although insect motion is inherently three-dimensional, certain components can be approximated as predominantly confined to the two-dimensional imaging plane when a sufficiently thick slice is used. In particular, abdominal pitching during tethered locomotion is largely restricted to the imaging plane. As a result, it produces coherent in-plane phase modulation and represents the primary contributor to observable ghosting artifacts in the reconstructed images. Three representative phases of this semi-periodic abdominal motion are illustrated in Fig. 4.10(b).

This section establishes a kinematic framework for characterizing gross body motion and its impact on the k -space signal, thereby guiding the motion-compensation strategies developed in this thesis. The analysis is restricted to inter-scan motion, where substantial displacement occurs primarily between successive phase-encoding steps during MRI acquisition. Generally, the echo time and signal sampling occur on the millisecond timescale, motion within these intervals is assumed to be negligible and is therefore not explicitly modeled in this study.

4.6.1 Multi-rigid approximation

Although the insect body undergoes continuous non-rigid deformation, most observable displacements are dominated by a few major segments, particularly the abdomen and head (see Fig. 4.8(b)). These segment dynamics can be effectively approximated by piecewise bulk translation and rotation, as discussed in Section 2.3.1. The dominant contributions to motion-related artifacts come from bulk translations and rotations of the major segments, whereas smaller-scale deformations have comparatively little effect on the reconstructed images.

In the most general case, motion is spatially varying and time-dependent, which can be described by a displacement field:

$$\mathbf{r}' = \mathbf{r} + \mathbf{u}(\mathbf{r}, t), \quad (4.1)$$

where $\mathbf{r} = (x, y, z)$ denotes the original spatial position and $\mathbf{u}(\mathbf{r}, t)$ is the position- and time-dependent displacement vector. The corresponding k -space signal, sampled along a trajectory $\mathbf{k}(t)$, is given by:

$$s(\mathbf{k}(t)) = \int_{\Omega} \rho(\mathbf{r}) e^{-i2\pi\mathbf{k}(t) \cdot (\mathbf{r} + \mathbf{u}(\mathbf{r}, t))} d\mathbf{r}, \quad (4.2)$$

with $\rho(\mathbf{r})$ representing the spin density. Because $\mathbf{u}(\mathbf{r}, t)$ varies across the object domain, the phase modulation becomes spatially non-uniform, leading to incoherent phase errors. These manifest as complex image blurring and non-linear ghosting that cannot be corrected by simple rigid-body translations.




Figure 4.9: Schematic representation of the multi-rigid approximation of tethered insect body motion, modeled as three articulated links connected by two joints (defined in the scanner coordinate system).

A useful simplification is to approximate the specimen (object) as a set of N independently moving rigid segments. While global translation and rotation were considered in Section 2.3.1, this approach extends the concept to segment-wise rigid motion. The object domain Ω is partitioned into disjoint regions Ω_n , each corresponding to a rigid segment with its own motion parameters; for example, head, thorax, and abdomen. Under this approximation, the measured k -space signal can be expressed as:

$$S(\mathbf{k}(t)) = \sum_{n=1}^N \int_{\Omega_n} \rho(\mathbf{r}) e^{-i2\pi\mathbf{k}(t) \cdot (\mathbf{R}_n(t) \mathbf{r} + \mathbf{T}_n(t))} d\mathbf{r}, \quad (4.3)$$

where $\mathbf{R}_n(t)$ denotes the rotation matrix and $\mathbf{T}_n(t)$ the translation vector of the n -th segment. This transition effectively replaces the infinite-dimensional displacement field $\mathbf{u}(\mathbf{r}, t)$ with

a discrete set of degrees of freedom. For a specimen composed of N rigid segments, the motion of any point $\mathbf{r}$ within a specific segment Ω_n is constrained to a shared affine mapping $\mathbf{r}' = \mathbf{R}_n(t)\mathbf{r} + \mathbf{T}_n(t)$. Here, $\mathbf{R}_n(t)$ denotes an orthogonal 3×3 rotation matrix that describes the angular orientation of the segment (e.g., abdominal pitching), while $\mathbf{T}_n(t)$ is a 3×1 translation vector representing the linear displacement of the segment's center of mass.

This distinction is physically significant in MRI: according to the Fourier Shift Theorem, the translation $\mathbf{T}_n(t)$ introduces a linear phase modulation in k -space, whereas the rotation $\mathbf{R}_n(t)$ necessitates a corresponding rotation of the k -space coordinates themselves, as discussed in Section 2.3.1. By decomposing the complex biological deformation into these segmental components, the motion correction problem is reduced from estimating a continuous vector field to tracking $6N$ parameters over time, allowing for localized phase and rotation compensation. In the simplified case of a motion in the 2D image slice as addressed in this work, these parameters further reduce to $3N$ (two translations and one in-plane rotation), significantly decreasing the computational complexity of modeling and developing the correction strategies.

4.6.1.1 Articulated kinematics: links and joints

The multi-rigid approximation can be physically interpreted as an articulated kinematic chain, similar to those used in robotics [105]. In this framework, the insect is modeled as a set of rigid *links* connected by *joints* (pivots), as illustrated in Fig. 4.9. For a tethered specimen, the thorax serves as the stationary base link. The head and abdomen are modeled as independent links; while their motion is biologically constrained to rotations at the joints, they are mathematically treated with six degrees of freedom (6-DOF) each to account for the resulting translational shifts of their respective centers of mass relative to the MRI isocenter.

The relationship between any two segments is defined by the coordinates of the joint center $\mathbf{r}_{joint}$. For the n -th segment, the transformation of a point $\mathbf{r}$ into the moving frame $\mathbf{r}'$ is constrained by its pivot point:

$$\mathbf{r}' = \mathbf{R}_n(t)(\mathbf{r} - \mathbf{r}_{pivot}) + \mathbf{r}_{pivot} + \mathbf{T}_n(t) \quad (4.4)$$

This formulation ensures that the rotation $\mathbf{R}_n(t)$ occurs around the anatomical pivot (e.g., the petiole or the neck of the insect), rather than the global origin of the MRI coordinate system. By treating the insect as a series of connected links, we can categorize the motion artifacts as specific “articulation errors”, where phase inconsistencies are concentrated in the moving link (the abdomen) while the base link (the thorax) remains phase-stable, which is evident when comparing the reconstructed images in Fig. 4.8(a) and Fig. 4.8(b).

4.6.1.2 Localized motion and k -space implications

When motion is confined to a small region of the specimen, such as the abdomen of the tethered insect, while the rest of the body remains stationary relative to the tether, the

k -space signal exhibits characteristic features. To analyze these localized phase artifacts, the translational component of motion is considered, denoted by the time-varying displacement vector $\Delta \mathbf{r}(t)$. This vector is equivalent to the translation term $T_n(t)$ introduced in the multi-rigid model (see Section 4.6.1), isolated here to highlight the effects of the Fourier Shift Theorem. Let $w_n(\mathbf{r})$ denote a spatial window representing the moving patch. If this patch undergoes a simple translation $\Delta \mathbf{r}(t)$, the measured k -space signal can be written as:

$$S(\mathbf{k}(t)) = \underbrace{\int w_n(\mathbf{r}) \rho(\mathbf{r}) e^{-i2\pi \mathbf{k}(t) \cdot (\mathbf{r} + \Delta \mathbf{r}(t))} d\mathbf{r}}_{\text{moving patch}} + \underbrace{\int (1 - w_n(\mathbf{r})) \rho(\mathbf{r}) e^{-i2\pi \mathbf{k}(t) \cdot \mathbf{r}} d\mathbf{r}}_{\text{static background}}. \quad (4.5)$$

In image space, multiplication by $w_n(\mathbf{r})$ corresponds to a convolution in k -space. The motion-induced phase modulation therefore spreads across k -space. However, since the convolution is weighted by the spatial support of w_n , the resulting ghosting artifacts appear globally but are strongest in regions corresponding to the moving segment in the reconstructed image. This is evident in the motion affected image in Fig. 4.8(b). This establishes a direct relationship between segmental motion and the resulting localized artifact patterns in the reconstructed images, thereby informing the motion-compensation strategies developed in this thesis.

4.7 Prospective gating

The true motion model within the selected sagittal 2D slice is inherently non-rigid. However, it can be reasonably approximated as a multi-rigid motion comprising three interconnected links, as discussed in the previous section. These can be referred to as the abdominal link, the thoracic link, and the head link. Under the current imaging configuration, the dominant motion artifacts arise from the pitching of the abdominal link (link 1) during locomotion. The thoracic link serves as the tethering point and therefore exhibits relatively minimal motion. In contrast, head motion is comparatively rapid and irregular; however, due to its small voxel volume, the resulting artifacts remain localized and do not significantly affect the surrounding regions.

Motion gating can mitigate artifacts by synchronizing data acquisition to specific phases of motion. The articulated link-and-joint model provides the necessary framework for a prospective gating strategy for the motion compensated μ MRI of the tethered insect locomotion. While the abdominal pitching motion is not strictly periodic, it recurs frequently and consistently within a reasonable time frame, enabling the effective use of prospective gating despite deviations from ideal periodicity. Implementing motion gating typically requires an external or internal surrogate signal that reliably reflects the underlying physiological motion. In the clinical and preclinical MRI, commonly used surrogates include respiratory belts, pressure sensors, navigator echoes, or electrophysiological signals, which provide a

low-dimensional proxy for the motion state. However, such signals often offer limited spatial specificity and may not adequately capture complex motion patterns such as abdominal pitching during locomotion, which represents a fundamentally different regime of motion from the internal physiological motion commonly addressed with gating. In addition, most conventional approaches, aside from MR navigator echoes, require some form of contact with the subject, which can be invasive and may compromise the goal of minimizing the constraints. To address these limitations, motion is estimated non-invasively in a markerless fashion using a computer vision system built upon the integrated optical imaging system, enabling a more spatially resolved and flexible surrogate for prospective gating.

4.7.1 Computer vision-based prospective gating

The integrated optical imaging system described in Section 4.3 enables real-time video acquisition of the active insect on the spherical treadmill. Basic image processing of these data allows the extraction of the insect’s motion trajectory, which forms the foundation for a motion-compensated acquisition strategy. However, reliable prospective motion compensation requires more than offline processing; it demands robust, automated, and low-latency motion estimation.

The Lucas–Kanade sparse optical flow method, as described in Section 2.3.4, can be used for real-time estimation of the insect’s abdominal pitching motion. This effectively transforms the *in situ* MR-compatible optical imaging setup to a comprehensive computer vision system designed for real-time motion estimation. A complete description of the algorithm for the keypoint tracking is given in Algorithm 1. The resulting motion estimate serves as a trigger for prospective gating during single-coil 2D Cartesian k -space line acquisition. Specifically, the computer vision module generates a gating signal based on the tracked motion of a predefined anatomical landmark. The tracking algorithm is implemented in OpenCV [81] using a sparse optical flow formulation. Unless stated otherwise, the posterior apex of the insect (Fig. 4.11(c)) is selected as the tracking keypoint throughout this chapter. While the present implementation focuses on single-point tracking, Algorithm 1 readily extends to multipoint tracking with minor modifications. Furthermore, a comprehensive velocity field could be obtained by adopting a dense optical flow formulation; however, achieving real-time performance with such an approach may pose computational challenges.

The current motion model requires only 2D motion information, corresponding to the MRI slice plane. The relative configuration of the abdominal and thoracic links defines the gating state. The tracking algorithm continuously compares the estimated pose with a predefined reference configuration and generates a trigger signal via a serially connected Arduino Uno microcontroller (Arduino AG, Ivrea, Italy) whenever the state falls within a specified tolerance window. This approach ensures that k -space lines are acquired in a spatially consistent manner, thereby minimizing the effects of gross abdominal motion on MRI reconstruction. The complete control strategy is summarized in Algorithm 2, with a




Figure 4.10: *In situ* computer vision system for real-time prospective gating. **(a)** Schematic of computer vision-based prospective gating. **(b)** Abdominal motion phases of the insect. (Adopted from [17].)

Algorithm 2: Computer vision-based prospective gating [17].

Input: I_k , Image frames captured from real time video

Output: TTL trigger signal to spectrometer console

Initialization: Define MRI parameters;

Capture first optical image frame I_1 ;

Select the reference position $r(x, y)$ and the body part keypoint $p_1(x, y)$ to be tracked;

Start MRI acquisition;

for counter = 1 **to** number of phase encoding **do**

 Optical flow computation based tracking (see Algorithm 1) ;

if tracked position $p_k(x, y) = r(x, y) \pm \epsilon$ **then**

 TTL trigger ON for next k -space line acquisition;

 Wait for one TR and continue;

else

 TTL trigger OFF;

schematic diagram provided in Fig. 4.10. The vision-based triggering pipeline introduces unavoidable latency due to the sequential processes of image acquisition, motion estimation, and trigger generation, which may lead to residual motion artifacts in the reconstructed images. The following section quantitatively characterizes both the motion detection accuracy of the computer vision algorithm and the total end-to-end latency of the triggering pipeline.

4.7.2 Computer vision system performance characterization

The proposed computer vision–based prospective gating strategy does not require cross-calibration procedures, which are commonly required in prospective motion correction (PMC) approaches in clinical MRI [36]. The triggering pulse generation pipeline compares the position of the defined keypoint between consecutive image frames within the same coordinate system, thereby eliminating the need for any coordinate transformation.




Figure 4.11: Euclidean distance error between optical flow algorithm detection and manual annotation for 100 images. **(a)** Y-position trajectory. **(b)** Error histogram. Sample images showing **(c)** small and **(d)** large detection error. (Adapted from [17].)

The accuracy of the optical flow–based algorithm for keypoint detection and tracking was evaluated by comparing its outputs with manually annotated ground truth data. For this purpose, 100 images were extracted from a 256 s-long video of an *in situ* behaving insect and manually annotated. The optical flow–based Algorithm 1 tracked the keypoints across these images. The resulting Y-position trajectories for both the algorithm and manual annotations are shown in Fig. 4.11(a).

Tracking error was quantified as the Euclidean distance between the algorithmically detected and manually annotated keypoint locations (see example images in Fig. 4.11(c) and (d)). Across the 100 images, the mean error was 2.04 ± 1.13 pixels (or $91 \mu\text{m} \pm 50 \mu\text{m}$), with the corresponding histogram shown in Fig. 4.11(b). To account for variability in manual annotation, the Euclidean distance between two independent sets of annotations of the same data was also calculated, resulting in 1.47 ± 0.85 pixels (or $66 \mu\text{m} \pm 38 \mu\text{m}$). These results demonstrate that the optical flow-based algorithm achieves near-human-level accuracy in keypoint detection.

The computer vision-based triggering system introduces inherent latency due to several factors, including the camera's limited frame rate, the extension cable between the camera and the computer, optical flow computation, serial communication with the Arduino, and the transistor-transistor logic (TTL) trigger to the spectrometer for a single MRI k -space scan. The total latency from acquisition of the first video frame to the TTL trigger of the first k -space line was measured to be around 43 ms.

4.8 MRI of the behaving insect

The baseline motion-free and motion-affected MRI experiments described in Section 4.5 highlight the limitations of conventional acquisition in the presence of tethered insect locomotion, and motivate the need for an active motion-compensation strategy. In this section, the developed vision-based prospective gating technique is employed to address this challenge.

Two primary imaging orientations were considered for the 2D MRI of the tethered insect. The first orientation corresponds to an axial section through the thoracic region, which typically exhibits minimal motion during locomotion. The second orientation corresponds to a sagittal section extending longitudinally through the entire body and is therefore highly susceptible to motion artifacts arising from abdominal movement. For this sagittal configuration, the developed computer vision system provides real-time motion estimates that are used to drive prospective gating for motion-compensated μMRI acquisition. Position references of the selected imaging planes are shown in Fig. 3.10(c).

4.8.1 Axial MRI of the thorax

All imaging experiments were performed using a RARE sequence unless stated otherwise. The spin echo-based sequence was chosen to reduce the severe susceptibility artifacts arising from the hard exoskeleton of the insect, as discussed in Section 3.3.2. For imaging the thoracic region, a short TE of 5.24 ms and a TR of 1000 ms were used. Images were acquired with a matrix size of 512×512 pixels over a FOV of $20 \text{ mm} \times 20 \text{ mm}$, resulting in an in-plane spatial resolution of $39 \mu\text{m} \times 39 \mu\text{m}$. The slice thickness was set to 1 mm, with

no signal averaging, a RARE factor of 4, and a total acquisition time of 128 s for a single image.

Figure 4.12(b) shows the thoracic MRI acquired when the insect was inactive, exhibiting minimal motion, as confirmed by the abdominal motion trajectory shown in Fig. 4.12(a). Figure 4.12(d) shows the corresponding thoracic MRI acquired while the insect was actively walking, with the associated motion trajectory shown in Fig. 4.12(c). Both images were acquired without triggering. No visible differences are observed between the two images, indicating that when the dominant motion occurs outside the excited slice, motion-free MRI images of the actively moving insect can be obtained without the need for active motion correction or gating strategies.




Figure 4.12: Effect of tethered locomotion on axial thorax MRI. **(a)** Motion trajectory of the inactive insect abdomen. **(b)** Thorax MRI of the inactive insect. **(c)** Motion trajectory of the walking insect abdomen. **(d)** Thorax MRI of the active insect. (Adapted from [17].)

4.8.2 Sagittal MRI of the insect body

For sagittal MRI acquisition, a short TE of 5 ms and a TR of 1000 ms were used. Images were acquired with a matrix size of 256×256 pixels over a FOV of $30 \text{ mm} \times 30 \text{ mm}$, resulting in an in-plane spatial resolution of $117 \mu\text{m} \times 117 \mu\text{m}$. The slice thickness was set to 1 mm, with no signal averaging and a RARE factor of 1, yielding a total acquisition time of 256 s for a single image. Figure 4.13(f) shows the sagittal MRI acquired without gating during

active locomotion, together with the corresponding abdominal motion trajectory shown in Fig. 4.13(e). The magnified view in Fig. 4.13(g) clearly illustrates the extent of motion artifacts induced by abdominal movement. In contrast, Figs. 4.13(j), (n), and (r) present sagittal MRI images of the same anatomical section acquired using prospective gating under different motion trajectories, shown in Figs. 4.13(i), (m), and (q). The corresponding magnified views in Figs. 4.13(k), (o), and (s) demonstrate a substantial reduction in motion-related artifacts.

As a reference for comparison, Fig. 4.13(b) shows the sagittal MRI of an inactive insect, with the corresponding magnified view in Fig. 4.13(c) and the associated motion trajectory shown in Fig. 4.13(a). As expected, minimal motion artifacts are observed in this case. Pixel intensity profiles extracted from the sagittal MRI images, shown in Figs. 4.13(d), (h), (l), (p), and (t), further quantify the extent of motion artifacts and demonstrate the effectiveness of prospective gating in mitigating these artifacts. A variation in total acquisition time is observed across the prospectively gated MRI datasets when compared with both the ungated active locomotion case and the inactive insect reference. This variation arises from non-uniform waiting periods between successive triggering events, which are strongly influenced by the locomotion behavior of the insect.

Observable differences in SNR between the thoracic MRI (Figs. 4.12(b) and (d)) and the abdominal MRI (Fig. 3.10(a)) arise from a combination of acquisition parameter choices and anatomical factors. Although comparable TR and TE values were used for both scans, the abdominal MRI exhibits higher SNR due to its larger voxel size, larger imaging volume, and the use of a single-echo acquisition with a RARE factor of 1. These factors collectively contribute to an increased signal amplitude.

In contrast, the thoracic MRI prioritizes higher spatial resolution, resulting in smaller voxel sizes and reduced signal per voxel. In addition, a RARE factor of 4 was employed, leading to longer echo trains and increased cumulative T_2 decay, which together reduce the effective SNR. For sagittal abdominal MRI, a RARE factor of 1 was intentionally selected to minimize sensitivity to motion between successive echoes during gated acquisition, since higher RARE factors involve multiple echoes per gating event and are therefore more susceptible to motion-induced phase inconsistencies. In the thoracic MRI experiments, motion gating was not required due to the minimal motion within the excited slice, allowing the use of a RARE factor of 4 as a standard acceleration strategy for RARE-based acquisitions. Although gating can, in principle, be combined with larger RARE factors, such implementations require sequence modifications and precise timing control to mitigate the increased motion sensitivity associated with longer echo trains.

