The Muon-Weighted Magnetic Field and Related Challenges in the Muon g-2 Experiment

by

Eva Kraegeloh

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Physics and Scientific Computing)
in the University of Michigan
2026

Doctoral Committee:

Professor Tim Chupp, Chair
Professor Myron Campbell
Professor Emanuel Gull
Professor Wolfgang Lorenzon
Professor Eric Michielssen

Eva Kraegeloh
krageloh@umich.edu
ORCID iD: 0000-0001-8512-7059

© Eva Kraegeloh 2026

DEDICATION

To Violet and Maddie,
you're going to run the world.

;

ACKNOWLEDGEMENTS

This thesis would not exist without the support of many people in both my professional and personal life. I want to start by thanking my advisor, Tim Chupp, for his patience and understanding, and unwavering belief in my competence as a physicist, even when I was questioning my life choices. I would not have crossed the finish line without his encouragement. I also owe a debt of gratitude to the entire Fermilab Muon g-2 collaboration, and a few people in particular: Saskia Charity and Simon Corrodi for always delivering the data I needed, and answering my barrage of questions whenever I tried to make sense of something in my head; Brynn MacCoy for trusting me to continue and expand on her work on the muon weighting, and commiserating about the unique challenges of being a woman in a physics PhD program with me; Joe Price and James Mott for their help with anything tracker-related; and especially Midhat Farooq and Alec Tewsley-Booth, without whom I would have felt lonely and lost in a big collaboration when I didn't know anyone. Similarly, my office mates David Aguillard and Henry Sottrel always manage to make me feel like the long commute to campus is worth it, and I am so grateful for the support they've shown me in this quest to finish my PhD.

I could not have completed this thesis without the people in my professional life, but neither could I have done it without my incredible friends and chosen family. Ciara, Haley, Christine, Antasia, Travis, Christie, and the rest of the polycule have encouraged me, listened to me vent, fed me, and so much more. I would not be here without them, without my community. It truly takes a village. And last but not least I want to thank my parents for being there whenever I need them, even though we are continents apart. I realize that I am very fortunate to know, deep in my soul, that they will have my back no matter what.

Funding acknowledgment

This work was supported in part by NSF Grants PHY-1812314, PHY-2110988, and PHY-2412368.

TABLE OF CONTENTS

DEDICATION ii

ACKNOWLEDGEMENTS iii

LIST OF FIGURES vii

LIST OF TABLES x

LIST OF APPENDICES xii

LIST OF ACRONYMS xiii

ABSTRACT xv

CHAPTER

1 Introduction 1

1.1 A brief history of Muon g-2 1
1.2 Fermilab Muon g-2 timeline 2
1.3 Contents of this work 2

2 Theoretical Background 5

2.1 Magnetic dipole moments 5
2.2 Standard Model prediction 6
2.2.1 Quantum electrodynamic contribution 8
2.2.2 Quantum chromodynamic contribution 9
2.2.3 Electroweak contribution 19
2.3 Current status of theory 20

3 Experimental Overview 24

3.1 Experimental setup 24
3.1.1 Muon production and delivery 24
3.1.2 Muon storage 25
3.2 Methodology 28
3.2.1 Measurement principle 28
3.2.2 Measurement of $\omega_{\mathbf{a}}$ 30
3.2.3 Measurement of $\omega_{\mathbf{p}}$ 36

4 Spatial Magnetic Field Fitting . . . 40
4.1 2D parameterization . . . 41
4.2 Linear regression fit . . . 46
4.3 3D toroidal parameterization . . . 48
4.3.1 Introduction to toroidal coordinates . . . 48
4.3.2 Solution to Laplace's equation . . . 50
4.3.3 Direct comparison to 2D polar solution . . . 52
4.4 Magnetic field fitting in toroidal coordinates . . . 57
4.4.1 Fit performance considerations . . . 57
4.4.2 Free parameters . . . 68
4.4.3 Simulated toroidal fields . . . 70
4.4.4 Toroidal fit results . . . 74
4.5 Validation of 2D multipole fits . . . 80
4.5.1 Effect of non-vertical field components . . . 81
4.5.2 Difference between 2D and 3D fits . . . 82
4.5.3 Azimuthal position interpolation . . . 83
5 Muon Weighted Field Average . . . 85
5.1 Magnetic field binning . . . 92
5.2 Beam data binning . . . 95
5.2.1 Run group formation . . . 95
5.2.2 Non-standard run groups . . . 97
5.2.3 Beam distribution reconstruction . . . 100
5.3 Beam profile extrapolation . . . 101
5.3.1 Momentum deviation . . . 105
5.3.2 Betatron and dispersion functions . . . 107
5.3.3 Radial closed orbit distortion . . . 108
5.4 Run-4/5/6 muon weighting results . . . 111
5.5 Run-4/5/6 muon weighting uncertainties . . . 113
5.5.1 Tracking uncertainties . . . 119
5.5.2 Beam dynamics inputs . . . 121
5.5.3 Calorimeter acceptance . . . 129
5.5.4 Field-related uncertainties . . . 130
5.5.5 Summary . . . 133
6 Learning-Based Track Reconstruction . . . 135
6.1 Analytic track reconstruction algorithm . . . 136
6.1.1 Hit pattern recognition . . . 136
6.1.2 Track finding . . . 138
6.2 A brief neural network primer . . . 140
6.3 Machine learning algorithm for Muon g-2 tracking . . . 142
6.4 Toy model . . . 143
6.4.1 Idealized detector setup . . . 145
6.4.2 Analytic likelihood . . . 146
6.4.3 Neural net . . . 147

6.4.4 Minimization 150
6.5 Towards a full-scale model 155
7 Summary 159
7.1 Conclusions 160
7.2 Outlook 160

APPENDICES 162

BIBLIOGRAPHY 197

LIST OF FIGURES

FIGURE

1.1 Accumulated data collected during the Fermilab Muon g-2 experiment 4

2.1 Feynman diagrams of a muon interacting with a magnetic field 7
2.2 Hadronic sector Feynman diagrams 10
2.3 Results for $a_\mu^{\pi^+\pi^-,\mathrm{LO}}$ in the energy region $0.6\,\mathrm{GeV} < \sqrt{s} < 0.88\,\mathrm{GeV}$ 12
2.4 Illustration of hadronic vacuum polarization (HVP) window weights 13
2.5 Quark flavor contributions to $a_\mu^{\mathrm{HVP,LO}}$ 15
2.6 Window contribution to $a_\mu^{\mathrm{HVP,LO}}$ 15
2.7 $a_\mu^{\mathrm{HLbL},\pi^0\text{-pole}}$ and $a_\mu^{\mathrm{HLbL,LO}}$ 18
2.8 Electroweak sector Feynman diagrams 20
2.9 $a_\mu^{\mathrm{HVP,LO}}$ from lattice-QCD 21
2.10 Comparison of a_μ^{exp} with select theory predictions 23

3.1 Schematic of the Fermilab Muon Campus 25
3.2 Dominant pion decay channel 26
3.3 Muon decay 26
3.4 Cross-section of the Muon g-2 magnetic storage ring 27
3.5 Top-down schematic view of the Muon g-2 storage ring 27
3.6 Positron number and asymmetry distributions, and figure-of-merit 32
3.7 Wiggle plot 34
3.8 Schematic view of a decay positron hitting the detectors 34
3.9 Side view of a straw tracker module 36
3.10 Cross-sectional view of trolley and fixed probe locations 39

4.1 2D multipole base functions 44
4.2 Illustration of bipolar coordinates. 49
4.3 Toroidal B_y shape for $n = 0$ 53
4.4 Direct comparison of 2D and 3D transverse B_y components. 54
4.5 Toroidal fit attempts for const. B_y 55
4.6 Tests of functions in $Q^n_{m-1/2}$ expression 60
4.7 Range of η and $\cosh\eta$ values inside the storage ring 61
4.8 Relative differences between analytical derivatives and numerical gradient . . . 64
4.9 Relative differences between each of the analytical derivatives 65
4.10 Singular value spectrum of $\mathbf{T}$ 66
4.11 F-test and χ^2_ν values for determining the number of transverse parameters . . . 71
4.12 2016 radial and longitudinal field measurements 71

4.13 Example fit results for two trolley probes . . . 76
4.14 Example fit results for STF1 . . . 77
4.15 Example fit results for STF5 . . . 78
4.16 Example fit results for STF6 . . . 79
4.17 Average RF-dataset beam distributions . . . 80