A relatively long repetition time was used across all imaging experiments to ensure near-complete T_1 relaxation, thereby enhancing signal intensity and image contrast, and to accommodate system latency between successive trigger signals during gated acquisition. Mild ghosting artifacts observed in the gated abdominal MRI are attributed to a combination of residual motion and system latency. This effect is particularly evident in the gated scan




Figure 4.13: Prospective gating for motion-compensated sagittal MRI of a behaving insect on a treadmill. **(a–d)** Inactive condition: motion trajectory, sagittal MRI, magnified abdominal view, and corresponding intensity profile showing minimal motion artifacts. **(e–h)** Behaving insect without gating, demonstrating pronounced motion artifacts. **(i–l)**, **(m–p)**, and **(q–t)** Prospective gating iterations 1–3, respectively, showing motion trajectories, MRI images, magnified abdominal views, and intensity profiles, showing consistent reduction of motion artifacts. (Adapted from [17]).

labeled "Gated 3" (Fig. 4.13(r)), which exhibits a pronounced motion profile. The triggering system latency of 43 ms implies that rapid motion occurring within this delay window can contribute to the persistence of residual motion artifacts.

4.9 Discussion

The *in situ* treadmill platform combined with the computer vision-based prospective gating significantly reduced motion-related artifacts, enabling MR imaging of a tethered, behaving insect. The accuracy and latency of the developed computer vision system were sufficient for reliable estimation of the dominant abdominal pitching motion. While adequate for the present application, performance could be further improved through faster imaging hardware and more efficient motion-estimation algorithms. Recent advances in computer vision, particularly in markerless animal tracking, have had a substantial impact across various domains of zoological research, with deep learning-based methods enabling precise behavioral analysis from video data [74, 106]. Although this chapter focused exclusively on motion estimation for MR artifact reduction, the MR-compatible computer vision system could be readily extended to quantitative behavioral analysis, either by enhancing the existing optical flow-based approach or by integrating deep learning-based methods.

This chapter primarily addressed gross abdominal motion, which represents the dominant source of movement during insect locomotion in the tethered configuration. Extension of the proposed approach to other body segments is feasible, provided that the motion can be reliably estimated and parameterized. Currently, leg movements do not produce a measurable impact on the MR signal within the selected imaging slices. Head motion, while relevant in the context of insect locomotion, appeared to be spatially localized and was therefore neglected in this study. The abdominal motion pattern was observed to be semi-periodic, resulting in variable waiting periods for prospective triggering. This is conceptually similar to respiratory gating in clinical MRI, where acquisition efficiency is reduced due to phase-matching constraints. Internal organ motion can also contribute to image blurring, though its effect is generally smaller than that of gross body movement and was not explicitly addressed here. Established MR-based motion estimation techniques, such as navigator echoes [37, 107, 108], could be integrated to monitor internal motion during acquisition, thereby enabling both prospective and retrospective motion correction.

An alternative to motion compensation strategies is to capture the motion by accelerating the MR acquisition beyond the temporal scale of the subject's movement. This generally relies on computational frameworks that reconstruct dynamic 2D or 3D MRI from rapid, continuous, non-gated acquisitions. Such methods typically leverage compressed sensing (CS), low-rank modeling, and stochastic optimization to manage severe undersampling and irregular motion, with demonstrated success in clinical applications like pulmonary and dynamic contrast-enhanced imaging [109]. While μ MRI imposes inherent constraints that limit direct translation of these techniques to insect locomotion, the experimental platform

is compatible with accelerated k -space sampling and advanced reconstruction strategies. This general idea will be explored in later chapters to achieve better temporal resolution, potentially enabling biomechanical studies of insect pose and locomotion, analogous to dynamic and real-time imaging paradigms in clinical and preclinical MRI [110, 111, 112].

The small voxel volumes of insect μ MRI impose fundamental sensitivity constraints, as discussed in Section 3.3.2. The use of ultra-high magnetic field strengths, such as the 15.2 T system employed here, partially alleviates these limitations by enhancing the SNR and spin sensitivity. Nevertheless, the proposed methodology is not intrinsically tied to this field strength and can be adapted to other high-field preclinical MRI platforms. Application to lower-field systems or platforms with limited gradient performance is also feasible, although it necessitates careful consideration of trade-offs among spatial resolution, SNR, and acquisition time. Some of these limitations may be mitigated through extended signal averaging or the use of custom hardware, including optimized RF and gradient coil designs.

Beyond prospective gating, the current framework can be naturally extended to a retrospective gating strategy. In such an approach, k -space would be oversampled to acquire multiple frames of the same slice during periods of motion, followed by post hoc synchronization or binning of k -space lines according to motion phase. This strategy can reduce motion artifacts without introducing irregular waiting periods and potentially enables dynamic MRI of the behaving insect, which is explored in the following chapter.

4.10 Conclusion

This chapter introduced the concept, design, and development of an *in situ* treadmill setup as a platform for μ MRI of actively behaving insects. The MR-compatible, 3D-printed treadmill assembly functions as a tethering platform, enabling stable and reproducible positioning required for MRI acquisition. In addition, the integrated computer vision system allows real-time extraction of motion information of the insect's locomotion. This information is subsequently used for prospective gating to achieve motion-compensated μ MRI acquisition and reconstruction. At present, the temporal resolution of this approach is limited by non-rigid, often irregular abdominal pitching motions that occur during tethered locomotion. Furthermore, prospective gating introduces residual motion artifacts. These limitations are addressed by the methods developed in the following chapter, further advancing the goal towards dynamic μ MRI of behaving insects.

5 | Residual motion artifact correction

The contents of this chapter are based on the author’s contributions to the article “Residual motion artifact removal enables dynamic μ MRI of a behaving *Pachnoda marginata*,” published in *Journal of Magnetic Resonance* (2025) [18].

In situ treadmill provides spatial consistency for MR acquisition with minimal constraints on the movement of the tethered, behaving organism. By confining motion largely within the MRI field of view (FOV) or the selected imaging slice, this configuration enables effective suppression of motion artifacts through the prospectively gated k -space acquisition. The dominant motion during tethered insect locomotion can be approximated as a set of multi-rigid body movements, primarily associated with abdominal pitching. While prospective gating substantially reduces the impact of gross body motion, internal organ motion persists and is not explicitly addressed by the gating strategy. Consequently, residual artifacts will degrade the MRI reconstruction quality, arising from these higher-order motion effects that include non-rigid spatial deformations and temporal variations beyond the assumptions of the gating model.

Additionally, behavioral variability of insect locomotion imposes practical limits on the temporal resolution achievable with prospective gating. Spatial consistency is achieved through position–state matching, which inherently favors motion-compensated imaging of discrete locomotor states (or poses) over continuous dynamic imaging. Retrospective gating offers a complementary alternative, which avoids the need for waiting periods for state-matching. The MRI k -space acquisition is synchronized with the optical imaging and employs a continuous acquisition scheme, with MRI frame repetitions specifically aligned to match the physiological timescale of the insect’s behavior. Motion-compensated MR image frames can be retrospectively recovered by aligning the insect motion states with the corresponding phase encoding steps, analogous to the retrospective cardiac and respiratory gating techniques used in clinical MRI [102, 103, 104].

In addition to limitations imposed by the assumed motion model (Section 4.6), both prospective and retrospective motion correction strategies that rely on external sensors or MR-based internal measurements for motion estimation inherently introduce uncertainty, which can also contribute to residual artifacts in the reconstructed images [113, 114]. Accurately modeling these residual artifacts is particularly challenging, as they arise from a complex

interplay of system latency, tracking errors, and unmodeled non-rigid motion components. This renders the problem highly ill-posed, especially in the context of insect locomotion. Nevertheless, residual motion artifacts are typically less severe than primary motion effects and tend to be spatially localized. Although the underlying dynamics that give rise to these artifacts are challenging to model explicitly due to their non-rigid and irregular nature, their image-domain effects can often be approximated by a simplified translational and rotational motion model. Given these characteristics, data-driven approaches are well-suited for addressing such problems, and have demonstrated considerable success in clinical MRI motion correction [34, 35]. Deep learning (DL)-based methods for MRI motion correction have rapidly advanced in recent years, and several comprehensive reviews summarize these developments [34, 115, 55]. Among these, supervised image-to-image translation approaches employing deep convolutional neural network (CNN) in the image magnitude domain provide a strong and widely adopted baseline, owing to their simplicity, effectiveness, and broad applicability in mitigating motion-induced artifacts [116, 35].

Deep learning-based approaches for MRI motion correction require extensive datasets for training, comprising paired images with and without motion artifacts. Acquiring such datasets from real experiments is often challenging, time-consuming, and costly. Consequently, most motion correction strategies rely on simulated data, frequently derived from large public clinical MRI repositories. In contrast, public MRI datasets for insects or even small animal vertebrate models are largely not readily available, primarily due to the use of highly specialized and species-specific experimental models that are not as standardized or widely adopted as clinical MRI protocols. As a result, obtaining a sufficiently large, motion-free ground truth dataset for insect MRI remains particularly laborious and expensive. Transfer learning can be an effective approach in addressing the data scarcity problem when training deep neural networks [68, 117], as discussed in Section 2.12. The model can be pre-trained with a source domain of motion-corrupted dataset simulated from a large publicly available brain MRI dataset [118], enabling the model to learn generalizable features associated with the motion artifacts. By fine-tuning this pre-trained model on a small target domain dataset of insect MRI with simulated motion artifacts, the network adapts to the specific structural and textural characteristics of insect anatomy while retaining useful feature representations from the source domain.

In summary, this chapter presents a fully retrospective approach for motion-corrected μ MRI of the actively behaving insect. The framework integrates retrospective gating to compensate for gross, multi-rigid abdominal motion with a deep learning-based method to correct residual motion artifacts arising from higher-order dynamics and measurement uncertainties. This combined strategy, as illustrated in Fig. 5.1, extends the potential temporal resolution and image fidelity achievable with μ MRI, forming a step toward the broader goal of dynamic μ MRI of tethered insect locomotion.




Figure 5.1: Schematic of the two-stage MRI motion compensation workflow for the tethered insect on a treadmill, showing **(a)** initial retrospective gating, followed by **(b)** deep learning-based residual motion correction. (Modified from [18].)

5.1 Retrospective gating

The motion model of the tethered insect considered here is essentially the same as that employed for the prospective gating in Section 4.6. The gross non-rigid body motion was approximated using a simplified multi-rigid representation, with a predominantly static horizontal segment corresponding to the thorax, and a dynamically moving abdominal segment. Three discrete representative motion states are illustrated in Fig. 4.10(b) and in Fig. 5.2(b), capturing the dominant abdominal postures observed during tethered locomotion.

The MRI experiment involved continuous acquisition of 2D k -space data from a selected sagittal slice over multiple frame repetitions. The imaging was performed using a RARE sequence, with the acquisition parameters including a TR of 250 ms, a TE of 5 ms, and an image matrix of 256×256 pixels with a FOV of $30 \text{ mm} \times 30 \text{ mm}$. The sequence was configured without signal averaging, with a RARE factor of one, a slice thickness of 1 mm, and 50 repetitions to acquire dynamic MRI frames. This setup resulted in a total acquisition time of 53 min 20 s and an in-plane image resolution of $117 \mu\text{m} \times 117 \mu\text{m}$. The total number of acquired phase encoding or k -space lines was an integer multiple of MRI frame repetitions.

Optical imaging acquisition was synchronized with the MRI acquisition. The 2D motion trajectory of the keypoint was extracted using the optical flow technique [17], as discussed in Section 2.3.4. The abdominal pitching motion was irregular and only semi-periodic, a consistent motion cycle could not be reliably defined; as a result, conventional phase-based gating was unreliable and could not ensure uniform sampling of phase-encoding lines for a given motion state. To address this limitation, the insect abdominal pose was discretized into position states via grid-based clustering of the motion trajectory, with cluster centroids defining representative states. Although the motion was temporally irregular, it was spatially constrained within a finite range (e.g., $\pm 5^\circ$ rotation of the abdomen link), making a discretization into spatial bins meaningful. This approach provides a simple and robust way to define motion states without assuming periodicity or a specific motion model and can be interpreted as a form of self-gating.

Figure 5.3(a) presents the motion trajectory of the annotated keypoint (posterior part of the insect abdomen) during k -space acquisition, with color-coding indicating the spatial clusters identified through grid-based categorization. Fourteen clusters were defined using a 4×4 equally spaced 2D spatial grid, excluding any grid cells that did not contain data points. Figure 5.3(b) displays a 2D X-Y plot of the same trajectory, where the cluster centroids define the position states for the gated MRI reconstruction. The standard deviation of points within each cluster, shown in Figure 5.3(c), served as a metric for the precision of each position state estimation. Phase encoding steps were then correlated with the insect abdominal pose position states during MRI acquisition, to generate individual k -space matrices for motion-compensated image reconstruction. Multiple k -space frames of the same position state were averaged for better SNR, while gaps in k -space were filled using linear interpolation. Figure 5.2 shows an overall schematic of the process, where Figure 5.2(a) represents the




Figure 5.2: Schematic representation of the *in situ* treadmill setup integrated with an MR-compatible computer vision system for motion estimation, enabling retrospective gating for motion-compensated MRI. **(a)** Computer vision system. **(b)** *In situ* treadmill. **(c)** Motion trajectory of the selected keypoint, with the phase encoding index of the MRI frames. **(d)** Phase encode reordering for position consistent MRI reconstruction. (Adopted from [18].)

in situ MR compatible computer vision system, Figure 5.2(b) shows the *in situ* treadmill setup, Figure 5.2(c) illustrates the motion trajectory of the selected keypoint calculated by the optical flow algorithm, and Figure 5.2(d) represents the sorted MRI k -space frames corresponding to the discrete position states, which upon DFT reconstruction gives motion compensated images. The overall pipeline is presented in Algorithm 3.

Algorithm 3: Retrospective gating. (Adopted from [18].)

Input :

- $\mathbf{K}(k_x, k_y, n)$: 3D k -space (frequency encoding $\times$ phase encoding $\times$ frames).
- $\mathbf{S}(n)$: Motion signal - sparse optical flow motion trajectory of keypoint.

Output: $\{\mathbf{I}_m(x, y)\}_{m=1}^M$: Motion-resolved MRI frames.

I. Bin frames by motion state

Cluster motion signal $\mathbf{S}(n)$ into M states $\{m_1, \dots, m_M\}$ using grid-based clustering.

II. Assign frames to motion bins

foreach frame n **do**

└ **Assign** $\mathbf{K}(:, :, n)$ to bin m where $\mathbf{S}(n) \in m$.

III. Reconstruct motion-resolved k -space

foreach motion bin m **do**

└ **Reconstruct** $\mathbf{K}_m(k_x, k_y) = \text{mean}\{\mathbf{K}(k_x, k_y, n) \mid n \in m\}$.

└ **if** $\mathbf{K}_m$ has missing k_y lines **then**

└└ **Interpolate** using:

└└└ - Linear interpolation.

IV. Generate motion-resolved images

foreach $\mathbf{K}_m$ **do**

└ $\mathbf{I}_m(x, y) = \text{FFT}^{-1}(\mathbf{K}_m(k_x, k_y))$.

5.1.1 Retrospective gating for the initial correction

Retrospective gating was implemented following the procedure outlined in Algorithm 3. Figure 5.3(d) presents the average of 50 MRI magnitude images, highlighting the extent of motion-induced artifacts throughout the acquisition. Motion is primarily localized to the abdomen and head, while the thorax exhibits minimal to no artifacts, and is consistent with the assumed motion model. A representative motion-corrupted frame from this dataset is shown in Fig. 5.3(e). Initial motion compensation achieved via gated reconstruction is illustrated in Fig. 5.3(f), while Fig. 5.3(g) provides a reference MRI frame of the same slice in the absence of large-scale abdominal motion. The multi-rigid abdominal position states of the tethered insect on the treadmill are visualized in Fig. 5.2(b). This experiment demonstrates the reconstruction of a single gated frame corresponding to a specific positional state. The method extends to the reconstruction of multiple frames representing the positional states

extracted through clustering, with performance determined by the motion dynamics of the behaving insect.

Both the gated frame (Fig. 5.3(f)) and the reference frame (Fig. 5.3(g)) exhibit residual blurring, primarily due to the physiological motion of the internal organs. The following section characterizes the nature and extent of this residual motion, and the subsequent sections present strategies to mitigate it.




Figure 5.3: Retrospective gating for motion compensated μ MRI. (a) X&Y-motion trajectory of the selected keypoint for the duration of MRI raw data acquisition, color-coding represents the different spatial clusters. (b) 2D X-Y plot of the same motion trajectory, with color-coding represents the grid-based clusters with the centroid defines the different position states for gating. (c) Standard deviation of distance from the centroid of the individual clusters. (d) Average magnitude image of 50 MRI frames showing the extend of motion-artifacts. (e) Sample motion-corrupted frame out of the reconstructed 50 frames without motion compensation. (f) Retrospectively gated MRI frame showing reduction in motion-artifacts. (g) Reference image with less motion artifacts. (Adopted from [18].)

5.2 MRI of the dormant insect: internal organ motion

The relative timescale of motion compared to the acquisition time is a major factor influencing the severity of motion artifacts. So far, the attention has primarily focused on the gross body motion during tethered locomotion. However, internal physiological organ motion represents another important source of artifacts and should be addressed when moving towards the goal of dynamic imaging. These motions span a wide range of spatial and temporal scales, depending on the organ and physiological process of interest, and are ubiquitous irrespective of overt behavior. The major contributors to these internal dynamics are hemolymph circulation throughout the body cavity, peristaltic digestive movements transporting food through the esophagus, crop, midgut, and hindgut, and tracheal compression associated with respiration. Under standard MRI acquisition settings, the cumulative effect of these processes manifests as blurring and ghosting artifacts. MRI experiments can be designed to explicitly probe the sensitivity to internal organ motion. While multiple factors influence the severity of image-domain artifacts, acquisition time serves as a useful aggregate measure, as it is intrinsically linked to several other MR parameters.

Tethered but dormant insect MRI follows the protocol described in the Animal handling Section 4.4. Two-dimensional sagittal-plane (position referene Fig. 5.4(b)) MR images were acquired using a RARE sequence with a TR of 500 ms and a TE of 5 ms. The images were acquired with a matrix size of 256×256 over a FOV of $30 \text{ mm} \times 30 \text{ mm}$, yielding an in-plane spatial resolution of $117 \mu\text{m} \times 117 \mu\text{m}$. A slice thickness of 3 mm was used, with no signal averaging and a RARE factor of four. A total of 20 repetitions were acquired to generate the dynamic MRI frames, resulting in a total acquisition time of 640 s. The motion trajectory of a selected keypoint within the digestive system is shown in Fig. 5.4(a), demonstrating the periodic nature of internal organ motion, while Fig. 5.4(c)–(e) reveal distinct phases of digestive system peristalsis. Although each frame required an acquisition time of 32 s for this particular MRI experiment, the dormant state of the insect featured relatively slow internal organ motion, making the reconstruction less prone to motion artifacts. This is not the case for an active insect, where both gross body motion and increased internal organ dynamics lead to pronounced image blurring, as shown in Fig. 4.13(f) and (g).

Figure 5.5(c) shows the same sagittal slice and was based on identical MRI parameters, except for a TR of 1000 ms and RARE factor of 4, resulting in an effective acquisition time of 64 s per frame. Here, the magnified image clearly shows motion artifacts caused by internal organ movement, which can be compared to Figure 5.5(b), acquired over 32 s for the same specimen and similar physiological state. This can be further validated with the following MRI of the same sagittal section, based on a much faster acquisition. The same sequence with much shorter TR of 75 ms, RARE factor of 10, and a partial Fourier acquisition of 160 k -space lines with homodyne reconstruction (using Paravision 360v3.3) resulted in an effective acquisition time of 1.2 s per frame. This relatively fast acquisition shows the corrugated structure of the insect midgut displayed in Figure 5.5(a), which is associated with peristaltic




Figure 5.4: Internal organ motion of a dormant insect. **(a)** Motion trajectory of the marked keypoint on the insect's digestive system estimated using the optical flow algorithm 1 (Time axis shown at 5× real-time). **(b)** MRI sagittal slice reference position. **(c)** Internal organ motion, phase 1. **(d)** Internal organ motion, phase 2. **(e)** Internal organ motion, phase 3. (Adapted from [18].)