5.1 Muon weighting illustration . . . 86
5.2 Flow chart of muon weighting steps and inputs . . . 92
5.3 Unweighted field vs. time . . . 94
5.4 Run group algorithm comparison for trolley pair 5A2 . . . 98
5.5 Run group algorithm comparison for trolley pair 5B1 . . . 99
5.6 Comparison of run group durations . . . 100
5.7 Radial and vertical acceptance functions for beam reconstruction . . . 102
5.8 Illustration of beam reconstruction corrections . . . 103
5.9 Example of beam extrapolation results . . . 106
5.10 Time-dependence of Run-3b COSY beam dynamics functions . . . 107
5.11 Betatron functions vs azimuth . . . 109
5.12 Radial dispersion function vs azimuth . . . 110
5.13 Muon weighting results for the noRF dataset . . . 114
5.14 Muon weighting results for the xRF dataset . . . 115
5.15 Muon weighting results for the xyRF5 dataset . . . 116
5.16 Muon weighting results for the xyRF6 dataset . . . 117
5.17 Varied beam function residuals . . . 122
5.18 Shift of mean vertical beam position Δy_{COD} . . . 122
5.19 CODs from ESQ misalignments . . . 126
5.20 Monte Carlo CODs for ESQ misalignments . . . 127
5.21 Hard-coded ESQ misalignment COD extrapolation . . . 128
5.22 Total uncertainty on $\langle \tilde{B} \rangle$ due to averaging in azimuthal bins . . . 132

6.1 Top-down schematic of tracker station 18 . . . 137
6.2 Top-down cross-sectional view of a single tracker module . . . 137
6.3 Basic structure of a fully connected neural network (NN) . . . 140
6.4 LLH estimator flow chart . . . 144
6.5 Analytic likelihood for toy event A . . . 147
6.6 Analytic likelihood for toy event B . . . 148
6.7 Analytic likelihood for toy event C . . . 149
6.8 Loss and accuracy during NN training . . . 151
6.9 NN approximated likelihood for toy event A . . . 152
6.10 NN approximated likelihood for toy event B . . . 153
6.11 NN approximated likelihood for toy event C . . . 154
6.12 NN minimization results . . . 156

C.1 STF input trolley runs . . . 176
C.2 FFTs of trolley runs . . . 178
C.3 Graphs of STF1 . . . 180

C.4 Graphs of STF2 181
C.5 Graphs of STF3 183
C.6 Graphs of STF4 184
C.7 Graphs of STF5 186
C.8 Graphs of STF6 188
C.9 Graphs of STF7 190

LIST OF TABLES

TABLE

1.1 Summary of experimental results a_μ^{exp} . . . 3

2.1 Summary of individual contributions to a_μ^{SM} . . . 23

3.1 External parameters used to calculate a_μ . . . 31

4.1 List of 2D multipole moments . . . 45
4.2 Summary of simulated toroidal fields . . . 73
4.3 Summary of $|\vec{B}| - B_y$ for each STF . . . 81
4.4 Results of 2D and 3D fits to simulated toroidal fields . . . 82
4.5 Results of 2D and 3D fits to trolley runs . . . 82
4.6 Test of 2D and 3D fits to interpolated data (simulation) . . . 84

5.1 Comparison of field multipole moments and beam multipole projections . . . 91
5.2 List of required hard-coded run group splits during Run-4/5/6 . . . 100
5.3 Sub-target trolley pairs . . . 101
5.4 Summary of momentum deviation and spread for RF datasets . . . 107
5.5 Summary of muon weighting results for noRF dataset . . . 111
5.6 Summary of muon weighting results for xRF dataset . . . 112
5.7 Summary of muon weighting results for xyRF5 dataset . . . 112
5.8 Summary of muon weighting results for xyRF6 dataset . . . 113
5.9 Uncertainty results for each tracker correction . . . 120
5.10 COD due to B_r uncertainties . . . 124
5.11 ESQ centroid displacements . . . 125
5.12 Corrections and uncertainties from CODs due to ESQ misalignment . . . 127
5.13 Calorimeter acceptance uncertainty estimate . . . 130
5.14 Muon weighted field uncertainty due to azimuthal averaging . . . 131
5.15 Summary of muon weighting uncertainties . . . 134

6.1 Simple toy model example events . . . 147
6.2 Network architecture of the simple toy model classifier . . . 150
6.3 Minimization results for analytic and NN LLH of toy events . . . 157

A.1 Parameters and operational conditions of the muon storage ring . . . 162
A.2 Definition of RF datasets . . . 162
A.3 Overview of different coordinate systems in the Muon g-2 experiment . . . 163

B.1 Transformation between polar and cartesian coordinates 165

D.1 Beam projection statistical uncertainty in sub-target run groups 194
D.2 Multipole averages in sub-target run groups 195
D.3 Statistical uncertainty limits for sub-target run groups 196

LIST OF APPENDICES

A Experiment Parameters and Properties 162

B Derivations and Detailed Calculations 164

C Simulated Toroidal Field Plots 175

D Run Group Uncertainty Estimates 191

LIST OF ACRONYMS

BNL Brookhaven National Laboratory

BSM beyond the Standard Model

CBO coherent betatron oscillation

COD closed orbit distortion

COM center-of-mass

CTAG **C**alorimeter **TAG**ged

CTAGs CTAG count

CVC conserved vector current

DAQ data acquisition

DQC data quality cuts

ESQ electrostatic quadrupole

EW electroweak

FFT fast Fourier transform

FID free induction decay

FNAL Fermi National Accelerator Laboratory (Fermilab)

HLbL hadronic light-by-light scattering

HV high voltage

HVP hadronic vacuum polarization

IB isospin-breaking

LD long distance

LHC Large Hadron Collider

LLH log-likelihood

LO leading-order

l-QCD lattice quantum chromodynamics (lattice-QCD)

MC Monte Carlo

ML machine learning

NLO next-to-leading-order

NNLO next-to-next-to-leading-order

NN neural network

NMR nuclear magnetic resonance

QCD quantum chromodynamic

QED quantum electrodynamic

RMS root mean square

RF radio frequency

SD short distance

SEM standard error of the mean

SiPM silicon photo-multiplier

SM Standard Model of Particle Physics

STF simulated toroidal field

SVD singular value decomposition

VTR virtual trolley run

ABSTRACT

The Fermilab Muon g-2 experiment measured the muon's anomalous magnetic moment a_μ to 127 parts per billion (ppb) precision after six years of data collection and analysis. The measurement was performed by circulating a polarized muon beam in a magnetic storage ring, and correlating the difference of the cyclotron and spin-precession frequencies with the strength of the magnetic field. This dissertation presents a detailed report on the muon-weighted average of the magnetic field and a comprehensive evaluation of the associated uncertainties for the second half of the experiment's active data collection period, which are critical in the determination of a_μ. Furthermore, this dissertation introduces a three-dimensional toroidal magnetic field parameterization uniquely suited for the storage ring geometry and develops an optimized fit algorithm based on the toroidal parameterization. A realistic magnetic field simulation methodology is derived that validates a 2D multipole parameterization used in the standard magnetic field fit for the Muon g-2 experiment. Lastly, a machine-learning algorithm is explored to estimate track reconstruction parameters in a likelihood-free inference framework. This work concludes that the muon-weighting approach, effectively a type of phase-space weighting, is relevant for any non-uniform distribution of particles affected by fields, while the toroidal parameterization, specific to this geometry, shows significant potential, and suggests several avenues for further study for the toroidal parameterization and the machine-learning track reconstruction algorithm.

CHAPTER 1

Introduction

The Standard Model of Particle Physics (SM) is at present the most comprehensive available theory to describe the natural forces and building blocks that make up the physical universe, and it has so far accurately described and predicted numerous phenomena, the most recent example being the prediction and subsequent discovery of the Higgs boson at the Large Hadron Collider (LHC) at CERN in 2012 [1, 54], as well as withstood many falsification attempts. Yet despite the SM now being self-contained, i.e. all required "ingredients" have been experimentally confirmed and all predictions of the model can be explained within it, there is also ample evidence that it is not complete, some of the most notable examples including the existence of dark matter, neutrino masses, and gravity, none of which are included in the SM. Searches for physics beyond the Standard Model (BSM) typically focus either on probing very high energy regimes for new exotic particles and interactions ("energy frontier"), or using extremely well controlled environments to test SM predictions with high accuracy ("precision frontier"). The Muon g-2 experiment is one such high precision test of a SM quantity, namely the g-factor of the muon, a property describing the interaction strength of the particle's magnetic moment with magnetic fields.