Figure 5.5: Effect of acquisition time on motion sensitivity in single-frame μ MRI of the dormant insect. **(a)** Short acquisition time of 1.2 s captures the corrugated structure of the midgut with minimal motion artifacts. **(b)** Intermediate acquisition time of 32 s shows moderate motion artifacts. **(c)** Long acquisition time of 64 s results in pronounced motion artifacts. (Adapted from [18].)

wave propagation along the tissue. In contrast, for slower acquisitions, the peristaltic motion averages over time, causing the midgut to appear as a smooth, cylindrical structure devoid of visible corrugation. The reduced signal intensity observed in the fast-acquisition images is attributed to the short repetition time, which limits longitudinal magnetization recovery between excitations and consequently lowers the overall signal.

In summary, while gating can compensate for gross body motion during tethered locomotion, residual motion artifacts persist and are often challenging to model explicitly. Internal physiological organ dynamics are a major contributor to these residual artifacts, with their severity depending strongly on the acquisition timescale. The deep learning-based correction described in the following sections addresses these limitations to achieve the final refined image reconstruction.

5.3 Deep learning for residual motion correction

Although the true residual motion is non-rigid, its effects are often spatially localized and can be approximated locally by simple translational and rotational perturbations. Motivated by this observation, rigid-body motion simulations with dense sampling of plausible motion parameters are used for network training. While this does not explicitly model non-rigid dynamics, the convolutional nature of the U-Net [65]-derived network used here enables

it to learn local corrective mappings that can generalize to more complex residual motion patterns.

A transfer learning-based approach was employed to train the model, addressing the data scarcity challenges inherent to insect μ MRI. The model was pre-trained on a larger brain MRI dataset and subsequently fine-tuned end-to-end on the smaller insect μ MRI dataset. Details on the deep learning problem formulation and implementation are provided in the following sections.




Figure 5.6: Transfer learning for residual motion artifact correction. **(a)** Pre-training the model with the brain MRI dataset. **(b)** Fine-tuning the model with a mixed dataset of brain and insect MRI images. (Adapted from [18].)

5.3.1 Datasets

Two ground truth motion-free datasets were employed for simulating motion-corrupted images. The first dataset consisted of a publicly available T1-weighted adult human brain MRI dataset [118], originally acquired using a 12-channel imaging coil. Coil sensitivity maps were normalized to produce a single complex-valued image set with a matrix size of 256×256 .

The second dataset was experimentally acquired and comprised sagittal MRI slices of inactive insects. All insect imaging was performed on the previously described 15.2 T Bruker Biospin MRI system with the 35 mm quadrature ^1H RF coil. Images were acquired with a RARE sequence using the following parameters: TR of 500 ms, TE of 5 ms, matrix size of 256×256 , and FOV of $30 \text{ mm} \times 30 \text{ mm}$, resulting in an in-plane resolution of $117 \mu\text{m} \times 117 \mu\text{m}$, slice

thickness of 1 mm, and no signal averaging. A linear scaling factor of 1.2 was applied to all insect MR magnitude images for visualization unless stated otherwise. Anatomical reference image highlighting major structures are shown in Figure 3.10. Although the same insect species (*Pachnoda marginata*) with similar external morphology was used across experiments, subtle variations in internal anatomy occur between specimens due to physiological differences and minor variations in acquisition parameters. These variations contribute to a more diverse and robust dataset for training the deep learning model.

5.3.2 Motion artifact simulation

Supervised training of deep neural networks for MRI motion correction typically requires a substantial dataset of motion-corrupted and motion-free MRI magnitude images [116, 35, 34]. This presents significant challenges, especially in the context of MRI for small invertebrate organisms. Simulating bulk rigid-body motion artifacts from motion-free k -space and image-space is an effective approach for generating training datasets, as commonly reported in the clinical MRI motion correction literature [34].

In-plane rigid-body motion was assumed in combination with a single-coil, sequential Cartesian k -space acquisition. Motion mostly happened between individual k -space line acquisitions, known as inter-view motion. Motion during the acquisition of the individual k -space lines over the period of the echo time were assumed to be negligible. According to the Fourier shift theorem, bulk translation motion induces a phase shift in the corresponding k -space line and by the Fourier rotation theorem, subject rotation results in corresponding rotation of the k -space line itself, as discussed in Section 2.3.1. Translational motion was simulated directly in k -space using phase modulation, as dictated by the Fourier shift theorem. In contrast, rotational motion was simulated in the image domain, following the methodology outlined in [119]. Implementing rotation directly in discretely sampled Cartesian k -space would require resampling onto a rotated, non-Cartesian grid and subsequent reconstruction using gridding or non-uniform FFT (NUFFT), which is computationally complex and prone to interpolation-induced artifacts. Image-space rotation therefore provides a simpler and sufficiently accurate approximation for the small rotational perturbations considered in this chapter.

The simulation pipeline was implemented in MATLAB (Mathworks Inc.), with the schematic representation provided in Fig. 5.7. Figure 5.7(a) represents the motion-free MRI magnitude image and the corresponding k -space magnitude. Figure 5.7(b) shows the k -space compilation process, where the image was rotated to varying degrees and the corresponding k -space was sampled, and similarly, the motion-free k -space was translated with a varying amount and sampled for a motion corrupted k -space. Standard Cartesian Fourier reconstruction of this compiled k -space gives the motion-corrupted MR image, shown in Fig. 5.7(c). Although the schematic diagram shows the rotation and translation applied simultaneously for easier representation, independent translation and rotations of the images are feasible and effective for model training.




Figure 5.7: Schematic diagram showing motion artifact simulation pipeline. **(a)** Motion-free MRI and k -space magnitude images. **(b)** k -space compilation based on image rotation and k -space translation. **(c)** Reconstructed MRI magnitude image showing motion artifacts. (Adapted from [18].)

The original brain MRI dataset [118] comprised 25 subjects, each with $256 \times 256 \times 170$ slices. To simulate motion artifacts, slices 11 to 160 were selected from each subject, resulting in a total of 3750 image pairs. Out of these, 3,000 image pairs were selected for the training set, and 400 image pairs were selected for the validation set. Translation motion artifacts were simulated by applying pseudo-random motion magnitudes within the normalized range $[-1, 1]$, scaled by three distinct motion severity levels (1, 2, 3), to phase encode lines $[20:120, 145:220]$ of the 256 total phase encoding lines. This resulted in spatial domain translations ranging from 1 to 10 pixels in both the X and Y directions. A separate rotational motion dataset was generated using the first 10 subjects out of the same 25 subjects, yielding 1500 image pairs, which were divided into 1400 training and 100 validation sets. Rotation magnitudes were pseudo-randomly selected between -1 and 1 pixel degrees applied to phase encoding lines $[80:120, 150:220]$ of the 256 lines. Final training set and validation set for the pre-training of the model resulted in 4400 and 500 image pairs respectively, and sample image-pairs are shown in Fig. 5.8. The motion artifact simulation of the small insect dataset followed the same pipeline, resulting in a training dataset of 500 image pairs and a validation dataset of 55 image pairs. Sample insect MRI motion-free and motion-corrupted image pairs are shown in Fig. 5.9. Dataset distribution details are provided in Appendix C.




Figure 5.8: Motion artifact simulation: brain MRI dataset. **(a)** Motion-free reference (Example 1). **(b)** Corresponding image corrupted by simulated translation motion. **(c)** Absolute difference between motion-free and motion-corrupted images, highlighting motion artifacts. **(d)** Translation motion trajectory used for simulation. **(e)** Motion-free reference (Example 2). **(f)** Corresponding image corrupted by simulated rotational motion. **(g)** Absolute difference between motion-free and motion-corrupted images. **(h)** Rotation motion trajectory used for simulation.




Figure 5.9: Motion artifact simulation: insect MRI dataset. **(a)** Motion-free reference (Example 1). **(b)** Corresponding image corrupted by simulated translation motion. **(c)** Absolute difference between motion-free and motion-corrupted images, highlighting motion artifacts. **(d)** Translation motion trajectory used for simulation. **(e)** Motion-free reference (Example 2). **(f)** Corresponding image corrupted by simulated rotational motion. **(g)** Absolute difference between motion-free and motion-corrupted images. **(h)** Rotation motion trajectory used for simulation.

5.3.3 Deep learning model architecture and problem formulation

The general principles of the transfer learning paradigm in deep learning were introduced in Section 2.4.2. Building upon this framework, motion artifact correction is formulated here as a supervised image-to-image regression task.

Let $\mathbf{x} \in \mathcal{X}$ denote an MR magnitude image degraded by motion artifacts, and let $\mathbf{y} \in \mathcal{Y}$ denote the corresponding motion-free reference image. The objective is to learn a parametric mapping $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, with trainable parameters θ , such that the network output $\hat{\mathbf{y}} = f_\theta(\mathbf{x})$ provides an accurate approximation of the ground-truth image $\mathbf{y}$. In this formulation, the model implicitly learns to invert the artifact-generating process and restore structural and intensity information that is degraded by motion.

Pretraining with brain MRI

Network pretraining was conducted using simulated motion-corrupted images derived from a publicly available brain MRI dataset, paired with their corresponding motion-free reference images. A U-Net-based convolutional neural network was trained in a supervised manner by minimizing the mean squared error (MSE) loss between the predicted and reference images,

$$\mathcal{L}_{\text{pretrain}}(\theta) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}_{\text{brain}}} [\|f_\theta(\mathbf{x}) - \mathbf{y}\|_2^2], \quad (5.1)$$

where $\mathcal{D}_{\text{brain}}$ denotes the brain MRI training dataset.

Fine-tuning with brain and insect MRI

Following pretraining, the network was fine-tuned using a combined dataset $\mathcal{D}_{\text{fine}}$ consisting of a subset of brain MRI data and a limited insect MRI dataset. Fine-tuning employed a weighted loss function combining structural similarity (SSIM [120]) and MSE,

$$\mathcal{L}_{\text{fine}}(\theta) = 0.8 \cdot (1 - \text{SSIM}(\mathbf{y}, \hat{\mathbf{y}})) + 0.2 \cdot \|f_\theta(\mathbf{x}) - \mathbf{y}\|_2^2. \quad (5.2)$$

This composite loss encourages preservation of perceptual image structure while maintaining pixel-wise fidelity.

Optimization

Model parameters θ were optimized using the Adam optimizer [121] by minimizing the fine-tuning loss,

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{\text{fine}}(\theta). \quad (5.3)$$

This optimization strategy enables adaptation of the pretrained network to the insect MRI domain while retaining generalizable features learned from the brain MRI data.

The deep learning model employed was based on a modified version of the standard U-Net architecture [65], originally developed for biomedical image segmentation. The network

follows an encoder–decoder structure with skip connections, enabling the simultaneous capture of local spatial details and global contextual information. The encoder consisted of five hierarchical levels with feature channel sizes of $[16, 32, 64, 128, 256]$, operating on a single-channel input image of size 256×256 . Each encoder level comprised a DoubleConv block with two 3×3 convolutional layers, each followed by batch normalization and ReLU activation. Spatial downsampling was performed using 2×2 max-pooling operations, progressively reducing spatial resolution while increasing feature dimensionality.

The bottleneck stage consisted of two convolutional layers with 512 feature channels to extract high-level feature representations. The decoder mirrored the encoder structure, employing 2×2 transposed convolutions with stride 2 for upsampling, followed by DoubleConv blocks for feature refinement. Skip connections from the corresponding encoder levels were incorporated to preserve spatial information and improve reconstruction fidelity. The final output layer used a 1×1 convolution to map the decoded features to a single-channel motion-corrected MR magnitude image of size 256×256 . The schematic of the architecture used is provided in Fig. 5.10.




Figure 5.10: Schematic of the U-Net [65]–based deep learning architecture used for image-domain motion artifact correction.

5.3.4 Model training and evaluation

For model **pre-training**, brain MRI dataset comprising simulated motion-corrupted and motion-free image pairs was used, with a training–validation split of 4400 and 500 images,

respectively. Image preprocessing and data augmentation included intensity normalization, random rotations, and horizontal and vertical flipping to improve model robustness and generalization. Model training was formulated as a pixel-wise image-to-image regression task using the MSE loss function, with a mini-batch size of 12. Optimization was performed using the Adam optimizer [121] with an initial learning rate of 1×10^{-4} . Early stopping was applied when the validation loss did not improve by more than $\Delta = 1 \times 10^{-4}$ for 20 consecutive epochs. Upon convergence, the learning rate was reduced to 1×10^{-5} , and training was completed after 72 epochs. All hyperparameters were selected empirically based on validation performance.

For model **fine-tuning**, a mixed dataset was employed consisting of 500 brain MRI image pairs (a subset of the pre-training dataset) and 500 insect MRI image pairs for training, along with a validation set of 80 image pairs (25 brain and 55 insect MRI). End-to-end fine-tuning was performed while retaining all training parameters from the pre-training stage, with the exception that the number of training epochs was limited to 30 to mitigate overfitting. A combined loss function incorporating MSE and structural similarity index measure (SSIM) [120] was used, with empirical weighting of 0.2 and 0.8, respectively, to balance pixel-wise fidelity and perceptual image quality, particularly for the small insect dataset. Model performance was subsequently evaluated using both simulated motion-corrupted images and real experimentally acquired motion-corrupted images, as presented in the following sections.

Hardware specifications:

All deep learning experiments were conducted on a workstation equipped with an NVIDIA RTX A5000 GPU featuring 24 GB of dedicated memory with CUDA support, along with 256 GB of system RAM. Model development and training were performed using Spyder 5.4.3 integrated development environment, Python 3.9.17, and the PyTorch framework [122], running on the Windows 10 operating system (Microsoft, 2019).

5.3.5 Image quality metrics

In MRI motion correction and related artifact reduction literature, image quality evaluation has traditionally relied on quantitative metrics such as SNR, contrast-to-noise ratio (CNR), MSE, and peak signal-to-noise ratio (PSNR). These measures quantify noise characteristics, contrast preservation, and voxel-wise intensity differences, and have been widely used in conventional reconstruction and artifact reduction studies. However, voxel-wise error metrics are highly sensitive to small spatial misalignments and smooth intensity variations, and therefore often correlate poorly with perceived anatomical fidelity, particularly in the presence of motion-induced structural distortions [123].

In clinical MRI, residual motion, geometric distortions, and bias-field variations frequently introduce sub-voxel shifts and intensity modulations that may have a limited impact on

diagnostic interpretability but can produce large errors in MSE- or RMSE-based evaluations. As a result, these metrics may over-penalize reconstructions that preserve anatomical structures while underrepresenting perceptually relevant degradations such as blurring. All of these are largely true for the case of preclinical small animal MRI and μ MRI as well.

To overcome the limitations of voxel-wise error metrics, structure-aware and perception-based measures, most notably the SSIM [120] and related variants, have become standard in modern MRI reconstruction and motion correction research. These metrics evaluate local luminance, contrast, and structural consistency, and therefore align more closely with radiological assessment than pure pixel-wise differences.

In addition, MRI motion correction scenarios do not always provide a perfectly aligned motion-free reference image, particularly for the case of *in vivo* μ MRI settings. This motivates the inclusion of complementary reference-free image quality measures that can assess structural sharpness and artifact levels directly from the reconstructed image. Accordingly, this thesis employs two structure-aware image quality metrics, SSIM and gradient entropy (GE), to quantitatively assess improvements following residual motion correction. SSIM captures structural similarity relative to a reference, while GE provides a reference-free assessment sensitive to edge preservation and motion-induced blurring. The formulation and implementation details of these metrics are described in the following sections.

5.3.5.1 Structural Similarity Index Measure (SSIM)

The SSIM [120] is a full-reference perceptual metric that evaluates image similarity by comparing luminance, contrast, and structural details between a reconstructed image and a reference. SSIM produces a score between -1 and 1 , where 1 indicates perfect similarity, 0 indicates no similarity, and -1 indicates perfect dissimilarity. It has been shown to correlate well with human visual perception. Mathematically, SSIM is defined as

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)}, \quad (5.4)$$

where x and y denote the reference and reconstructed MRI images, respectively; μ_x and μ_y are their mean intensities; σ_x^2 and σ_y^2 are the corresponding variances; and σ_{xy} is the covariance between x and y . The constants C_1 and C_2 stabilize the metric in regions of low intensity or contrast.

5.3.5.2 Gradient Entropy (GE)

GE is a reference-free metric that quantifies the randomness of image gradients, providing an estimate of edge richness and structural definition. GE is particularly useful for evaluating motion correction in experimental MRI data, where a motion-free reference may not be available. It is defined as,

$$\text{GE} = - \sum_i P(g_i) \log P(g_i), \quad (5.5)$$

where g_i represents the gradient magnitude values, and $P(g_i)$ denotes the normalized probability distribution. The gradient magnitudes are computed as,

$$|\nabla I| = \sqrt{I_x^2 + I_y^2}, \quad (5.6)$$

with I_x and I_y being the image gradients in the horizontal and vertical directions, respectively. These are calculated by convolving the image with Sobel or equivalent edge-detection filters. Lower GE values correspond to more structured edges and higher image quality in the context of MRI motion correction [124, 125]. SSIM and GE were computed using custom Python scripts with OpenCV [81] v4.10.0, scikit-image [126] v0.21.0, and SciPy [127] v1.11.1.

5.3.6 Motion correction on simulated data

The deep learning model, fine-tuned using the small insect dataset, was used to correct simulated motion artifacts in a test dataset comprising 100 insect MRI images. Figure 5.11(a) presents a representative motion-corrupted image, while the corresponding motion corrected output from the model is shown in Fig. 5.11(b). For reference, the ground truth motion-free image is displayed in Fig. 5.11(c). For all cases, the rightmost panel shows the absolute difference map with respect to the motion-free reference image.

The effectiveness of the correction across the dataset is summarized in Fig. 5.12(a), which shows the histogram of SSIM values before and after correction. The mean SSIM improved from 0.668 ± 0.080 (uncorrected) to 0.847 ± 0.055 (corrected), indicating a substantial reduction in motion artifacts. Further validation using the GE metric is presented in Figure 5.12(b), where the mean reduced from 2.985 ± 0.190 before correction to 2.231 ± 0.143 after correction, confirming the enhanced image quality achieved by the model.

5.3.7 Motion correction on experimental data

The performance of the fine-tuned model described in Section 5.3.6 was evaluated on experimental MRI data containing real motion artifacts. Figure 5.13(a) shows a retrospectively gated MR image frame, providing the initial motion compensation as described in Section 5.1.1, with a corresponding GE value of 3.154. Residual motion artifacts remain evident in this reconstruction, which were subsequently addressed using the deep learning-based correction. The corresponding motion-corrected image is shown in Fig. 5.13(b), where the GE decreased to 2.381, indicating a clear reduction in residual distortions.

The severity of motion artifacts was directly related to the activity level of the insect during k -space acquisition. For cases with pronounced multi-rigid motion, involving substantial abdominal movements, initial compensation via retrospectively gated reconstruction was required, as shown in Fig. 5.13(a). Additionally, less severe motion-corrupted images, as shown in Fig. 5.13(c) could also be directly corrected using the same deep learning model, which is shown in Fig. 5.13(d), with the GE reduced from 3.280 to 2.143, demonstrating




Figure 5.11: Deep learning motion correction on simulated dataset. (a) Simulated motion-corrupted test image, magnified view, absolute difference with the reference. (b) Motion-corrected image with the fine-tuned model, magnified view, absolute difference with the reference. (c) Motion-free reference for comparison, magnified view, absolute difference with the reference. (Adapted from [18].)