1.1 A brief history of Muon g-2

The muon magnetic moment was measured for the first time in 1957 with 5% precision as part of a measurement on parity and charge conservation [77]. Shortly thereafter, the first dedicated experiment to measure the anomalous magnetic moment of the muon, a_μ, began at CERN [72], producing a 2% precision result in 1961[52], and an improved 0.4% precision result a year later [53]. Advances in accelerator technology made the production of higher energy muons possible, which lead to the proposal of a muon storage ring experiment in 1962 [72], and another six published results from the CERN experiment between 1962 and 1979. A proposal for a new iteration of the Muon g-2 experiment with improved precision

at Brookhaven National Lab (BNL) was developed starting in 1984 [72], for which a storage ring magnet was designed specifically optimized to create an extremely uniform field in the muon storage area, the same magnet used for the experiment that is the focus of this work and described in more detail in chapter 3. Data taking began in 1998, after extensive work on developing new strategies to measure the magnetic field inside the storage ring, new detector arrangements, and precise control of systematic effects, and resulted in a total of four measurements of the muon's anomalous magnetic moment. By the time the final result of the BNL experiment was published, a seeming discrepancy between theoretical predictions and experimentally determined values had emerged, culminating in a 2.9σ difference after taking a weighted average of all experiments [72], hinting at a potential for the presence of BSM physics. This generated substantial interest in the topic, and motivation to repeat the experiment with even higher precision, to confidently determine whether there was a substantial disagreement with the theoretical prediction. The Muon g-2 experiment at Fermilab (FNAL) is the (currently) final iteration of this project, and has very successfully measured the anomalous magnetic moment with a precision of 127 ppb (thereby surpassing its precision goal), as well as confirmed the previous experiments' values. Table 1.1 gives an overview of the historical development of muon g-2 results over the years. Interestingly, the discrepancy between theory and experiment has recently been mostly resolved by a change of the theoretical prediction, as new and more precise calculations and different analysis techniques became available. More details on this will be presented in chapter 2.

1.2 Fermilab Muon g-2 timeline

The Fermilab Muon g-2 experiment analyzed a total of more than 300 billion positrons over its six run periods between 2018 and 2023, with the majority of the data collected during run periods 4-6 from December 2020 to April 2023, as shown in fig. 1.1. Independent results were released after the Run-1 analysis, the Run-2/3 analysis, and the final result after the Run-4/5/6 analysis was concluded, summarized as well in table 1.1.

1.3 Contents of this work

The information above provides context and motivation for the importance of precision tests of the SM in general, and the Fermilab Muon g-2 experiment in particular. Chapter 2 will summarize the theoretical foundation for magnetic moments of elementary particles, and the contributions to the muon magnetic moment from different forces, to give a number for the predicted value of a_μ in the SM. This thesis will then focus on the Fermilab Muon

Year	Experiment	**result (uncertainty)**	**precision**	Δa_μ
1961	CERN [52]	$1\,145(22) \cdot 10^{-6}$	2%	0.9σ
1962	CERN [53]	$1\,162(5) \cdot 10^{-6}$	0.4%	0.6σ
1966	CERN [71]	$1\,165(3) \cdot 10^{-6}$	0.3%	0.2σ
1968	CERN [25, 26]	$116\,616(31) \cdot 10^{-8}$	266 ppm	0.9σ
1975	CERN (1974 data) [27]	$1\,165\,895(27) \cdot 10^{-9}$	23 ppm	0.5σ
1979	CERN (1974-76 data) [28]	$11\,659\,240(85) \cdot 10^{-10}$	7.3 ppm	0.3σ
2006	BNL [31]	$116\,592\,080(63) \cdot 10^{-11}$	540 ppb	$2.2 - 2.7\sigma$
2021	FNAL (Run 1) [2]	$116\,592\,040(54) \cdot 10^{-11}$	463 ppb	3.3σ
2023	FNAL (Run 2+3) [5]	$116\,592\,057(25) \cdot 10^{-11}$	214 ppb	4.97σ*
2025	FNAL (Run 4-6) [7]	$1\,165\,920\,710(162) \cdot 10^{-12}$	139 ppb	1.06σ*
2025	FNAL (Run 1-6) [7]	$1\,165\,920\,705(148) \cdot 10^{-12}$	127 ppb	1.06σ*
2025	Experimental average [7]	1 165 920 715(145)	124 ppb	1.1σ*

Table 1.1: Summary of dedicated Muon g-2 experimental results with uncertainties. Only BNL + FNAL contribute to the experimental average. $\Delta a_\mu = a_\mu^{\text{exp}} - a_\mu^{\text{SM}}$ at the time of publication of the respective experimental result.
Values marked with * should be taken with caution, as significant tension has been found in some contributions to the theory value, and SM predictions are likely to change again in the future. (For more details, see chapter 2.)