Figure 5.12: Quantitative evaluation of deep learning motion correction on simulated dataset. **(a)** Histogram of SSIM values for test images before and after correction. **(b)** Histogram of GE values for test images before and after correction. (Adapted from [18].)

effective artifact mitigation even without prior gating. Figure 5.13(e) summarizes the performance across the dataset, presenting a bar graph of gradient entropy values before and after deep learning correction for both gated and non-gated MR reconstructions, highlighting the consistent improvement in image quality achieved by the model.

5.4 Discussion and conclusion

External motion tracking and subsequent gating primarily handle gross multi-rigid-body motion; fine-scale non-rigid dynamics of the insect abdomen and internal organ motion frequently introduce residual artifacts. Deep learning-based correction substantially reduces these residuals, highlighting the benefit of combining classical gating with learned image-domain approaches. While the current pipeline is effective for the presented problem, it is not yet a universal solution for the full spectrum of complex motion dynamics arising in tethered insect locomotion. In principle, an end-to-end deep learning model could capture fully non-rigid dynamics, but its effectiveness depends on accurate modeling of insect locomotion for the generation of supervised training datasets.

It has been shown that the transfer learning-based framework can address the data scarcity challenges of motion-correction in μ MRI, and the applicability can be readily extended to other related problems, such as correcting various imaging artifacts or performing image denoising. While the fine-tuning presented in this chapter focused on a single species (*Pachnoda marginata*), the methodology is inherently adaptable across multiple species, offering a generalizable approach for mitigating motion artifacts in behaving animals.

Retrospective gating offers a distinct advantage over prospective gating by reducing the impact of measurement uncertainties and system latencies on the reconstructed image. Unlike prospective approaches, which require precise triggering for every k -space line,




Figure 5.13: Deep learning (DL) motion correction on experimental dataset. (a) Retro-spectively gated image with residual motion artifacts. (b) DL motion-corrected image. (c) Non-gated image with less severe motion artifacts. (d) DL motion-corrected image. (e) Bar graph of GE values before and after DL correction. (Adapted from [18].)

retrospective gating continuously acquires data and assigns motion states during post-processing. Because multiple frames can be recovered subject to the motion dynamics, the effective temporal resolution is, in principle, superior to that of prospective gating. However, the method has inherent limitations. It requires oversampling of k -space to capture sufficient data across all motion states, which increases total acquisition time. Furthermore, binning k -space lines into discrete motion phases can create gaps that require interpolation, potentially introducing additional artifacts. The resulting temporal resolution is also highly dependent on the specimen's dynamics, often producing snapshots at selected motion states rather than truly continuous dynamic imaging. This reduces its suitability for studying rapid or continuously evolving behavioral motion.

These limitations motivate the development of a continuous μ MRI acquisition framework, which is introduced in the following chapter as a more general approach for imaging tethered locomotion.

6 | Time resolved μ MRI

The contents of this chapter are based on the author's contributions to the article “*In situ* time-resolved motion of a tethered *Pachnoda marginata*, AI-correlated using μ MRI and optical imaging”, published in *Journal of Magnetic Resonance* (2026) [19].

The *in situ* treadmill platform and the methodologies developed in the previous chapters represent progressive steps toward the broader objective of enabling MRI of the behaving organism within a μ MRI framework. Prospective gating provided controlled, phase-specific “snapshot” imaging of distinct motion states, offering initial access to anatomical structure under defined locomotor conditions. However, establishing meaningful links between spontaneous behavior, or more broadly, the animal's natural state, and its internal anatomy or physiology ultimately requires truly continuous dynamic imaging. This led to the adoption of retrospective gating, which in principle can capture a broader spectrum of motion states (or abdominal strides) in the tethered insect locomotion in a single experiment. By retrospectively sorting continuously acquired data and incorporating deep learning–based residual correction, this approach refined the quality of the reconstructed images, thereby improving the robustness to motion inconsistencies.

Nevertheless, a fundamental limitation remained: insect locomotion is often irregular or semi-periodic, and fundamentally non-rigid. Conventional gating-based MRI methods depend on motion periodicity, and distribute data across motion bins, inherently limiting temporal resolution and obscuring transient events. As a result, they struggle to faithfully represent continuous body dynamics during naturalistic behavior. Fast dynamic MRI techniques offer a promising alternative, as demonstrated in various preclinical and clinical applications [110, 111]. Advances in acquisition hardware and reconstruction algorithms have enabled real-time imaging of rapid physiological processes, making dynamic MRI increasingly suitable for studies requiring both high spatial and temporal resolution [112, 128, 129]. In this chapter, the feasibility of accelerating dynamic 2D μ MRI within the constraints of the quadrature transceiver coil (no multi-channel sensitivity encoding) in a tethered, active insect is investigated, with the specific goal of visualizing abdominal stride patterns associated with locomotion.

Deep learning (DL)–based methods have emerged as state-of-the-art solutions for accelerated MRI, spanning the pipeline from acquisition design to image reconstruction [130,

131]. Among the many proposed architectures, U-Net [65]-derived supervised image-to-image translation networks remain strong and widely adopted baselines for undersampled reconstruction. However, data scarcity remains a persistent challenge when considering the adoption of supervised deep learning methods for various problems in MRI and even more so in μ MRI, as discussed in the previous chapter. However, transfer learning-based frameworks have proven effective in this context, specifically for the case of residual motion correction. This chapter extends that approach to the problem of undersampled dynamic MRI reconstruction.

The degree of undersampling, and consequently the achievable acceleration in dynamic MRI, is tightly coupled to the desired spatiotemporal resolution. Excessive undersampling can degrade reconstruction quality, particularly when deep learning is applied only as a post-hoc artifact correction step. k -space interpolation techniques such as **view-sharing (VS)** can partially mitigate this limitation by leveraging information from neighboring acquisitions, thereby improving the apparent temporal resolution in rapid imaging applications [21, 132]. However, aggressive view-sharing can introduce temporal blurring and motion-related artifacts when inconsistencies arise with the true underlying dynamics. To address this limitation, a hybrid reconstruction framework is proposed in which view-sharing first augments the undersampled data to increase effective data density per frame, followed by a deep learning-based refinement step to correct residual artifacts and improve fidelity of the dynamic images. This enhances the temporal resolution necessary for *in vivo* μ MRI characterization of abdominal pitching motion in tethered *Pachnoda marginata*.

In parallel, the integrated MR-compatible computer vision system (see Section 4.3 and Section 2.3.4) enables real-time capture of tethered locomotion dynamics and can be directly synchronized with time-resolved μ MRI frames. This multimodal acquisition framework facilitates direct correlation between μ MRI and optical imaging measurements, and provides a foundation for future investigations into the behavioral and biomechanical mechanisms underlying tethered locomotion.

6.1 Deep learning for MRI acceleration

Fast imaging in MRI is fundamentally constrained by the physics of magnetic resonance signal generation and spatial encoding. The phase encoding step is typically the most time-consuming component in sequential MRI acquisitions, as each phase-encoding step requires a separate repetition of the pulse sequence. Consequently, most MRI acceleration techniques focus on improving sampling efficiency by reducing the number of phase-encoding steps, for shortening the total acquisition time. This raises a central question in accelerated imaging: how much k -space information can be omitted while still preserving a faithful representation of the underlying anatomy. This trade-off defines the limits of undersampled MRI acquisition.

For the simplest case of single-coil Cartesian sampled MRI, image reconstruction is commonly performed using a direct inverse Fourier transform. However, when k -space is undersampled, this reconstruction becomes an ill-posed problem, leading to characteristic artifacts such as blurring or aliasing in the reconstructed images. The severity of these artifacts increases with higher acceleration factors and limits the achievable imaging speed with conventional reconstruction frameworks.

Deep learning-based approaches have emerged as an effective strategy for addressing undersampling artifacts in MRI. As reviewed extensively in the literature [130, 131], this problem can be formulated in several ways. One approach treats undersampled reconstruction as a direct inverse problem, in which a neural network learns a mapping from undersampled k -space to artifact-free images. A complementary and widely adopted strategy formulates the task as an image-to-image regression problem, where the network learns to transform an artifact-corrupted image into a clean reference image, effectively acting as a post-processing step. The method developed in this chapter follows the latter approach, leveraging deep learning-based image-domain correction to mitigate undersampling artifacts in accelerated MRI. Methodologically, this framework closely parallels the residual motion artifact correction strategy introduced in Chapter 5.

As encountered previously, the limited availability of insect μ MRI data poses a significant challenge for training deep learning models. Transfer learning therefore provides a practical and effective solution (Section 2.12). Pretraining on brain MRI datasets, which exhibit rich morphological diversity and high contrast, enables the network to learn generalizable representations of undersampling-related artifacts. These learned features can subsequently be refined through fine-tuning on insect μ MRI data, allowing adaptation to the distinct anatomical and textural characteristics of insect imaging. In all cases, undersampling effects must be synthetically introduced to generate paired datasets for the model training, validation, and testing. This simulation strategy for the dataset generation is detailed in the following sections.

6.1.1 MRI undersampling simulation

k -space-based simulations were used to generate undersampled MRI data for model training. The ground truth dataset used to generate the paired training images was drawn from a subset of the insect and brain MRI datasets [118] described in Section 5.3.1 of Chapter 5, consisting of 3,750 brain MRI slices and 200 inactive insect μ MRI slices.

For the brain dataset, the 3,750 fully sampled MR images of size 256×256 pixels served as ground truth. Each image was undersampled once using an independent variable-density mask, producing a total of 3,750 paired undersampled and fully sampled images. Undersampling was applied along the phase-encoding direction, with the central 15 k -space lines fully sampled and the peripheral lines randomly sampled with variable density, resulting in an overall sampling rate of approximately 13–33%.




Figure 6.1: Schematic representation of single coil variable density Cartesian undersampling MRI simulation. (a) Fully-sampled MRI. (b) Fully-sampled k -space. (c) Undersampling mask. (d) Undersampled k -space. (e) Zero-filling (ZF) MRI reconstruction. (Adapted from [19].)

For the insect dataset, 200 fully sampled MR images of size 256×256 pixels served as ground truth. Each image was undersampled multiple times using independent masks, producing a total of 500 paired undersampled and fully sampled images. In each k -space, the central 24 phase-encode lines were fully sampled, while the peripheral lines were undersampled with random variable-density pattern, resulting in an overall sampling fraction of 20–50%. The difference in central k -space sampling between the brain and insect datasets reflects the SNR and feature content: brain MRI has higher SNR and more structural detail, allowing fewer fully sampled lines, whereas the μ MRI insect data requires more central lines to preserve essential information. The full pipeline is illustrated in Fig. 6.1: (a) fully sampled magnitude image, (b) corresponding k -space magnitude, (c) undersampling mask, (d) undersampled k -space magnitude, and (e) zero-filled Fourier reconstruction magnitude image.

For model evaluation, a simulated insect test set consisting of 150 image pairs was retrospectively undersampled at effective acceleration factors of $R = 2\times$, $3\times$, and $5\times$. In all cases, additive complex Gaussian noise, scaled relative to the signal intensity, was applied to the fully sampled k -space prior undersampling to simulate realistic measurement conditions and enhance model robustness. This is particularly important because undersampling inherently reduces the SNR, resulting in noisier reconstructions.

6.1.2 Deep learning model training

The transfer learning-based training strategy (see Section 2.4.2 and Fig. 5.6) follows the supervised image-to-image regression framework described in Section 5.3.3. The same U-Net [65]-derived architecture, loss functions, quantitative image evaluation metrics,

and system setup are used without modification. Paired undersampled and fully sampled images served as input and target, respectively. The primary difference lies in the training parameters, which are summarized below.

For **pre-training**, the model was trained on the brain MRI dataset consisting of paired undersampled and fully sampled images, using a 3,000/750 training–validation split. Images were intensity-normalized and augmented through rotations and flips to enhance robustness. Training was formulated as a pixel-wise regression task using mean squared error (MSE) loss, with a batch size of 12 and the Adam optimizer [121] initialized with a learning rate of 1×10^{-4} . Early stopping with a patience of 20 epochs ($\Delta = 1 \times 10^{-4}$) was applied, and the learning rate was gradually reduced to 1×10^{-8} when the validation loss plateaued around epoch 303. The model architecture and all the reported hyperparameters were determined empirically through experimentation.

For end-to-end **fine-tuning**, training was performed on a mixed dataset comprising 500 brain MR image pairs (a subset of the pre-training dataset) and 500 insect MR image pairs, using an 800/200 training–validation split. All training hyperparameters were maintained as in pre-training, except for a reduced learning rate of 1×10^{-7} and a maximum of 15 epochs to mitigate overfitting. A combined MSE–SSIM loss function (0.2:0.8) was employed to balance pixel-level accuracy with perceptual image quality. Model performance was subsequently evaluated on both simulated and experimentally acquired undersampled images in the following sections.

6.1.3 Model evaluation on simulated data

The fine-tuned model was evaluated on a set of 150 simulated undersampled insect MR images at acceleration factors of 2×, 3×, and 5×, using the structural similarity index measure (SSIM) and gradient entropy (GE) as quantitative metrics. Ground truth k -space data were retrospectively undersampled following the procedure described in Section 6.1.1. ZF reconstructions for the respective acceleration factors are shown in Fig. 6.2(a–c), with the corresponding DL corrections displayed in Fig. 6.2(e–g). The fully sampled ground truth image is presented in Fig. 6.2(d). Quantitative image quality metrics are summarized in Table 6.1 and visualized as bar graphs in Fig. 6.3(a) and (b).

Variable-density undersampling primarily introduced blurring artifacts, which were consistently reduced by the deep learning correction across all evaluated acceleration factors. At higher acceleration rates, the reconstructions became progressively smoother, likely due to over-regularization by the network. This behavior is consistent with the limited information content inherent to single-coil (non-parallel imaging configuration) acquisitions, where aggressive undersampling restricts the recoverable spatial detail. Based on this trade-off between artifact suppression and spatial fidelity, an acceleration factor of 2.5× was selected for the experimental acquisitions described in the following sections.




Figure 6.2: DL correction of undersampling artifacts for the simulated dataset ($n = 150$). (a–c) ZF reconstructions at 2 $\times$, 3 $\times$, and 5 $\times$ acceleration, respectively. (d) Fully sampled ground truth image. (e–g) DL-corrected reconstructions for 2 $\times$, 3 $\times$, and 5 $\times$ acceleration. (Adapted from [19].)

Acceleration	GE		SSIM	
	ZF	DL	ZF	DL
2 $\times$	3.230 ± 0.051	2.154 ± 0.086	0.6461 ± 0.0709	0.7585 ± 0.0707
3 $\times$	3.453 ± 0.070	2.191 ± 0.095	0.5846 ± 0.0808	0.7447 ± 0.0699
5 $\times$	3.594 ± 0.087	2.284 ± 0.110	0.5361 ± 0.0853	0.7184 ± 0.0802

Table 6.1: GE and SSIM values for ZF and DL-corrected reconstructions of the simulated dataset at different acceleration factors (mean $\pm$ SD, $n = 150$).




Figure 6.3: Quantitative assessment of deep learning-based undersampling artifact correction across different acceleration factors. **(a)** Mean GE. **(b)** Mean SSIM. (Adapted from [19].)

6.2 Retrospective view-sharing for dynamic MRI

The experimental setup used throughout this thesis employs a quadrature transceiver coil (no multi-channel sensitivity encoding) configuration, which is typical in preclinical small animal MRI systems. Even with the incorporation of learned priors via the deep learning model, the fundamental Nyquist sampling constraint limits the degree of acceleration achievable in this setting. In contrast, clinical MRI systems benefit from parallel array coils with sensitivity encoding, which both improve SNR and enable higher acceleration factors using either classical or deep learning-based reconstruction methods. In essence, the achievable acceleration is dictated by the information content of k -space.

Classical interpolation methods can also effectively increase the apparent data density per frame in dynamic MRI k -space. One such technique is view-sharing, which reconstructs missing k -space lines by incorporating information from neighboring temporal frames. View-sharing is commonly combined with undersampling strategies to reduce scan time [21, 132, 133], thereby enhancing the effective temporal resolution in dynamic imaging. This approach has proven particularly valuable in applications where rapid acquisition is critical, such as cardiac imaging and contrast-enhanced studies [134, 135].

However, aggressive view-sharing expands the temporal footprint of each frame by incorporating k -space data from distant time points. When subject motion is faster than the sharing window, these temporal inconsistencies manifest as blurring and motion-related artifacts in the final reconstruction. Integrating compressed sensing (CS) reconstruction with view-sharing has been shown to reduce the temporal footprint and associated blurring. Restricting view-sharing to a small region of k -space, typically the most recent neighboring lines and reconstructing the remaining undersampled data with CS shortens the effective

temporal footprint without compromising image quality [136, 137]. Here, this concept is extended by replacing the CS reconstruction with a deep learning module, enabling rapid reconstruction with improved artifact suppression. The implementation details of this framework are described in the following sections.

6.2.1 k -space undersampling and view-sharing: implementation

Cartesian undersampling (non-parallel imaging configuration) was employed to acquire single-slice 2D sagittal MRI temporal frames. The k -space sampling strategy followed the schematic shown in Fig. 6.4. For each 200×200 k -space frame, the central 60 phase-encode lines were fully sampled to capture low-frequency information, complemented by 20 peripheral phase-encode lines randomly selected, with each of the 7 consecutive frames having a unique peripheral sampling pattern. This variable-density random sampling scheme, with a fully sampled k -space center, favors retention of low-frequency content such that undersampling artifacts predominantly manifest as blurring rather than aliasing.




Figure 6.4: Schematic representation of retrospective view-sharing with DL-corrected reconstruction. (a) Previous k -space. (b) Current k -space. (c) Next k -space. (d) View-shared MRI. (e) DL-corrected MRI. (Adapted from [19].)

A RARE sequence available in *Paravision 360v3* was modified to support acquisition with a custom phase-encoding table for each frame. Acquisition parameters were as follows: TR = 150 ms, TE = 5 ms, RARE factor = 10, FOV = 30 mm $\times$ 30 mm, undersampled matrix

size = 80×200 , in-plane spatial resolution = $150 \mu\text{m} \times 150 \mu\text{m}$, slice thickness = 0.7 mm, 140 repetitions, no signal averaging, and an acquisition time of 1.2 s per frame. The total acquisition time was 168 s, corresponding to an effective acceleration factor of $R = 2.5$. The sequence parameters and undersampling design were selected to favor better spatial resolution and temporal efficiency for depicting internal morphology, rather than optimizing for conventional tissue contrast. The resulting atypical contrast is adequate for resolving the structural features examined in this chapter.

Retrospective view-sharing was implemented by incorporating peripheral k -space lines from the two frames preceding and the two frames following each target frame (i.e., frames $n - 2$ through $n + 2$) to increase the effective data density per frame (Fig. 6.4(a–c) shows the case for $n - 1$ and $n + 1$). The resulting k -space was then processed using the discrete fast Fourier transform to produce the reconstructed MR images (Fig. 6.4(d)). The view-sharing reconstruction has been implemented using custom MATLAB (MathWorks, Natick, MA, USA, R2024b) scripts. The resulting reconstructions contained undersampling artifacts, which were subsequently corrected using the deep learning model as a downstream processing step (Fig. 6.4(e)).

6.2.2 Dynamic MRI: view-sharing with deep learning correction

The processing pipeline begins with undersampled k -space acquisition of a single 2D slice over 140 repetitions, producing dynamic image frames with an effective acquisition time of 1.2 s per frame, as described in Section 6.2.1. The k -space magnitude of a representative frame is shown in Fig. 6.5(a). The corresponding ZF reconstruction is presented in Fig. 6.5(b), followed by the DL-corrected ZF reconstruction in Fig. 6.5(c). For comparison, the view-shared k -space magnitude of the same frame is shown in Fig. 6.5(f), with its ZF reconstruction in Fig. 6.5(g) and the corresponding DL-corrected VS reconstruction in Fig. 6.5(h). The intensity–time profiles corresponding to the 140 frames for each of the above cases (i.e., Fig. 6.5(b), (c), (g), and (h)) are shown in Fig. 6.5(d), (e), (i), and (j), respectively.