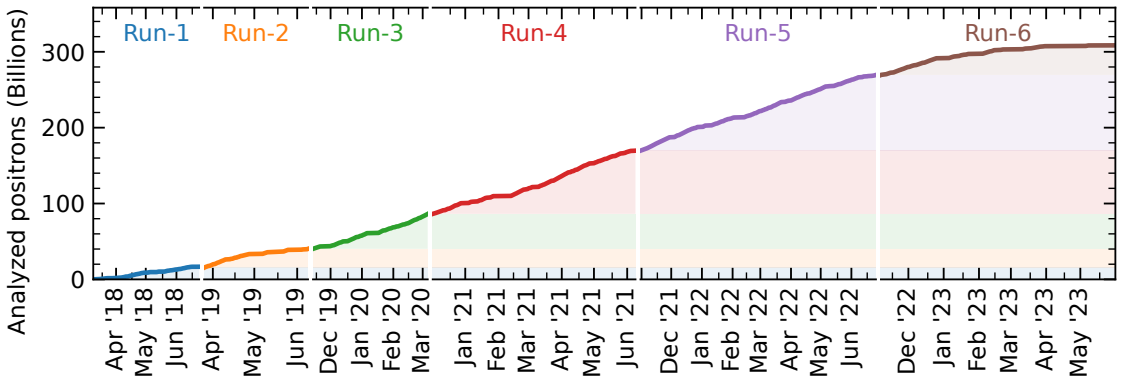


Figure 1.1: Accumulated data collected during the Fermilab Muon g-2 experiment, represented as positron events in the final dataset during the six run periods, courtesy of the Muon g-2 collaboration.

g-2 experiment, whose measurement principle and components are described in chapter 3. Chapter 4 goes into significant detail with respect to the analysis of magnetic field data, which is an integral part of the measurement of a_μ, and develops a secondary analysis method used for validation and systematic tests. The main contribution of this work is chapter 5, which reports the process and results of determining the average of the magnetic field weighted by the muon distribution for the second half of the data collection period (Run-4/5/6), and carefully examines its systematic uncertainties.

CHAPTER 2

Theoretical Background

2.1 Magnetic dipole moments

Classically, a magnetic dipole moment $\vec{\mu}$ is the lowest order in the multipole expansion of the magnetic vector potential of a current distribution $\vec{J}$, defined as [100]

$$\vec{\mu} = \frac{1}{2}\int (\vec{x} \times \vec{J}) \mathrm{d}^3x, \tag{2.1}$$

with the standard cartesian 3-dimensional infinitesimal volume element d^3x and position in space $\vec{x}$.
For a discrete current distribution, such as an ensemble of identical charged particles with charge q, mass m, and momentum $\vec{p_i}$ at position $\vec{x_i}$, the magnetic dipole moment can be written in terms of the total angular momentum $\vec{L} = \sum_i \vec{x_i} \times \vec{p_i}$ as

$$\vec{\mu} = \frac{q}{2m}\vec{L}. \tag{2.2}$$

Within an external magnetic field $\vec{B}$, a magnetic dipole moment exhibits observable behaviors that are due to the torque $\vec{\tau}$ and potential energy U, defined as [100]

$$\vec{\tau} = \vec{\mu} \times \vec{B} = -\frac{q}{2m}\vec{B} \times \vec{L}, \tag{2.3}$$

$$U = -\vec{\mu} \cdot \vec{B} = -\frac{q}{2m}\vec{L} \cdot \vec{B}. \tag{2.4}$$

These equations imply that, in the absence of other forces, the magnetic dipole moment will align itself with the axis of the external field to minimize potential energy, and otherwise the torque will cause a precession with frequency

$$\vec{\omega} = -\frac{q}{2m}\vec{B}, \tag{2.5}$$

such that $\vec{\tau} = \vec{\omega} \times \vec{L}$.

For quantum mechanical charged particles, such as muons, with intrinsic angular momentum (spin) $\vec{S}$, the magnetic dipole moment must be aligned with its spin axis and is defined analogously to eq. 2.2 as [100, 131, 128]

$$\vec{\mu} = g\frac{q}{2m}\vec{S}, \tag{2.6}$$

with the addition of a dimensionless factor g, a measure of how strongly a particle couples to magnetic fields.

The interaction Hamiltonian of such a particle with a magnetic field $\vec{B}$ is

$$\mathcal{H} = -\vec{\mu} \cdot \vec{B} = -g\frac{q}{2m}\vec{S} \cdot \vec{B}. \tag{2.7}$$

The energy eigenstates of this Hamiltonian are separated by $\Delta E = \hbar\omega$, and any spin component not aligned with the magnetic field will lead to a precession with frequency [131]

$$\omega = g\frac{qB}{2m}. \tag{2.8}$$

Both the Hamiltonian and the precession frequency ω are consistent with the classical result for $g = 1$.