Dynamic 2D RARE images were acquired with a repetition time of 150 ms. The relatively short repetition time limits full longitudinal magnetization recovery between successive excitations, leading to frame-to-frame signal intensity fluctuations, as observed in Fig. 6.5(d) and (e). Temporal smoothing introduced by the view-sharing step reduces these intensity oscillations (Fig. 6.5(i)), while subsequent DL-based correction restores spatial sharpness and structural detail, as evident in Fig. 6.5(j).

As ground-truth dynamic images are not available, reconstruction quality is evaluated using the GE metric, which is a reference-free measure. Lower GE values indicate better MR image quality. The GE values for the sample frames are provided directly in the figures. As expected, the VS reconstruction exhibits lower GE than the ZF reconstruction, and the DL correction further reduces GE for both cases. Although the added peripheral k -space lines improve spatial resolution, their lower SNR introduces additional noise into




Figure 6.5: DL-based correction of undersampling artifacts following VS reconstruction ($n = 140$ dynamic MR images). **(a)** Undersampled k -space. **(b)** ZF reconstruction. **(c)** DL-corrected ZF reconstruction. **(d)** Image–time profile along the blue line in (b) for ZF reconstruction. **(e)** Corresponding image–time profile for the DL-corrected reconstruction. **(f)** VS k -space. **(g)** VS reconstruction. **(h)** DL-corrected VS reconstruction. **(i)** Image–time profile for the VS reconstruction. **(j)** Corresponding image–time profile for the DL-corrected VS reconstruction. (Adapted from [19].)

the reconstruction. The DL model effectively suppresses this noise while preserving the resolution gain.

Figure 6.6(a) shows the GE distribution, computed using kernel density estimation (KDE), for ZF reconstructions and their DL-corrected counterparts, where DL correction reduces mean GE from 3.384 ± 0.214 to 2.421 ± 0.199 . Similarly, Fig. 6.6(b) shows the GE distribution for VS reconstructions before and after DL correction, with mean GE improving from 3.221 ± 0.092 to 2.228 ± 0.055 . Notably, VS reconstructions exhibit lower GE than ZF even prior to DL correction, and the combination of VS and DL correction yields the lowest GE overall, demonstrating improved image quality.




Figure 6.6: Comparison of GE distributions. **(a)** GE distribution for ZF and DL-corrected reconstructions. **(b)** GE distribution for VS and DL-corrected reconstructions. (Adapted from [19].)

6.2.3 Correlated dynamic μ MRI and optical imaging

The integrated MR-compatible camera system described in Section 4.3 provides real-time optical images of external morphology, temporally synchronized with DL-corrected dynamic μ MRI frames depicting internal anatomy (see Fig. 6.7). Dynamic MRI acquisition parameters are provided in Section 6.2.1. Intensity-time profiles along a selected line are shown for both modalities in Fig. 6.7(a) and (b). Representative frames illustrating different stride phases of the insect's abdominal pitching during tethered locomotion, with corresponding frame indices ($n = 140$) and times relative to the first displayed frame, are shown in Fig. 6.7(c) and (d). Due to behavioral variability, gross abdominal movements are intermittent, whereas internal organ motion dominates the dynamic sequence. Observed instances of gross abdominal motion are indicated in Fig. 6.7(a) and (b).

This demonstrates that μ MRI can resolve distinct abdominal pitching dynamics in a live, tethered insect. The internal morphological dynamics revealed by MRI can be directly correlated with external abdominal movements captured by the integrated optical imaging system. This multimodal framework enables non-invasive, simultaneous visualization of

internal and external abdominal kinematics during locomotion, providing access to dynamic information that is not attainable with conventional optical imaging alone.




Figure 6.7: Dynamic *in vivo* correlated μ MRI and optical imaging of the *in situ* behaving insect. (a) Intensity–time profile extracted from the dynamic optical image sequence ($n = 140$). (b) Intensity–time profile extracted from the dynamic μ MRI sequence ($n = 140$). Representative dynamic frames illustrating different stride phases of the insect’s abdominal pitching during tethered locomotion, with corresponding frame indices and times (relative to the first displayed frame): (c) Optical image frames. (d) μ MRI frames (green line: reference position). (Adopted from [19].)

6.3 Discussion and conclusion

Dynamic and real-time MRI are typically performed using gradient recalled echo (GRE) or echo planar imaging (EPI)-based sequences, often combined with advanced non-Cartesian trajectories. In this chapter, however, the RARE (or Fast Spin-Echo) sequence was used to overcome the low SNR and severe susceptibility artifacts introduced by the beetle’s dense chitinous exoskeleton and internal air pockets, which are limitations commonly encountered in many invertebrate models at high-field μ MRI [13, 12], which has been discussed in Section 3.3.2. The spin-echo refocusing train in RARE provides intrinsically high SNR and robust suppression of signal loss and geometric distortion arising from steep susceptibility gradients at exoskeleton-tissue interfaces, whereas GRE-based non-Cartesian methods such as spiral and radial imaging are generally more prone to T_2^* -related dephasing and off-resonance blurring or distortion under these conditions.

View-sharing was incorporated as a preprocessing step to enhance the apparent information content of the undersampled k -space, thereby improving the subsequent deep learning

correction and the effective temporal fidelity of the reconstructed dynamic images. Here, the motion model evolves slowly enough that minor view-sharing artifacts can be neglected, and the deep learning model focuses solely on mitigating undersampling artifacts. While aggressive view-sharing may lead to temporal blurring, explicitly modeling it within the deep learning framework can further improve artifact correction and the fidelity of the reconstructed images.

In principle, end-to-end deep learning-based MRI reconstruction models are capable of implicitly learning temporal correlations present in dynamic data to enhance temporal resolution [138]. However, this requires large, representative datasets, which are difficult to obtain for non-rigid, irregular or unpredictable motion, such as the case of tethered insect locomotion. View-sharing offers a classical, relatively motion-agnostic strategy to fill missing k -space lines using information from neighboring frames, supplying the network with an improved initial estimate, and reducing the temporal footprint. Importantly, this method is not restricted to RARE sequences and can be applied to other imaging protocols with minimal adjustment.

More broadly, the reconstruction framework developed in this chapter illustrates a hybrid paradigm in which classical signal processing techniques and learning-based methods are combined to achieve improved performance under practical constraints. Rather than relying solely on data-driven learning to infer all temporal structure, model-based or interpolation-driven steps, such as view-sharing introduce explicit prior assumptions about temporal redundancy directly into the reconstruction pipeline. This can stabilize learning, reduce the burden on the network, and improve generalization when training data are limited, as in the case of insect μ MRI. A variety of related strategies may be explored for similar problems. For example, Kalman filter-based approaches, which have been used to model motion dynamics in time-resolved imaging, provide a model-based, estimation-theoretic framework for incorporating temporal state evolution and measurement uncertainty [139, 140]. While view-sharing relies on implicit temporal interpolation between neighboring frames, Kalman filtering explicitly models temporal state evolution and measurement uncertainty, enabling recursive estimation of dynamic processes. Such methods have been successfully applied to accelerated Cartesian dynamic cardiac imaging, demonstrating improved temporal fidelity and robustness to undersampling by incorporating motion dynamics directly into the reconstruction process [141].

Integrating these state-space modeling approaches with deep learning-based reconstruction could therefore offer advantages over purely interpolation-based schemes. By combining physics-informed temporal modeling with data-driven spatial priors, such hybrid frameworks have the potential to better preserve rapid or transient events while maintaining stability under high acceleration factors. This may be particularly beneficial in scenarios involving structured but partially unpredictable motion, where simple view-sharing may introduce temporal blurring or fail to capture abrupt changes. Although motivated here by the constraints of insect μ MRI, these challenges and methodological considerations are

not unique to this domain. Similar trade-offs between undersampling, motion complexity, data scarcity, and reconstruction robustness arise in many areas of clinical dynamic MRI, including cardiac, abdominal, and fetal imaging. As such, the hybrid strategies discussed here reflect broader themes in modern MRI reconstruction, where the integration of physics-based modeling and learning-based methods is increasingly important for reliable imaging under non-ideal and data-limited conditions [142, 143].

In summary, the methodology developed in this chapter is not intended to achieve state-of-the-art high-speed dynamic MRI performance, and the resulting temporal resolution of 1.2 s per frame at $2.5\times$ acceleration is therefore not directly comparable to advanced clinical research protocols, which benefit from larger voxels and parallel coil arrays. The significantly smaller voxel volume of the insect μ MRI limits the SNR even at the ultra-high field of 15.2 T, and consequently, the attainable acceleration rate. As a result, the proof-of-concept demonstrated here is suited to imaging the relatively slow abdominal dynamics of the *Pachonoda marginata*, which may not hold for a smaller, faster model organism such as *Drosophila melanogaster*. To address such cases, future work could explore strategies for improving temporal resolution, drawing on approaches established in preclinical and clinical dynamic MRI to accommodate rapid and complex motion [111, 112, 144]. Even modest custom parallel arrays may provide the necessary acceleration to track dynamic processes *in vivo*. When paired with advanced reconstruction, this approach offers a viable path toward high-fidelity imaging of rapid behaviors in the constrained environment of a μ MRI scanner.

7 | ***In vivo* MRS: towards functional characterization**

The fundamental goal of “MRI of behaving organism” in this thesis represents a subset of the broader goal of functionally characterizing a live, behaving model organism using magnetic resonance. The methodologies developed in the preceding chapters enable both snapshot imaging during active tethered locomotion and time-resolved imaging, establishing a foundation for probing physiological processes beyond static morphology. While the primary focus thus far has been on biomechanical characterization of insect locomotion, these advances naturally motivate extension toward more direct functional and biochemical measurements.

An overarching objective is the acquisition of molecular signatures associated with behavior. Magnetic resonance spectroscopy (MRS), which shares the same physical basis as MRI and underpins its contrast mechanisms, provides localized chemical or metabolic information that is inaccessible to conventional imaging alone. In principle, spatially localized *in vivo* MRS experiments enable interrogation of tissue-specific biochemical states during behavior.

As an initial step toward this goal, preliminary *in vivo* MRS experiments were conducted to assess the feasibility of such measurements in a tethered insect positioned on the treadmill. These experiments are not intended as comprehensive functional studies, but rather as a proof of viability for extending the developed experimental framework toward future functional and spectroscopic investigations.

7.1 *In vivo* localized ^1H MRS

In vivo MRS can be performed using a standard, non-localized approach, in which the entire sample is excited and the resulting free induction decay (FID) is acquired without any spatial specificity. While simple, this approach provides only a global signal and can suffer from inhomogeneities arising from tissue heterogeneity and magnetic susceptibility variations. In contrast, localized MRS exploits imaging gradients to selectively excite and acquire signals from a defined voxel or anatomical region, enabling spectroscopic measurements from specific tissues of interest. Although localization typically reduces the overall SNR due to

the smaller excited volume, it offers significant advantages by minimizing line broadening from magnetic field inhomogeneities and potentially provide tissue-specific biochemical information [145, 146].

Localized ^1H MRS experiments were performed using a point resolved spectroscopy (PRESS) sequence (ParaVision 360 v3.3) on the Bruker preclinical MRI system. All measurements were conducted using the *in situ* treadmill setup, with the quadrature transmit–receive RF coil configuration, consistent with the imaging experiments described throughout the thesis. The insect was alive and exhibited periods of active behavior throughout the experiments, and detailed animal handling protocols are outlined in Section 4.4.

7.1.1 Spectroscopic sequence and parameters

Spatial localization was achieved using a standard PRESS scheme consisting of a 90° excitation pulse followed by two orthogonal 180° refocusing pulses. TE was set to 16 ms, implemented as sub-echo times of $\text{TE}_1 = 8.03$ ms and $\text{TE}_2 = 7.97$ ms. The TR was 1000 ms, providing a balance between SNR and total acquisition time while limiting excessive saturation effects. Each spectrum was acquired using 50 signal averages, and the experiment was repeated for 60 repetitions, resulting in a total acquisition time of approximately 3000 s.

Spectroscopic data were acquired in one dimension with a spectral matrix size of 2048 points. A spectral bandwidth of 13.16 kHz was used, corresponding to a dwell time of $38\ \mu\text{s}$ and a nominal spectral resolution of approximately 3.2 Hz. The transmitter frequency was referenced to the water resonance at 4.7 ppm. Water suppression was not applied during acquisition in order to preserve the full spectral content and facilitate post hoc assessment of spectral quality and frequency stability.

A single cubic voxel of dimensions $2 \times 2 \times 2\ \text{mm}^3$ was selected within the thorax region of the insect, as shown in Fig. 7.1(a) and (b). Crusher gradients and spoiler pulses were applied to minimize signal contamination from outside the prescribed volume. First-order shimming was performed automatically over the selected voxel prior to data acquisition to optimize magnetic field homogeneity. Transmitter reference power and frequency calibration were carried out using the system’s automated adjustment routines. No dynamic shimming or navigator-based frequency correction was applied during the spectroscopy acquisition.

Raw FID data were saved for offline processing and analysis in MATLAB, no real-time spectral fitting or metabolite quantification was performed during acquisition. For each measurement, the time-domain signals were averaged, followed by an FFT to convert the signals into the frequency domain. The resulting spectra were visualized as magnitude spectra without applying phase or baseline correction. The primary objective was to evaluate signal stability over extended acquisition times. This approach captures overall amplitude fluctuations while remaining insensitive to minor phase variations induced by internal motion, thereby providing a robust measure of motion stability in the tethered insect.




Figure 7.1: *In vivo* localized ^1H MRS of a tethered insect. **(a)** Insect on the *in situ* treadmill. **(b)** Anatomical reference indicating voxel placement within the thorax. **(c)** Magnitude spectrum highlighting water and lipid resonances. **(d)** Time-resolved stack plot illustrating spectral stability over a 50 min duration.

Figure 7.1(a) illustrates the placement of the localized spectroscopy voxel within the thorax. The corresponding ^1H MRS magnitude spectrum is shown in Fig. 7.1(b), dominated by the water resonance at approximately 4.7 ppm and broad lipid-associated resonances. Figure 7.1(c) presents a time-resolved stack plot of spectra acquired over a 50 min period, demonstrating spectral stability throughout the acquisition. The primary objective of this experiment was to qualitatively assess whether normal tethered locomotion induces large-scale signal fluctuations or instability in localized *in vivo* MRS measurements. Consistent with the approach demonstrated by Schilling *et al.* [147], ^1H NMR can separate water and lipid resonances in living insects and non-invasively monitor fat distribution; however, in contrast to their fixed and cold-anesthetized preparations, the present experiment demonstrates that meaningful ^1H measurements can also be obtained in actively tethered insects.

Although localized ^1H MRS provides high sensitivity and serves as a practical starting point for assessing motion robustness in tethered insects, its biochemical specificity is limited by the dominant water signal and broad lipid resonances. Consequently, the current measurements primarily demonstrate the feasibility and stability of localized spectroscopy during active behavior rather than detailed metabolic characterization. In future studies, the same experimental framework could be extended to alternative nuclei commonly employed in small-animal MRS, such as ^{31}P , which offers direct insight into energy metabolism and intracellular pH, or ^{13}C for probing metabolic pathways. While the SNR for these nuclei is expected to be substantially lower than for ^1H , optimized acquisition strategies or even hyperpolarization techniques could be explored to enhance sensitivity, presenting an ambitious yet promising avenue for future work on dynamic metabolic measurements in *in vivo* behaving insects.

8 | Conclusions and outlook

This chapter summarizes and reflects on the major scientific contributions of the thesis and outlines future directions for research in this area.

8.1 Conclusions

While the physical principles underlying clinical MRI and μ MRI are fundamentally the same, their practical implementations and scientific motivations may differ. Clinical MRI is generally optimized for imaging larger anatomical regions, whereas μ MRI targets much smaller specimens and finer spatial scales, where sensitivity and spatial resolution pose greater technical challenges. In addition, μ MRI often relies on higher static magnetic field strengths and stronger gradients to achieve microscopic resolution, which can introduce increased susceptibility effects, field inhomogeneities, and RF-related constraints. These differences in operating field regimes further complicate experimental design and quantitative interpretation at the microscale. As discussed by Glover and Mansfield [148], early developments in magnetic resonance microscopy (MRM) (or μ MRI) were constrained by SNR limitations that hindered its adoption as a widespread microscopic imaging tool.

Introducing motion at this scale further complicates the problem. At the microscopic spatial resolutions, even small displacements can constitute a significant fraction of a voxel dimension, leading to phase errors, blurring, and signal loss. High-resolution μ MRI often requires longer acquisition times and stronger gradients, and the measurements are proportionally more sensitive to motion-induced inconsistencies than in conventional clinical imaging. This places additional demands on hardware design, sequence development, and reconstruction strategies. Nevertheless, these constraints do not represent a fundamental limit, and innovations in methodology and experimental design can expand the scope and impact of μ MRI, as demonstrated in this thesis.

While the treadmill platform established the physical interface required for μ MRI of tethered insect locomotion, the methodological developments transformed this experimental paradigm into a viable imaging approach. In particular, motion management strategies played a central role. Prospective gating, although not ideal for irregular motion, demonstrated the feasibility of snapshot imaging during active locomotion. Building on this,

retrospective gating was implemented to relax acquisition constraints and enable more flexible reconstruction, including the possibility of dynamic imaging. Residual motion-related artifacts following gating were further mitigated using the deep learning-based strategy. Given the complexity of motion at the microscale, where displacements can approach the spatial resolution and interact nonlinearly with encoding fields, precise motion parameter estimation and analytical correction remain highly challenging. In this context, data-driven methods provide a viable route to overcome microscale motion challenges, advancing μ MRI toward capturing the intrinsic dynamics of active biological systems.

Truly dynamic imaging was demonstrated in the final part of this thesis through a proof-of-concept 2D time-resolved μ MRI of abdominal pitching during tethered locomotion. Fast imaging enabled visualization of internal motion at temporal resolutions sufficient to capture natural biomechanical dynamics of the sun beetle. Simultaneously, the integrated MR-compatible computer vision system provided complementary information on the correlated external morphology, allowing simultaneous observation of internal and external motion. This integrated framework enables future exploration of the interplay between internal physiological processes and externally expressed behaviors in active small organisms. In addition, the preliminary localized *in vivo* MRS experiments demonstrate the feasibility of acquiring stable ^1H signals from a tethered insect, indicating that the developed μ MRI framework can, in principle, be extended toward future spectroscopic investigations of active organisms.

In summary, the key contributions of this thesis are the development of an *in situ* treadmill paradigm for μ MRI of a tethered, actively behaving insect, along with associated methodologies for motion compensation and accelerated imaging. By extending μ MRI beyond purely structural characterization to frameworks compatible with natural activity, this work demonstrates the feasibility of studying dynamic biological processes and highlights opportunities for future functionally relevant investigations in active organisms.

8.2 Outlook

Several methodological and application-oriented directions can be pursued as future directions for this work. From a technical perspective, the current treadmill setup is not yet optimized for sensitivity or coil filling factor. Further improvements in custom RF coil design, sample positioning, and overall system integration could enhance the SNR. Such developments should ideally be guided by clearly defined biological questions, ensuring that hardware optimization aligns with the potential functional or structural measurements.

The commercial quadrature coil used throughout this thesis was not designed with parallel imaging capabilities. Although implementing parallel imaging at the μ MRI scale presents significant technical challenges, particularly due to geometric constraints, decoupling challenges between coil elements, and the need to preserve adequate SNR at high spatial resolution, there are potential pathways toward achieving it. Parallel imaging could be realized through the development of a modest multi-element receive array specifically tailored to the size and morphology of the insect specimen. Such designs would require careful electromagnetic optimization to ensure sufficient spatial variation in coil sensitivities, effective decoupling between elements, and low noise correlation. When combined with sensitivity-based reconstruction approaches, this may allow moderate acceleration, albeit with the expected SNR and g-factor penalties inherent to parallel imaging. While demanding in terms of hardware design and validation, these developments represent a feasible direction for improving temporal efficiency in future μ MRI implementations, especially for the case of dynamic imaging.

Tethered insect locomotion, and its interaction with the fields (B_0 , B_1 , and gradients) can also influence image quality and quantitative accuracy, but this has not yet been systematically investigated. Motion-induced changes in coil loading, susceptibility effects, or field perturbations could introduce subtle biases, particularly in quantitative measurements. Explicit electromagnetic modeling may help characterize these interactions and assess their impact, especially if future studies aim to extract quantitative parameters such as relaxation times or even metabolite concentrations. A more detailed understanding of field–tissue interactions at the insect scale would therefore contribute to more robust experimental design and more reliable interpretation of μ MRI measurements.

The preliminary localized *in vivo* MRS experiments highlight a promising direction for future studies. Although technically ambitious, developing a more robust MRS platform and workflow for active insects, with improved sensitivity and stability, could enable non-invasive monitoring of metabolic changes associated with development, behavior, or environmental perturbations. Motion is only one of the challenges; achieving sufficient sensitivity for X-nuclei MRS presents an even greater limitation. Future approaches may need to explore hyperpolarization techniques, which have increasingly been applied in preclinical small animal models to enhance the detectability of low-abundance metabolites.

Deep learning methods are increasingly shaping medical imaging and computer vision, with strategies such as transfer learning and physics-informed modeling helping to address data scarcity challenges. μ MRI may benefit from these advances, particularly in artifact removal, denoising, reconstruction and image analysis. However, as in clinical MRI, data-driven approaches must remain grounded in the underlying physical principles to ensure robustness and quantitative reliability, considering the limited data availability and experimental variability inherent to μ MRI.

The sun beetle (*Pachnoda marginata*) primarily served as a practical model to establish and validate the methodological framework underlying the thesis. While effective for demonstrating feasibility, the approaches developed here could be extended to more widely studied model organisms, such as *Drosophila melanogaster*. Moreover, the *in situ* treadmill paradigm is not inherently restricted to insects. With appropriate adaptation, the conceptual framework may prove useful, or serve as inspiration, for studies involving other small animal models commonly used in preclinical MRI, particularly in contexts where controlled activity or behavior during imaging is of interest.

Overall, the work presented in this thesis provides a methodological starting point. Continued refinement of instrumentation, motion-adaptive strategies, modeling, and multimodal integration may allow the approach to mature into a versatile tool for dynamic non-invasive studies of small living systems.

Bibliography

- [1] Johannes D Seelig, M Eugenia Chiappe, Gus K Lott, Anirban Dutta, Jason E Osborne, Michael B Reiser, and Vivek Jayaraman. “Two-photon calcium imaging from head-fixed *Drosophila* during optomotor walking behavior”. In: *Nature methods* 7.7 (2010), pp. 535–540.
- [2] Arne F Meyer, Jasper Poort, John O’Keefe, Maneesh Sahani, and Jennifer F Linden. “A head-mounted camera system integrates detailed behavioral monitoring with multichannel electrophysiology in freely moving mice”. In: *Neuron* 100.1 (2018), pp. 46–60.
- [3] Polina Anikeeva, Aaron S Andalman, Ilana Witten, Melissa Warden, Inbal Goshen, Logan Grosenick, Lisa A Gunaydin, Loren M Frank, and Karl Deisseroth. “Optetrode: a multichannel readout for optogenetic control in freely moving mice”. In: *Nature neuroscience* 15.1 (2012), pp. 163–170.
- [4] Tomy dos Santos Rolo, Alexey Ershov, Thomas van de Kamp, and Tilo Baumbach. “In vivo X-ray cine-tomography for tracking morphological dynamics”. In: *Proceedings of the National Academy of Sciences* 111.11 (2014), pp. 3921–3926.
- [5] Thomas van de Kamp, A Ershov, T dos Santos Rolo, A Riedel, and T Baumbach. “Insect imaging at the ANKA synchrotron radiation facility”. In: *Entomologie heute* 25 (2013), pp. 147–160.
- [6] Maxim Zaitsev, Julian Maclaren, and Michael Herbst. “Motion artifacts in MRI: A complex problem with many partial solutions”. In: *Journal of Magnetic Resonance Imaging* 42.4 (2015), pp. 887–901.
- [7] Andre Van Der Kouwe and Jalal B Andre. *Motion correction in MR: correction of position, motion, and dynamic field changes*. Vol. 6. Academic Press, 2022.
- [8] Zhifeng Liang, Glenn DR Watson, Kevin D Alloway, Gangchea Lee, Thomas Neuberger, and Nanyin Zhang. “Mapping the functional network of medial prefrontal cortex by combining optogenetics and fMRI in awake rats”. In: *Neuroimage* 117 (2015), pp. 114–123.

- [9] Jean A King, Timothy S Garelick, Mathew E Brevard, Wei Chen, Tara L Messenger, Timothy Q Duong, and Craig F Ferris. "Procedure for minimizing stress for fMRI studies in conscious rats". In: *Journal of neuroscience methods* 148.2 (2005), pp. 154–160.
- [10] Zhe Han, Wenjing Chen, Xifan Chen, Kaiwei Zhang, Chuanjun Tong, Xiaoxing Zhang, Chengyu T Li, and Zhifeng Liang. "Awake and behaving mouse fMRI during Go/No-Go task". In: *Neuroimage* 188 (2019), pp. 733–742.
- [11] Alan Jasanoff and Phillip Z. Sun. "In vivo magnetic resonance microscopy of brain structure in unanesthetized flies". In: *Journal of Magnetic Resonance* 158.1-2 (2002), pp. 79–85.
- [12] Alexander Ziegler, Martin Kunth, Susanne Mueller, Christian Bock, Rolf Pohmann, Leif Schröder, Cornelius Faber, and Gonzalo Giribet. "Application of magnetic resonance imaging in zoology". In: *Zoomorphology* 130 (2011), pp. 227–254.
- [13] Stefan Wecker, Thomas Hörschemeyer, and Mathias Hoehn. "Investigation of insect morphology by MRI: Assessment of spatial and temporal resolution". In: *Magnetic Resonance Imaging* 20.1 (2002), pp. 105–111.
- [14] Christian Bock, Markus Frederich, Rolf-M Wittig, and Hans-O Pörtner. "Simultaneous observations of haemolymph flow and ventilation in marine spider crabs at different temperatures: a flow weighted MRI study". In: *Magnetic Resonance Imaging* 19.8 (2001), pp. 1113–1124.
- [15] Christian Bock, Franz-Josef Sartoris, and Hans-Otto Pörtner. "In vivo MR spectroscopy and MR imaging on non-anaesthetized marine fish: techniques and first results". In: *Magnetic Resonance Imaging* 20.2 (2002), pp. 165–172.
- [16] A Van der Linden, M Verhoye, Hans-Otto Pörtner, and Christian Bock. "The strengths of in vivo magnetic resonance imaging (MRI) to study environmental adaptational physiology in fish". In: *Magnetic Resonance Materials in Physics, Biology and Medicine* 17 (2004), pp. 236–248.
- [17] Ajmal Chenakkara, Mazin Jouda, Ulrike Wallrabe, and Jan G Korvink. "Motion compensated magnetic resonance imaging of an active sun beetle using an in situ treadmill". In: *Scientific Reports* 15.1 (2025), p. 40340.
- [18] Ajmal Chenakkara, Mazin Jouda, Ulrike Wallrabe, and Jan G Korvink. "Residual motion artifact removal enables dynamic μ MRI of a behaving *Pachnoda marginata*". In: *Journal of Magnetic Resonance* (2025), p. 107954.
- [19] Ajmal Chenakkara, Mazin Jouda, Ulrike Wallrabe, and Jan G Korvink. "In situ time-resolved motion of a tethered *Pachnoda marginata*, AI-correlated using μ MRI and optical imaging". In: *Journal of Magnetic Resonance* 389 (2026), p. 108089.

-
- [20] Robert W Brown, Y-C Norman Cheng, E Mark Haacke, Michael R Thompson, and Ramesh Venkatesan. *Magnetic resonance imaging: physical principles and sequence design*. Hoboken, NJ: John Wiley & Sons, 2014.
 - [21] Matt A Bernstein, Kevin F King, and Xiaohong Joe Zhou. *Handbook of MRI pulse sequences*. Burlington, MA: Elsevier, 2004.
 - [22] Malcolm H Levitt. *Spin dynamics: Basics of nuclear magnetic resonance*. Chichester: John Wiley & Sons, 2008.
 - [23] Paul T Callaghan. *Principles of nuclear magnetic resonance microscopy*. Oxford: Clarendon Press, 1993.
 - [24] W Thomas Dixon. “Simple proton spectroscopic imaging.” In: *Radiology* 153.1 (1984), pp. 189–194.
 - [25] Novaksean. *Normal axial T2-weighted MR image of the brain*. Wikimedia Commons, CC BY-SA 4.0. 2015. URL: https://commons.wikimedia.org/wiki/File:Normal_axial_T2-weighted_MR_image_of_the_brain.jpg.
 - [26] J. Hennig, A. Nauerth, and H. Friedburg. “RARE imaging: A fast imaging method for clinical MR”. In: *Magnetic Resonance in Medicine* 3.6 (1986), pp. 823–833.
 - [27] Felix Bloch. “Nuclear induction”. In: *Physical review* 70.7-8 (1946), p. 460.
 - [28] Edward M Purcell, Henry Cutler Torrey, and Robert V Pound. “Resonance absorption by nuclear magnetic moments in a solid”. In: *Physical review* 69.1-2 (1946), p. 37.
 - [29] Herman Y Carr and Edward M Purcell. “Effects of diffusion on free precession in nuclear magnetic resonance experiments”. In: *Physical review* 94.3 (1954), p. 630.
 - [30] Edward O Stejskal and John E Tanner. “Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient”. In: *The journal of chemical physics* 42.1 (1965), pp. 288–292.
 - [31] Christian Federau and Daniel Gallichan. “Motion-correction enabled ultra-high resolution in-vivo 7T-MRI of the brain”. In: *PloS one* 11.5 (2016), e0154974.
 - [32] F Godenschweger, U Kägebein, D Stucht, U Yarach, A Sciarra, R Yakupov, F Lüsebrink, P Schulze, and O Speck. “Motion correction in MRI of the brain”. In: *Physics in Medicine and Biology* 61.5 (2016).
 - [33] Bastiaan Driehuys, John Nouis, Alexandra Badea, Elizabeth Bucholz, Ketan Ghaghada, Alexandra Petiet, and Laurence W Hedlund. “Small animal imaging with magnetic resonance microscopy”. In: *ILAR journal* 49.1 (2008), pp. 35–53.
 - [34] Veronika Spieker, Hannah Eichhorn, Kerstin Hammernik, Daniel Rueckert, Christine Preibisch, Dimitrios C Karampinos, and Julia A Schnabel. “Deep learning for retrospective motion correction in MRI: a comprehensive review”. In: *IEEE Transactions on Medical Imaging* 43.2 (2023), pp. 846–859.

- [35] Thomas Küstner, Karim Armanious, Jiahuan Yang, Bin Yang, Fritz Schick, and Sergios Gatidis. “Retrospective correction of motion-affected MR images using deep learning frameworks”. In: *Magnetic resonance in medicine* 82.4 (2019), pp. 1527–1540.
- [36] M. Zaitsev, C. Dold, G. Sakas, J. Hennig, and O. Speck. “Magnetic resonance imaging of freely moving objects: prospective real-time motion correction using an external optical motion tracking system”. In: *NeuroImage* 31.3 (2006), pp. 1038–1050.
- [37] R. L. Ehman and J. P. Felmlee. “Adaptive technique for high-definition MR imaging of moving structures”. In: *Radiology* 173.1 (1989), pp. 255–263.
- [38] Hope W Korin, Joel P Felmlee, Richard L Ehman, and Stephen J Riederer. “Adaptive technique for three-dimensional MR imaging of moving structures.” In: *Radiology* 177.1 (1990), pp. 217–221.
- [39] Bruno Madore, Aaron T Hess, Adam MJ van Niekerk, Daniel C Hoinkiss, Patrick Hucker, Maxim Zaitsev, Onur Afacan, and Matthias Günther. “External hardware and sensors, for improved MRI”. In: *Journal of Magnetic Resonance Imaging* 57.3 (2023), pp. 690–705.
- [40] Melvyn B Ooi, Murat Aksoy, Julian Maclaren, Ronald D Watkins, and Roland Bammer. “Prospective motion correction using inductively coupled wireless RF coils”. In: *Magnetic resonance in medicine* 70.3 (2013), pp. 639–647.
- [41] JL Ackerman, MC Offutt, RB Buxton, TJ Brady, et al. “Rapid 3D tracking of small RF coils”. In: *Proceedings of the 5th annual meeting of SMRM, Montreal, Canada*. 1986, p. 1131.
- [42] Lei Qin, Peter van Gelderen, John Andrew Derbyshire, Fenghua Jin, Jongho Lee, Jacco A de Zwart, Yang Tao, and Jeff H Duyn. “Prospective head-movement correction for high-resolution MRI using an in-bore optical tracking system”. In: *Magnetic Resonance in Medicine* 62.4 (2009), pp. 924–934.
- [43] Julian Maclaren, Brian SR Armstrong, Robert T Barrows, KA Danishad, Thomas Ernst, Colin L Foster, Kazim Gumus, Michael Herbst, Ilja Y Kadashevich, Todd P Kusik, et al. “Measurement and correction of microscopic head motion during magnetic resonance imaging of the brain”. In: *PloS one* 7.11 (2012), e48088.
- [44] Jessica Schulz, Thomas Siegert, Enrico Reimer, Christian Labadie, Julian Maclaren, Michael Herbst, Maxim Zaitsev, and Robert Turner. “An embedded optical tracking system for motion-corrected magnetic resonance imaging at 7T”. In: *Magnetic Resonance Materials in Physics, Biology and Medicine* 25 (2012), pp. 443–453.
- [45] Phillip DiGiacomo, Julian Maclaren, Murat Aksoy, Elizabeth Tong, Mackenzie Carlson, Bryan Lanzman, Syed Hashmi, Ronald Watkins, Jarrett Rosenberg, Brian Burns, et al. “A within-coil optical prospective motion-correction system for brain imaging at 7T”. In: *Magnetic resonance in medicine* 84.3 (2020), pp. 1661–1671.

-
- [46] Brian C Andrews-Shigaki, Brian SR Armstrong, Maxim Zaitsev, and Thomas Ernst. “Prospective motion correction for magnetic resonance spectroscopy using single camera retro-grate reflector optical tracking”. In: *Journal of magnetic resonance imaging* 33.2 (2011), pp. 498–504.
 - [47] Jakob M. Slipsager, Andreas H. Ellegaard, Stefan L. Glimberg, Rasmus R. Paulsen, M. Dylan Tisdall, Paul Wighton, André Van Der Kouwe, Lisbeth Marner, Otto M. Henriksen, Ian Law, and Oline V. Olesen. “Markerless motion tracking and correction for PET, MRI, and simultaneous PET/MRI”. In: *PLoS ONE* 14.4 (2019).
 - [48] Johan Berglund, Adam van Niekerk, Henric Rydén, Tim Sprenger, Enrico Avventi, Ola Norbeck, Stefan L. Glimberg, Oline V. Olesen, and Stefan Skare. “Prospective motion correction for diffusion weighted EPI of the brain using an optical markerless tracker”. In: *Magnetic Resonance in Medicine* 85.3 (2021), pp. 1427–1440.
 - [49] Robert Frost, Paul Wighton, F Işık Karahanoğlu, Richard L Robertson, P Ellen Grant, Bruce Fischl, M Dylan Tisdall, and André van der Kouwe. “Markerless high-frequency prospective motion correction for neuroanatomical MRI”. In: *Magnetic resonance in medicine* 82.1 (2019), pp. 126–144.
 - [50] Bruce D Lucas and Takeo Kanade. “An iterative image registration technique with an application to stereo vision”. In: *IJCAI’81: 7th international joint conference on Artificial intelligence*. Vol. 2. 1981, pp. 674–679.
 - [51] Berthold KP Horn and Brian G Schunck. “Determining optical flow”. In: *Artificial intelligence* 17.1-3 (1981), pp. 185–203.
 - [52] E. H. Moore. “On the reciprocal of the general algebraic matrix”. In: *Bulletin of the American Mathematical Society* 26.9 (1920), pp. 394–395.
 - [53] Roger Penrose. “A generalized inverse for matrices”. In: *Proceedings of the Cambridge Philosophical Society* 51.3 (1955), pp. 406–413.
 - [54] Chang Min Hyun, Hwa Pyung Kim, Sung Min Lee, Sungchul Lee, and Jin Keun Seo. “Deep learning for undersampled MRI reconstruction”. In: *Physics in Medicine & Biology* 63.13 (2018), p. 135007.
 - [55] Zhaolin Chen, Kamlesh Pawar, Mevan Ekanayake, Cameron Pain, Shenjun Zhong, and Gary F Egan. “Deep learning for image enhancement and correction in magnetic resonance imaging—state-of-the-art and challenges”. In: *Journal of Digital Imaging* 36.1 (2023), pp. 204–230.
 - [56] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. Cambridge, MA: MIT Press, 2016.
 - [57] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. “Deep learning”. In: *nature* 521.7553 (2015), pp. 436–444.

- [58] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. “Attention is all you need”. In: *Advances in neural information processing systems* 30 (2017).
- [59] Christopher M Bishop and Hugh Bishop. *Deep learning: Foundations and concepts*. Cham, Switzerland: Springer Nature, 2023.
- [60] Christopher M. Bishop. *Pattern recognition and machine learning*. New York, NY: Springer, 2006.
- [61] Frank Rosenblatt. “The perceptron: a probabilistic model for information storage and organization in the brain.” In: *Psychological review* 65.6 (1958), p. 386.
- [62] George Cybenko. “Approximation by superpositions of a sigmoidal function”. In: *Mathematics of control, signals and systems* 2.4 (1989), pp. 303–314.
- [63] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. “Learning representations by back-propagating errors”. In: *nature* 323.6088 (1986), pp. 533–536.
- [64] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. “Gradient-based learning applied to document recognition”. In: *Proceedings of the IEEE* 86.11 (2002), pp. 2278–2324.
- [65] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. “U-net: Convolutional networks for biomedical image segmentation”. In: *International Conference on Medical image computing and computer-assisted intervention*. Springer. 2015, pp. 234–241.
- [66] Sinno Jialin Pan and Qiang Yang. “A survey on transfer learning”. In: *IEEE Transactions on knowledge and data engineering* 22.10 (2009), pp. 1345–1359.
- [67] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. “A comprehensive survey on transfer learning”. In: *Proceedings of the IEEE* 109.1 (2020), pp. 43–76.
- [68] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. “How transferable are features in deep neural networks?” In: *Advances in neural information processing systems* 27 (2014).
- [69] Andrew M Leifer, Christopher Fang-Yen, Marc Gershow, Mark J Alkema, and Aravinthan DT Samuel. “Optogenetic manipulation of neural activity in freely moving *Caenorhabditis elegans*”. In: *Nature methods* 8.2 (2011), pp. 147–152.
- [70] Guillaume Aubry and Hang Lu. “A perspective on optical developments in microfluidic platforms for *Caenorhabditis elegans* research”. In: *Biomicrofluidics* 8.1 (2014).
- [71] Markus Reischl, Mazin Jouda, Neil MacKinnon, Erwin Fuhrer, Natalia Bakhtina, Andreas Bartschat, Ralf Mikut, and Jan G. Korvink. “Motion prediction enables simulated MR-imaging of freely moving model organisms”. In: *PLoS Computational Biology* 15.12 (2019), pp. 1–16.