2.2 Standard Model prediction

The generalized S-matrix element of the scattering amplitude $i\mathcal{M}$ of an incoming muon with (4-vector) momentum p from an external field A_μ, resulting in an outgoing muon with momentum p' (depicted in fig. 2.1c), is [128]

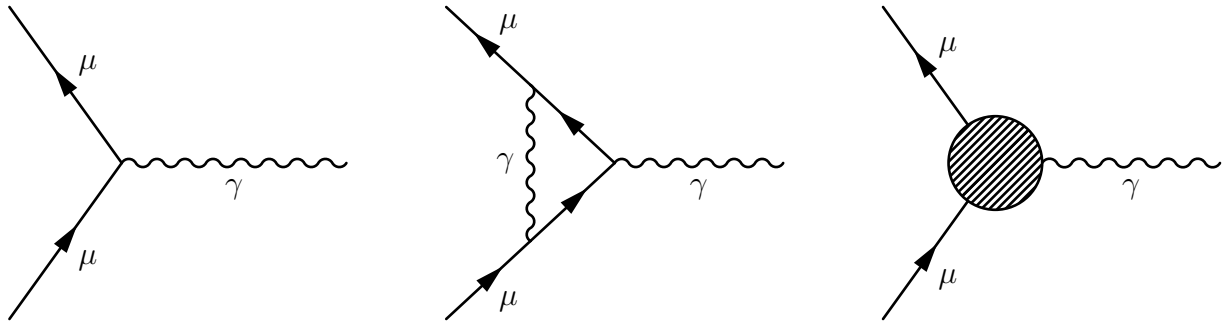
$$i\mathcal{M}(2\pi)\delta(p^{0'} - p^0) = -ie\bar{u}(p')\Gamma^\mu(p', p)u(p) \cdot \tilde{A}_\mu(p' - p). \tag{2.9}$$

The delta function[1] $\delta(p^{0'} - p^0)$ on the left hand side ensures conservation of energy (with $p^0 = E$ in 4-vector notation). $u(p)$ and $\bar{u}(p')$ are the Dirac spinors of the incoming and

[1]The Dirac delta function is defined as

$$\delta(x) = \begin{cases} 0 & \text{if } x \neq 0, \\ \infty & \text{if } x = 0, \end{cases}$$

such that $\int_{-\infty}^{\infty} \delta(x)\mathrm{d}x = 1$ and $\int_{-\infty}^{\infty} f(x)\delta(x)\mathrm{d}x = f(0)$.



(a) Dirac vertex/tree level diagram ($g = 2$)

(b) QED 1-loop correction/Schwinger term ($g = 2 + \frac{\alpha}{4\pi}$)

(c) General vertex diagram, including all corrections ($g = 2(1 + a)$)

Figure 2.1: Feynman diagrams of a muon interacting with a magnetic field

outgoing muon respectively, $\tilde{A}_\mu$ is the Fourier transform of A_μ, $-ie$ is the coupling constant with electric charge e, and the vertex function Γ^μ reduces to a single Dirac gamma matrix γ^μ at tree level. Assuming parity conservation and Lorentz invariance, Γ^μ must depend only on p, p', γ^μ and numbers or physical constants such as mass m and electric charge e and can be expressed in the form [128]

$$\Gamma^\mu(p', p) = \gamma^\mu F_1(q^2) + \frac{i\sigma^{\mu\nu} q_\nu}{2m} F_2(q^2), \tag{2.10}$$

where $\sigma^{\mu\nu} = \frac{i}{2}[\gamma^\mu, \gamma^\nu]$, $q = p' - p$ and $q^2 = m^2 - 2p' \cdot p$. $F_1(q^2)$ and $F_2(q^2)$ are called *form factors*. By considering the case of a purely scalar electrostatic potential $A_\mu(x) = (\Phi(\vec{x}), \vec{0})$, it can be shown that $F_1(0)$ is just the electric charge in units of e [128]. Similarly, the case of a static vector potential $A_\mu(x) = (0, \vec{A}(\vec{x}))$ leads to an expression for the expected value of the magnetic moment

$$\langle \vec{\mu} \rangle = \frac{e}{m}[F_1(0] + F_2(0)]\xi'^\dagger \frac{\vec{\sigma}}{2}\xi = \frac{e}{m}[F_1(0] + F_2(0)]\left\langle \vec{S} \right\rangle, \tag{2.11}$$