-
- [72] Hart AG, Bowtell RW, Wenseleers T, and Ratnieks FLW. “Magnetic resonance imaging in entomology: a critical review”. In: *Journal of Insect Science* 3.1 (2003), p. 5.
- [73] Persian Poet Gal. *A diagram of the anatomy of the Coccinellidae (ladybug)*. Wikimedia Commons, CC BY 3.0. 2007. URL: [https://commons.wikimedia.org/wiki/File:Coccinellidae_\(Ladybug\)_Anatomy.svg](https://commons.wikimedia.org/wiki/File:Coccinellidae_(Ladybug)_Anatomy.svg).
- [74] Alexander Mathis, Pranav Mamidanna, Kevin M. Cury, Taiga Abe, Venkatesh N. Murthy, Mackenzie Weygandt Mathis, and Matthias Bethge. “DeepLabCut: markerless pose estimation of user-defined body parts with deep learning”. In: *Nature Neuroscience* 21.9 (2018), pp. 1281–1289.
- [75] Glenn Jocher, Ayush Chaurasia, and Jing Qiu. *Ultralytics YOLOv8*. Version 8.0.0. 2023. URL: <https://github.com/ultralytics/ultralytics>.
- [76] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. *Microsoft COCO: Common Objects in Context*. 2015. arXiv: 1405.0312 [cs.CV].
- [77] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. “You only look once: Unified, real-time object detection”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 779–788.
- [78] Julian Maclaren, Rainer Boegle, Michael Herbst, Jürgen Hennig, and Maxim Zaitsev. “Head pose prediction for prospectively-corrected EPI during rapid subject motion”. In: *Proceedings of the 18th Annual Meeting of ISMRM, Stockholm, Sweden*. 2010, p. 5031.
- [79] Julian Maclaren, Kuan J Lee, Chaiya Luengviriyaya, Oliver Speck, and Maxim Zaitsev. “Combined prospective and retrospective motion correction to relax navigator requirements”. In: *Magnetic Resonance in Medicine* 65.6 (2011), pp. 1724–1732.
- [80] Rudolph Emil Kalman. “A New Approach to Linear Filtering and Prediction Problems”. In: *Transactions of the ASME—Journal of Basic Engineering* 82.Series D (1960), pp. 35–45.
- [81] G. Bradski. “The OpenCV Library”. In: *Dr. Dobb’s Journal of Software Tools* (2000).
- [82] Gary A Kane, Gonçalo Lopes, Jonny L Saunders, Alexander Mathis, and Mackenzie W Mathis. “Real-time, low-latency closed-loop feedback using markerless posture tracking”. In: *Elife* 9 (2020), e61909.
- [83] Nathan White, Cooper Roddey, Ajit Shankaranarayanan, Eric Han, Dan Rettmann, Juan Santos, Josh Kuperman, and Anders Dale. “PROMO: real-time prospective motion correction in MRI using image-based tracking”. In: *Magnetic Resonance in Medicine* 63.1 (2010), pp. 91–105.

- [84] Julian Maclaren, Oliver Speck, Jürgen Hennig, and MA Zaitsev. “Kalman filtering framework for prospective motion correction”. In: *Proceedings of the 17th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Honolulu, Hawaii*. 2009, p. 4602.
- [85] Xinzhou Li, Yu-Hsiu Lee, Samantha Mikael, James Simonelli, Tsu-Chin Tsao, and Holden H Wu. “Respiratory motion prediction using fusion-based multi-rate Kalman filtering and real-time golden-angle radial MRI”. In: *IEEE Transactions on Biomedical Engineering* 67.6 (2019), pp. 1727–1738.
- [86] Daniel S Weller, Douglas C Noll, and Jeffrey A Fessler. “Prospective motion correction for functional MRI using sparsity and Kalman filtering”. In: *Wavelets and Sparsity XV*. Vol. 8858. SPIE. 2013, pp. 567–576.
- [87] Matthias C Wapler, Frederik Testud, Patrick Hucker, Jochen Leupold, Dominik von Elverfeldt, Maxim Zaitsev, and Ulrike Wallrabe. “MR-compatible optical microscope for in-situ dual-mode MR-optical microscopy”. In: *PLoS One* 16.5 (2021), e0250903.
- [88] Terry W Belke and Jason P Wagner. “The reinforcing property and the rewarding aftereffect of wheel running in rats: a combination of two paradigms”. In: *Behavioural processes* 68.2 (2005), pp. 165–172.
- [89] Fritz-Olaf Lehmann. “Matching spiracle opening to metabolic need during flight in *Drosophila*”. In: *Science* 294.5548 (2001), pp. 1926–1929.
- [90] John O Holloszy and Frank W Booth. “Biochemical adaptations to endurance exercise in muscle”. In: *Annual review of physiology* 38.1 (1976), pp. 273–291.
- [91] Frank Loesche and Michael B Reiser. “An inexpensive, high-precision, modular spherical treadmill setup optimized for *Drosophila* experiments”. In: *Frontiers in Behavioral Neuroscience* 15 (2021), p. 689573.
- [92] Mihaela Jekic, Eric L Foster, Michelle R Ballinger, Subha V Raman, and Orlando P Simonetti. “Cardiac function and myocardial perfusion immediately following maximal treadmill exercise inside the MRI room”. In: *Journal of Cardiovascular Magnetic Resonance* 10.1 (2008), p. 3.
- [93] Eric L Foster, John W Arnold, Mihaela Jekic, Jacob A Bender, Vijay Balasubramanian, Paaladinesh Thavendiranathan, Jennifer A Dickerson, Subha V Raman, and Orlando P Simonetti. “MR-compatible treadmill for exercise stress cardiac magnetic resonance imaging”. In: *Magnetic resonance in medicine* 67.3 (2012), pp. 880–889.
- [94] Terufumi Fujiwara, Margarida Brotas, and M Eugenia Chiappe. “Walking strides direct rapid and flexible recruitment of visual circuits for course control in *Drosophila*”. In: *Neuron* 110.13 (2022), pp. 2124–2138.

-
- [95] Hansjürgen Dahmen, Verena L Wahl, Sarah E Pfeffer, Hanspeter A Mallot, and Matthias Wittlinger. “Naturalistic path integration of Cataglyphis desert ants on an air-cushioned lightweight spherical treadmill”. In: *Journal of Experimental Biology* 220.4 (2017), pp. 634–644.
- [96] David I Hoult and RE Richards. “The signal-to-noise ratio of the nuclear magnetic resonance experiment”. In: *Journal of Magnetic Resonance (1969)* 24.1 (1976), pp. 71–85.
- [97] WA Edelstein, GH Glover, CJ Hardy, and RW Redington. “The intrinsic signal-to-noise ratio in NMR imaging”. In: *Magnetic resonance in medicine* 3.4 (1986), pp. 604–618.
- [98] Yi Chen, Patricia Pais-Roldan, Xuming Chen, Michael H Frosz, and Xin Yu. “MRI-guided robotic arm drives optogenetic fMRI with concurrent Ca²⁺ recording”. In: *Nature Communications* 10.1 (2019), p. 2536.
- [99] Matthias C Wapler, Jochen Leupold, Iulius Dragonu, Dominik von Elverfeld, Maxim Zaitsev, and Ulrike Wallrabe. “Magnetic properties of materials for MR engineering, micro-MR and beyond”. In: *Journal of Magnetic Resonance* 242 (2014), pp. 233–242.
- [100] Karl-Heinz Herrmann, Clemens Gärtner, Daniel Güllmar, Martin Krämer, and Jürgen R Reichenbach. “3D printing of MRI compatible components: why every MRI research group should have a low-budget 3D printer”. In: *Medical engineering & physics* 36.10 (2014), pp. 1373–1380.
- [101] Murat Aksoy, Christoph Forman, Matus Straka, Stefan Skare, Samantha Holdsworth, Joachim Hornegger, and Roland Bammer. “Real-time optical motion correction for diffusion tensor imaging”. In: *Magnetic resonance in medicine* 66.2 (2011), pp. 366–378.
- [102] VM Runge, JA Clanton, CL Partain, and AE James Jr. “Respiratory gating in magnetic resonance imaging at 0.5 Tesla”. In: *Radiology* 151.2 (1984), pp. 521–523.
- [103] Richard L Ehman, MT McNamara, M Pallack, H Hricak, and CB Higgins. “Magnetic resonance imaging with respiratory gating: techniques and advantages”. In: *American journal of Roentgenology* 143.6 (1984), pp. 1175–1182.
- [104] Peter Lanzer, C Barta, EH Botvinick, HU Wiesendanger, G Modin, and CB Higgins. “ECG-synchronized cardiac MR imaging: method and evaluation”. In: *Radiology* 155.3 (1985), pp. 681–686.
- [105] M.W. Spong, S. Hutchinson, and M. Vidyasagar. *Robot modeling and control*. Hoboken, NJ: Wiley, 2005.
- [106] Talmo D Pereira, Diego E Aldarondo, Lindsay Willmore, Mikhail Kislin, Samuel S-H Wang, Mala Murthy, and Joshua W Shaevitz. “Fast animal pose estimation using deep neural networks”. In: *Nature methods* 16.1 (2019), pp. 117–125.

- [107] Edward Brian Welch, Armando Manduca, Roger C Grimm, Heidi A Ward, and Clifford R Jack Jr. “Spherical navigator echoes for full 3D rigid body motion measurement in MRI”. In: *Magnetic resonance in medicine* 47.1 (2002), pp. 32–41.
- [108] M Dylan Tisdall, Aaron T Hess, Martin Reuter, Ernesta M Meintjes, Bruce Fischl, and André JW van der Kouwe. “Volumetric navigators for prospective motion correction and selective reacquisition in neuroanatomical MRI”. In: *Magnetic resonance in medicine* 68.2 (2012), pp. 389–399.
- [109] Frank Ong, Xucheng Zhu, Joseph Y Cheng, Kevin M Johnson, Peder EZ Larson, Shreyas S Vasanawala, and Michael Lustig. “Extreme MRI: Large-scale volumetric dynamic imaging from continuous non-gated acquisitions”. In: *Magnetic resonance in medicine* 84.4 (2020), pp. 1763–1780.
- [110] Sajan Goud Lingala, Brad P Sutton, Marc E Miquel, and Krishna S Nayak. “Recommendations for real-time speech MRI”. In: *Journal of Magnetic Resonance Imaging* 43.1 (2016), pp. 28–44.
- [111] Michael Burdumy, Louisa Traser, Fabian Burk, Bernhard Richter, Matthias Echter-nach, Jan G Korvink, Jürgen Hennig, and Maxim Zaitsev. “One-second MRI of a three-dimensional vocal tract to measure dynamic articulator modifications”. In: *Journal of Magnetic Resonance Imaging* 46.1 (2017), pp. 94–101.
- [112] Krishna S Nayak, Yongwan Lim, Adrienne E Campbell-Washburn, and Jennifer Steeden. “Real-time magnetic resonance imaging”. In: *Journal of Magnetic Resonance Imaging* 55.1 (2022), pp. 81–99.
- [113] Julian Maclaren, Oliver Speck, Daniel Stucht, Peter Schulze, Jürgen Hennig, and Maxim Zaitsev. “Navigator accuracy requirements for prospective motion correc-tion”. In: *Magnetic Resonance in Medicine* 63.1 (2010), pp. 162–170.
- [114] Benjamin Zahneisen, Brian Keating, and Thomas Ernst. “Propagation of calibra-tion errors in prospective motion correction using external tracking”. In: *Magnetic Resonance in Medicine* 72.2 (2014), pp. 381–388.
- [115] Yuchou Chang, Zhiqiang Li, Gulfam Saju, Hui Mao, and Tianming Liu. “Deep learning-based rigid motion correction for magnetic resonance imaging: a survey”. In: *Meta-Radiology* (2023), p. 100001.
- [116] Patricia M. Johnson and Maria Drangova. “Conditional generative adversarial net-work for 3D rigid-body motion correction in MRI”. In: *Magnetic Resonance in Medicine* 82.3 (2019), pp. 901–910.
- [117] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. “Decaf: A deep convolutional activation feature for generic visual recognition”. In: *International conference on machine learning*. PMLR. 2014, pp. 647–655.

-
- [118] Roberto Souza, Oeslle Lucena, Julia Garrafa, David Gobbi, Marina Saluzzi, Simone Appenzeller, Letícia Rittner, Richard Frayne, and Roberto Lotufo. “An open, multi-vendor, multi-field-strength brain MR dataset and analysis of publicly available skull stripping methods agreement”. In: *NeuroImage* 170 (2018), pp. 482–494.
- [119] Mohammed A. Al-masni, Seul Lee, Jaeuk Yi, Sewook Kim, Sung Min Cho, Young Hun Choi, and Dong Hyun Kim. “Stacked U-Nets with self-assisted priors towards robust correction of rigid motion artifact in brain MRI”. In: *NeuroImage* 259.March (2022), p. 119411.
- [120] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. “Image quality assessment: from error visibility to structural similarity”. In: *IEEE transactions on image processing* 13.4 (2004), pp. 600–612.
- [121] Diederik P Kingma and Jimmy Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
- [122] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. “Pytorch: An imperative style, high-performance deep learning library”. In: *Advances in neural information processing systems* 32 (2019).
- [123] Melanie Dohmen, Mark A Klemens, Ivo M Baltruschat, Tuan Truong, and Matthias Lenga. “Similarity and quality metrics for MR image-to-image translation”. In: *Scientific Reports* 15.1 (2025), p. 3853.
- [124] Elisa Marchetto, Hannah Eichhorn, Daniel Gallichan, Julia A Schnabel, and Melanie Ganz. “Agreement of image quality metrics with radiological evaluation in the presence of motion artifacts”. In: *Magnetic Resonance Materials in Physics, Biology and Medicine* (2025), pp. 1–12.
- [125] Kieran P McGee, Armando Manduca, Joel P Felmlee, Stephen J Riederer, and Richard L Ehman. “Image metric-based correction (autocorrection) of motion effects: analysis of image metrics”. In: *Journal of Magnetic Resonance Imaging* 11.2 (2000), pp. 174–181.
- [126] Stefan Van der Walt, Johannes L Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D Warner, Neil Yager, Emmanuelle Gouillart, and Tony Yu. “scikit-image: image processing in Python”. In: *PeerJ* 2 (2014), e453.
- [127] Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. “SciPy 1.0: fundamental algorithms for scientific computing in Python”. In: *Nature methods* 17.3 (2020), pp. 261–272.
- [128] Christopher J Hardy, Robert D Darrow, Manojkumar Saranathan, Randy O Giaquinto, Yudong Zhu, Charles L Dumoulin, and Paul A Bottomley. “Large field-of-view real-time MRI with a 32-channel system”. In: *Magnetic Resonance in Medicine* 52.4 (2004), pp. 878–884.

- [129] Jesse Hamilton, Dominique Franson, and Nicole Seiberlich. “Recent advances in parallel imaging for MRI”. In: *Progress in nuclear magnetic resonance spectroscopy* 101 (2017), pp. 71–95.
- [130] Shanshan Wang, Zhenghang Su, Leslie Ying, Xi Peng, Shun Zhu, Feng Liang, Dagan Feng, and Dong Liang. “Accelerating magnetic resonance imaging via deep learning”. In: *2016 IEEE 13th international symposium on biomedical imaging (ISBI)*. IEEE. 2016, pp. 514–517.
- [131] Reinhard Heckel, Mathews Jacob, Akshay Chaudhari, Or Perlman, and Efrat Shimron. “Deep learning for accelerated and robust MRI reconstruction”. In: *Magnetic Resonance Materials in Physics, Biology and Medicine* 37.3 (2024), pp. 335–368.
- [132] Michael Markl and Jürgen Hennig. “Phase contrast MRI with improved temporal resolution by view sharing: k-space related velocity mapping properties”. In: *Magnetic Resonance Imaging* 19.5 (2001), pp. 669–676.
- [133] Ergys Subashi, Li Feng, Yilin Liu, Scott Robertson, Paul Segars, Bastiaan Driehuys, Christopher R Kelsey, Fang-Fang Yin, Ricardo Otazo, and Jing Cai. “View-sharing for 4D magnetic resonance imaging with randomized projection-encoding enables improvements of respiratory motion imaging for treatment planning in abdominothoracic radiotherapy”. In: *Physics and imaging in radiation oncology* 25 (2023), p. 100409.
- [134] Luminita A Tudorica, Karen Y Oh, Nicole Roy, Mark D Kettler, Yiyi Chen, Stephanie L Hemmingson, Aneela Afzal, John W Grinstead, Gerhard Laub, Xin Li, et al. “A feasible high spatiotemporal resolution breast DCE-MRI protocol for clinical settings”. In: *Magnetic Resonance Imaging* 30.9 (2012), pp. 1257–1267.
- [135] Frank R Korosec, Richard Frayne, Thomas M Grist, and Charles A Mistretta. “Time-resolved contrast-enhanced 3D MR angiography”. In: *Magnetic resonance in medicine* 36.3 (1996), pp. 345–351.
- [136] Evan Levine, Bruce Daniel, Shreyas Vasanawala, Brian Hargreaves, and Manojkumar Saranathan. “3D Cartesian MRI with compressed sensing and variable view sharing using complementary poisson-disc sampling”. In: *Magnetic resonance in medicine* 77.5 (2017), pp. 1774–1785.
- [137] Stanislas Rapacchi, Yutaka Natsuaki, Adam Plotnik, Simon Gabriel, Gerhard Laub, J Paul Finn, and Peng Hu. “Reducing view-sharing using compressed sensing in time-resolved contrast-enhanced magnetic resonance angiography”. In: *Magnetic resonance in medicine* 74.2 (2015), pp. 474–481.
- [138] Qiaoying Huang, Yikun Xian, Dong Yang, Hui Qu, Jingru Yi, Pengxiang Wu, and Dimitris N Metaxas. “Dynamic MRI reconstruction with end-to-end motion-guided network”. In: *Medical Image Analysis* 68 (2021), p. 101901.

-
- [139] Uygur Sumbul, Juan M Santos, and John M Pauly. “A practical acceleration algorithm for real-time imaging”. In: *IEEE transactions on medical imaging* 28.12 (2009), pp. 2042–2051.
- [140] Chenlu Qiu, Wei Lu, and Namrata Vaswani. “Real-time dynamic MR image reconstruction using Kalman filtered compressed sensing”. In: *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE. 2009, pp. 393–396.
- [141] Xue Feng, Michael Salerno, Christopher M Kramer, and Craig H Meyer. “Kalman filter techniques for accelerated Cartesian dynamic cardiac imaging”. In: *Magnetic Resonance in Medicine* 69.5 (2013), pp. 1346–1356.
- [142] Burhaneddin Yaman, Seyed Amir Hossein Hosseini, Steen Moeller, Jutta Ellermann, Kâmil Uğurbil, and Mehmet Akçakaya. “Self-supervised physics-based deep learning MRI reconstruction without fully-sampled data”. In: *2020 IEEE 17th international symposium on biomedical imaging (ISBI)*. IEEE. 2020, pp. 921–925.
- [143] Kerstin Hammernik, Thomas Küstner, Burhaneddin Yaman, Zhengnan Huang, Daniel Rueckert, Florian Knoll, and Mehmet Akçakaya. “Physics-driven deep learning for computational magnetic resonance imaging: Combining physics and machine learning for improved medical imaging”. In: *IEEE signal processing magazine* 40.1 (2023), pp. 98–114.
- [144] Li Feng, Robert Grimm, Kai Tobias Block, Hersh Chandarana, Sungheon Kim, Jian Xu, Leon Axel, Daniel K Sodickson, and Ricardo Otazo. “Golden-angle radial sparse parallel MRI: combination of compressed sensing, parallel imaging, and golden-angle radial sampling for fast and flexible dynamic volumetric MRI”. In: *Magnetic resonance in medicine* 72.3 (2014), pp. 707–717.
- [145] Bernard Lanz, Alireza Abaei, Olivier Braissant, In-Young Choi, Cristina Cudalbu, Pierre-Gilles Henry, Rolf Gruetter, Firat Kara, Kejal Kantarci, Phil Lee, et al. “Magnetic resonance spectroscopy in the rodent brain: experts’ consensus recommendations”. In: *NMR in Biomedicine* 34.5 (2021), e4325.
- [146] Samira Kabli, Herman P Spaink, Huub JM De Groot, and A Alia. “In vivo metabolite profile of adult zebrafish brain obtained by high-resolution localized magnetic resonance spectroscopy”. In: *Journal of Magnetic Resonance Imaging* 29.2 (2009), pp. 275–281.
- [147] Franz Schilling, Kai Dworschak, Reinhard Schopf, Ralph Kühn, Steffen J Glaser, and Axel Haase. “Non-invasive lipid measurement in living insects using NMR microscopy”. In: *Journal of Experimental Biology* 215.18 (2012), pp. 3137–3141.
- [148] Paul Glover and Sir Peter Mansfield. “Limits to magnetic resonance microscopy”. In: *Reports on progress in physics* 65.10 (2002), pp. 1489–1511.

Publications

Journal articles

- [1] **Ajmal Chenakkara**, Mazin Jouda, Ulrike Wallrabe, Jan G Korvink. “Motion compensated magnetic resonance imaging of an active sun beetle using an in situ treadmill”. In: *Scientific Reports* 15.1 (2025), p. 40340.
- [2] **Ajmal Chenakkara**, Mazin Jouda, Ulrike Wallrabe, Jan G Korvink. “Residual motion artifact removal enables dynamic μ MRI of a behaving *Pachnoda marginata*”. In: *Journal of Magnetic Resonance* (2025), p. 107954.
- [3] **Ajmal Chenakkara**, Mazin Jouda, Ulrike Wallrabe, Jan G Korvink. “*In situ* time-resolved motion of a tethered *Pachnoda marginata*, AI-correlated using μ MRI and optical imaging”. In: *Journal of Magnetic Resonance* 389 (2026), p. 108089.
- [4] Dilara Faderl, **Ajmal Chenakkara**, Mazin Jouda, Neil MacKinnon, Alvar D Gossert, Jan G Korvink. “Accelerated Screening of Protein–Ligand Interactions via Parallel T2-Weighted 19F-MRI”. In: *Analytical Chemistry* 96.24 (2024), pp. 9859–9865.
- [5] Moritz Becker, Yen-Tse Cheng, Achim Voigt, **Ajmal Chenakkara**, Mengjia He, Sören Lehmkuhl, Mazin Jouda, Jan G Korvink. “Artificial intelligence-driven shimming for parallel high field nuclear magnetic resonance”. In: *Scientific Reports* 13.1 (2023), p. 17983.

Conference contributions

- [1] **Ajmal Chenakkara**, Mazin Jouda, Ulrike Wallrabe, Jan G. Korvink. “Towards in-vivo NMR metabolomics of free organism”. In: *43rd FGMR Annual Discussion Meeting, Karlsruhe*. Poster presented. 2022.
- [2] **Ajmal Chenakkara**, Mazin Jouda, Ulrike Wallrabe, Jan G. Korvink. “Treadmill MRI – a platform for motion compensated imaging of behaving organisms”. In: *Proc. 66th Experimental NMR Conference (ENC-ISMAR), Asilomar Conference Center, Pacific Grove, CA*. Poster accepted (presented by JG Korvink). 2025.

Appendix

A Data availability

The relevant code and data supporting this thesis are publicly available at <https://github.com/chenakkara>.

B Imaging parameters

FLASH Parameters: TE/TR = 4 ms/200 ms, slice thickness = 1 mm, matrix = 256×256 , FOV = 30×30 mm, in-plane resolution = 117×117 μm , FA = 30° , BW = 50 kHz, NEX = 1, acquisition time = 51.2 s.

RARE Parameters: TE/TR = 5.24 ms/1000 ms, slice thickness = 1 mm, matrix = 256×256 , FOV = 30×30 mm, in-plane resolution = 117×117 μm , RARE factor = 1, NEX = 1, acquisition time = 256 s.

C DL motion correction: simulated dataset distribution




Figure C.1: Deep learning-based residual motion correction: training dataset distribution **(a)** Pretraining brain MRI translation motion simulation. **(b)** Pretraining brain MRI rotation motion simulation. **(c)** Finetuning insect MRI translation motion simulation. **(d)** Finetuning insect MRI rotation motion simulation.

List of Figures

2.1	Nuclear spin magnetization and the energy diagram. (a) Nuclear spins at thermal equilibrium in the absence of an external magnetic field. (b) In the presence of a static magnetic field B_0 , a slight population imbalance results in a net longitudinal magnetization M_0	9
2.2	RF excitation and relaxation of net nuclear magnetization in the rotating frame of reference. (a) A 90° RF excitation rotates the net longitudinal magnetization M_0 into the transverse (x – y) plane. (b–d) Magnetization relaxing back to equilibrium, with (e) showing the FID signal. (f) Recovery of longitudinal magnetization, characterized by the longitudinal relaxation time T_1 . (g) Decay of transverse magnetization, characterized by the transverse relaxation time T_2	11
2.3	Spin echo refocusing.	13
2.4	Gradient echo refocusing.	13
2.5	Spatial encoding in MRI. (a) Slice selection. (b) Phase and frequency encoding for k -space sampling, and subsequent image reconstruction. (Brain MRI from [25]).	16
2.6	Spin echo pulse sequence.	20
2.7	RARE pulse sequence.	20
2.8	Gradient echo pulse sequence.	21
2.9	Sparse optical flow tracking: keypoint moves from position 1 to position 2 under brightness constancy assumption.	28
2.10	Artificial neural networks. (a) Single-layer perceptron model. (b) Multilayer perceptron with two hidden layers.	31
2.11	CNN processes input MRI through convolutional layers with activation function to extract hierarchical features, interleaved pooling to reduce resolution, and final fully connected layers for task specific output.	32
2.12	Transfer learning pipeline.	34
3.1	Prospective motion correction in μ MRI: The imaging volume dynamically tracks the subject's motion, enabling MR imaging of freely moving specimens in an <i>in silico</i> study. (Adopted from Reischl et al. [71]).	38
3.2	External morphology of a ladybug. (Modified from [73]).	41

3.3	Deep learning–based motion tracking. (a) Schematic illustration of the <i>ex situ</i> insect tracking setup. (b) Schematic of the deep learning–based tracking pipeline, showing fine-tuning of a pre-trained model using custom insect data.	43
3.4	<i>Ex situ</i> tracking validation on test dataset ($n = 100$). (a) Model-estimated (x, y) position trajectory of the insect. (b) Histogram of the Euclidean distance error between model-predicted keypoints and manual annotations, showing a mean error of 7 ± 3 pixels. (c) Pose (θ) estimation results obtained from the model. (d) Representative test images illustrating keypoint detection with both small and large localization errors (green crosses: manual annotations; red crosses: model detections).	44
3.5	Deep learning–based tracking of ladybug body pose in the <i>ex situ</i> setup. Sampled video frames with timestamps show the motion trajectory.	45
3.6	Kalman filter–based state estimation under temporal subsampling. (a) Comparison of the insect’s X-position trajectory obtained from the deep learning–based tracker and the corresponding Kalman filter state estimates for an effective temporal interval of 0.27 s. (b) MAE between the detected positions and Kalman estimates as a function of the temporal interval. Error bars indicate the standard error of the mean.	46
3.7	Treadmill-based experimental setups in the literature. (a) Head-fixed <i>Drosophila</i> walking on a spherical treadmill during whole-cell patch-clamp recording with optogenetic stimulation [94]. (b) An air-cushioned spherical treadmill used for navigation studies in <i>Cataglyphis</i> desert ants [95].	49
3.8	External morphology of a sun beetle.	50
3.9	Tethered sun beetle on a spherical treadmill.	51
3.10	Anatomical reference image of <i>Pachnoda marginata</i> . (a) Sagittal MRI showing the abdomen. (b) Axial MRI showing the thorax. (c) MRI slice position. (Modified from [17].)	52
3.11	Sagittal MRI of sun beetle using (a) FLASH, and (b) RARE (anatomical reference).	54
4.1	Design and development of treadmill setup. (a) Integrated treadmill with optical imaging system inside 15.2 T magnet. (b) Spherical treadmill setup inside the ^1H RF coil. (Modified from [17].)	60
4.2	Design and development of treadmill setup. (a–d) CAD rendering of the treadmill assembly. Treadmill with active beetle (e) <i>in situ</i> and (f) <i>ex situ</i> ^1H RF volume coil. (Adapted from [17].)	60
4.3	CAD renderings of treadmill design iterations. (a) Initial design with a single air vent. (b) Revised design with six air vents for improved airflow stability. (c) Wireframe model of the six–air-vent design highlighting internal geometry.	61

4.4	Design and development of the <i>in situ</i> MR-compatible optical imaging system. (a) Perspective view. (b) Side view. (c) Webcam and housing. (d) Exploded view. (Adapted from [17].)	62
4.5	Schematic diagram of the <i>in situ</i> treadmill with an integrated computer vision system, a platform for performing MRI of behaving insects. (Adopted from [17]).	63
4.6	ΔB_0 distribution with and without the camera system, and with camera + active shim for a water sample. (a) Central axial slice. (b) Coronal slice [z-axis]. (Adapted from [17].)	64
4.7	Animal handling and tethering procedure. (a) The beetle is placed in a small ventilated plastic container and immobilized via cold-anesthesia. (b) Spherical treadmill setup used for tethered locomotion. (c) The removable PLA-based tether is prepared, and a single drop of UV-curable adhesive is applied at the designated attachment site. (d) The cold-anesthetized beetle is tethered following adhesive curing. (e) The tethered beetle is transferred onto the spherical treadmill. (f) The treadmill with the tethered beetle is slid into the RF volume coil. (g) Final positioning of the setup for MRI experiments.	65
4.8	Effect of locomotion-induced motion artifacts on MRI reconstruction during tethered experiments. (a) Reference image acquired while the insect was alive but in a quiescent state, resulting in minimal motion. (b) Motion-corrupted image showing artifacts predominantly arising from abdominal and head movements.	67
4.9	Schematic representation of the multi-rigid approximation of tethered insect body motion, modeled as three articulated links connected by two joints (defined in the scanner coordinate system).	69
4.10	<i>In situ</i> computer vision system for real-time prospective gating. (a) Schematic of computer vision-based prospective gating. (b) Abdominal motion phases of the insect. (Adopted from [17].)	73
4.11	Euclidean distance error between optical flow algorithm detection and manual annotation for 100 images. (a) Y-position trajectory. (b) Error histogram. Sample images showing (c) small and (d) large detection error. (Adapted from [17].)	74
4.12	Effect of tethered locomotion on axial thorax MRI. (a) Motion trajectory of the inactive insect abdomen. (b) Thorax MRI of the inactive insect. (c) Motion trajectory of the walking insect abdomen. (d) Thorax MRI of the active insect. (Adapted from [17].)	76

4.13	Prospective gating for motion-compensated sagittal MRI of a behaving insect on a treadmill. (a–d) Inactive condition: motion trajectory, sagittal MRI, magnified abdominal view, and corresponding intensity profile showing minimal motion artifacts. (e–h) Behaving insect without gating, demonstrating pronounced motion artifacts. (i–l) , (m–p) , and (q–t) Prospective gating iterations 1–3, respectively, showing motion trajectories, MRI images, magnified abdominal views, and intensity profiles, showing consistent reduction of motion artifacts. (Adapted from [17]).	78
5.1	Schematic of the two-stage MRI motion compensation workflow for the tethered insect on a treadmill, showing (a) initial retrospective gating, followed by (b) deep learning–based residual motion correction. (Modified from [18].)	83
5.2	Schematic representation of the <i>in situ</i> treadmill setup integrated with an MR-compatible computer vision system for motion estimation, enabling retrospective gating for motion-compensated MRI. (a) Computer vision system. (b) <i>In situ</i> treadmill. (c) Motion trajectory of the selected keypoint, with the phase encoding index of the MRI frames. (d) Phase encode reordering for position consistent MRI reconstruction. (Adopted from [18].)	85
5.3	Retrospective gating for motion compensated μ MRI. (a) X&Y-motion trajectory of the selected keypoint for the duration of MRI raw data acquisition, color-coding represents the different spatial clusters. (b) 2D X-Y plot of the same motion trajectory, with color-coding represents the grid-based clusters with the centroid defines the different position states for gating. (c) Standard deviation of distance from the centroid of the individual clusters. (d) Average magnitude image of 50 MRI frames showing the extent of motion-artifacts. (e) Sample motion-corrupted frame out of the reconstructed 50 frames without motion compensation. (f) Retrospectively gated MRI frame showing reduction in motion-artifacts. (g) Reference image with less motion artifacts. (Adopted from [18].)	87
5.4	Internal organ motion of a dormant insect. (a) Motion trajectory of the marked keypoint on the insect’s digestive system estimated using the optical flow algorithm 1 (Time axis shown at 5 $\times$ real-time). (b) MRI sagittal slice reference position. (c) Internal organ motion, phase 1. (d) Internal organ motion, phase 2. (e) Internal organ motion, phase 3. (Adapted from [18].)	89
5.5	Effect of acquisition time on motion sensitivity in single-frame μ MRI of the dormant insect. (a) Short acquisition time of 1.2 s captures the corrugated structure of the midgut with minimal motion artifacts. (b) Intermediate acquisition time of 32 s shows moderate motion artifacts. (c) Long acquisition time of 64 s results in pronounced motion artifacts. (Adapted from [18].)	90

5.6	Transfer learning for residual motion artifact correction. (a) Pre-training the model with the brain MRI dataset. (b) Fine-tuning the model with a mixed dataset of brain and insect MRI images. (Adapted from [18].)	91
5.7	Schematic diagram showing motion artifact simulation pipeline. (a) Motion-free MRI and k -space magnitude images. (b) k -space compilation based on image rotation and k -space translation. (c) Reconstructed MRI magnitude image showing motion artifacts. (Adapted from [18].)	93
5.8	Motion artifact simulation: brain MRI dataset. (a) Motion-free reference (Example 1). (b) Corresponding image corrupted by simulated translation motion. (c) Absolute difference between motion-free and motion-corrupted images, highlighting motion artifacts. (d) Translation motion trajectory used for simulation. (e) Motion-free reference (Example 2). (f) Corresponding image corrupted by simulated rotational motion. (g) Absolute difference between motion-free and motion-corrupted images. (h) Rotation motion trajectory used for simulation.	94
5.9	Motion artifact simulation: insect MRI dataset. (a) Motion-free reference (Example 1). (b) Corresponding image corrupted by simulated translation motion. (c) Absolute difference between motion-free and motion-corrupted images, highlighting motion artifacts. (d) Translation motion trajectory used for simulation. (e) Motion-free reference (Example 2). (f) Corresponding image corrupted by simulated rotational motion. (g) Absolute difference between motion-free and motion-corrupted images. (h) Rotation motion trajectory used for simulation.	95
5.10	Schematic of the U-Net [65]–based deep learning architecture used for image-domain motion artifact correction.	97
5.11	Deep learning motion correction on simulated dataset. (a) Simulated motion-corrupted test image, magnified view, absolute difference with the reference. (b) Motion-corrected image with the fine-tuned model, magnified view, absolute difference with the reference. (c) Motion-free reference for comparison, magnified view, absolute difference with the reference. (Adapted from [18].)	101
5.12	Quantitative evaluation of deep learning motion correction on simulated dataset. (a) Histogram of SSIM values for test images before and after correction. (b) Histogram of GE values for test images before and after correction. (Adapted from [18].)	102
5.13	Deep learning (DL) motion correction on experimental dataset. (a) Retrospectively gated image with residual motion artifacts. (b) DL motion-corrected image. (c) Non-gated image with less severe motion artifacts. (d) DL motion-corrected image. (e) Bar graph of GE values before and after DL correction. (Adapted from [18].)	103

6.1	Schematic representation of single coil variable density Cartesian undersampling MRI simulation. (a) Fully-sampled MRI. (b) Fully-sampled k -space. (c) Undersampling mask. (d) Undersampled k -space. (e) ZF MRI reconstruction. (Adapted from [19].)	108
6.2	DL correction of undersampling artifacts for the simulated dataset ($n = 150$). (a–c) ZF reconstructions at $2\times$, $3\times$, and $5\times$ acceleration, respectively. (d) Fully sampled ground truth image. (e–g) DL-corrected reconstructions for $2\times$, $3\times$, and $5\times$ acceleration. (Adapted from [19].)	110
6.3	Quantitative assessment of deep learning-based undersampling artifact correction across different acceleration factors. (a) Mean GE. (b) Mean SSIM. (Adapted from [19].)	111
6.4	Schematic representation of retrospective view-sharing with DL-corrected reconstruction. (a) Previous k -space. (b) Current k -space. (c) Next k -space. (d) View-shared MRI. (e) DL-corrected MRI. (Adapted from [19].)	112
6.5	DL-based correction of undersampling artifacts following VS reconstruction ($n = 140$ dynamic MR images). (a) Undersampled k -space. (b) ZF reconstruction. (c) DL-corrected ZF reconstruction. (d) Image–time profile along the blue line in (b) for ZF reconstruction. (e) Corresponding image–time profile for the DL-corrected reconstruction. (f) VS k -space. (g) VS reconstruction. (h) DL-corrected VS reconstruction. (i) Image–time profile for the VS reconstruction. (j) Corresponding image–time profile for the DL-corrected VS reconstruction. (Adapted from [19].)	114
6.6	Comparison of GE distributions. (a) GE distribution for ZF and DL-corrected reconstructions. (b) GE distribution for VS and DL-corrected reconstructions. (Adapted from [19].)	115
6.7	Dynamic <i>in vivo</i> correlated μ MRI and optical imaging of the <i>in situ</i> behaving insect. (a) Intensity–time profile extracted from the dynamic optical image sequence ($n = 140$). (b) Intensity–time profile extracted from the dynamic μ MRI sequence ($n = 140$). Representative dynamic frames illustrating different stride phases of the insect’s abdominal pitching during tethered locomotion, with corresponding frame indices and times (relative to the first displayed frame): (c) Optical image frames. (d) μ MRI frames (green line: reference position). (Adopted from [19].)	116
7.1	<i>In vivo</i> localized ^1H MRS of a tethered insect. (a) Insect on the <i>in situ</i> treadmill. (b) Anatomical reference indicating voxel placement within the thorax. (c) Magnitude spectrum highlighting water and lipid resonances. (d) Time-resolved stack plot illustrating spectral stability over a 50 min duration.	121

- C.1 Deep learning-based residual motion correction: training dataset distribution **(a)** Pretraining brain MRI translation motion simulation. **(b)** Pretraining brain MRI rotation motion simulation. **(c)** Finetuning insect MRI translation motion simulation. **(d)** Finetuning insect MRI rotation motion simulation. 144

List of Tables

- 6.1 GE and SSIM values for ZF and DL-corrected reconstructions of the simulated dataset at different acceleration factors (mean $\pm$ SD, $n = 150$). . . . 110

Acknowledgments

My heartfelt thanks to Prof. Jan Korvink for the opportunity to work on this exciting project and for being a wonderful advisor. His breadth of interests, knowledge, enthusiasm, and constant stream of exciting ideas inspired my own view of science as a fun and exploratory pursuit, and motivated me to explore some bold directions in my project. I truly believe that all the monthly meetings shaped me both personally and professionally. I would like to thank Prof. Ulrike Wallrabe for being my second advisor and for always offering warmth, guidance, and encouragement throughout my research journey.

I would like to thank Dr. Mazin Jouada for being a great supervisor, and I feel fortunate to have been part of his group at IMT. His office was always open to everyone, fostering welcoming and thoughtful discussions on many topics. I will especially cherish the brainstorming conversations we had over coffee at the bistro.

I would like to thank all the former and current members of Mazin's group; Moritz, Nourhan, Hossein, Mohammad, Anthony, Junlin, Alex, and Mehrdad, for making my time here such a pleasant experience. Collaborating academically with some of them also had a positive impact on my experience and reinforced my view of science as a collective pursuit.

I am grateful to all the former and current colleagues who supported me and made this long journey enjoyable. Special thanks to Dilara, Paco, Julia, Leon, Emil, Jing, Omar, Neil, Ajeya, and Shraddha. Thanks to Anthony, Han, and Max for being great office mates, and for all the interesting conversations. I am also thankful to all the IMT staff, and I apologize if I missed someone inadvertently.

I would like to thank my friends from CUSAT, Kochi; especially Tony, Sara, Bhavana, and Gopika for their support and companionship. I am also thankful to my childhood friends for their lifelong friendship and support.

I am deeply grateful and indebted to my parents and sisters for their unconditional support and encouragement throughout this journey. I am forever grateful to my wife for her patience, love, and understanding, and I am truly excited for all that lies ahead.

Finally, I would like to acknowledge that AI-based tools were used for language polishing and stylistic improvements.