with the two-component Pauli spinors ξ, $\xi'^\dagger$ and the vector of Pauli matrices $\vec{\sigma}$ forming the expected value of the spin operator $\left\langle \vec{S} \right\rangle$. Comparing this to eq. 2.6, g can be written in terms of the form factors as [128]

$$g = 2[F_1(0) + F_2(0)] = 2 + 2F_2(0). \tag{2.12}$$

Since $\Gamma^\mu = \gamma^\mu$ at tree level, it is clear from eq. 2.10 that at tree level $F_1(0) = 1$ and $F_2(0) = 0$, i.e. $g = 2$. However, higher order corrections from loop diagrams shift the value of g slightly away from 2. The *anomalous magnetic moment* a is defined as the deviation of g from 2,

$$a = \frac{g-2}{2} = F_2(0). \tag{2.13}$$

The value of the muon anomalous magnetic moment a_μ can be predicted in the SM by calculating the loop correction contributions in the quantum electrodynamic (QED), electroweak (EW) and hadronic sectors [20]

$$a_\mu^{\mathrm{SM}} = a_\mu^{\mathrm{QED}} + a_\mu^{\mathrm{EW}} + a_\mu^{\mathrm{had}}. \tag{2.14}$$

This effort is spearheaded by the Muon g-2 Theory Initiative, which monitors and evaluates the results of several different groups and experiments, and carefully combines them into a single SM value for a_μ^{SM}. The most recent SM prediction was published in the 2025 White Paper [16], updating the previous paper from 2020 [20].

2.2.1 Quantum electrodynamic contribution

QED sector corrections include all interactions comprised of only leptons and photons, and are proportional to the QED coupling constant $\alpha = \frac{e^2}{4\pi}$ (fine structure constant), such that a_μ^{QED} can be written as [20]

$$a_\mu^{\mathrm{QED}} = \sum_n \left(\frac{\alpha}{\pi}\right) A^{(2n)}. \tag{2.15}$$

Each $A^{(2n)}$ can also be split into a mass-independent term, and terms that depend on the lepton mass ratios $\frac{m_\mu}{m_e}$ (muon to electron mass ratio) and $\frac{m_\mu}{m_\tau}$ (muon to tau mass ratio). The dominant correction in the QED sector is the Schwinger term (shown in fig. 2.1b), first calculated by Julian Schwinger in 1948 [132, 128],

$$F_2(0) = \frac{\alpha}{2\pi} + \mathcal{O}(\alpha^2) \approx 0.0011614, \tag{2.16}$$

leading to $A^{(2)} = \frac{1}{2}$. Mass dependent terms don't exist at second order, since they arise from Feynman diagrams with at least one closed fermion loop.

All terms up to $\mathcal{O}(\alpha^5)$ have been calculated using perturbation theory either analytically or numerically [21, 22, 20]. The final value for a_μ^{QED} also depends on the value of the fine structure constant α, which is in itself a quantity determined by experimental

results. Three competitive results for α exist at the moment: Two values, obtained from interferometry measurements of the atomic mass of cesium (Cs-133) [127] and rubidium (Rb-87) [123],

$$\alpha^{-1}(\mathrm{Cs}) = 137.035\ 999\ 045(27), \tag{2.17}$$

$$\alpha^{-1}(\mathrm{Rb}) = 137.035\ 999\ 2052(97), \tag{2.18}$$

differ by about 5.5σ [16]. A third value for α comes from a comparison of the measured and predicted anomalous magnetic moment of the electron a_e [70, 16],

$$\alpha^{-1}(a_e) = 137.035999163(15), \tag{2.19}$$

which lands in between the two interferometry results, slightly closer to the one obtained from rubidium measurements. It is worth noting that the determination of the SM prediction of a_e relies heavily on calculations that are either the same (for mass-independent terms) or performed by similar algorithms (for mass-dependent terms) as in the calculation of a_μ^{QED} and is therefore strongly correlated to it [20].

Due to the statistically significant differences between the three α values, they cannot simply be averaged. The respective results for a_μ^{QED} are [16]:

$$a_\mu^{\mathrm{QED}}(\mathrm{Cs}) = 116\ 584\ 718.932(104) \cdot 10^{-11}, \tag{2.20}$$

$$a_\mu^{\mathrm{QED}}(\mathrm{Rb}) = 116\ 584\ 718.795(102) \cdot 10^{-11}, \tag{2.21}$$

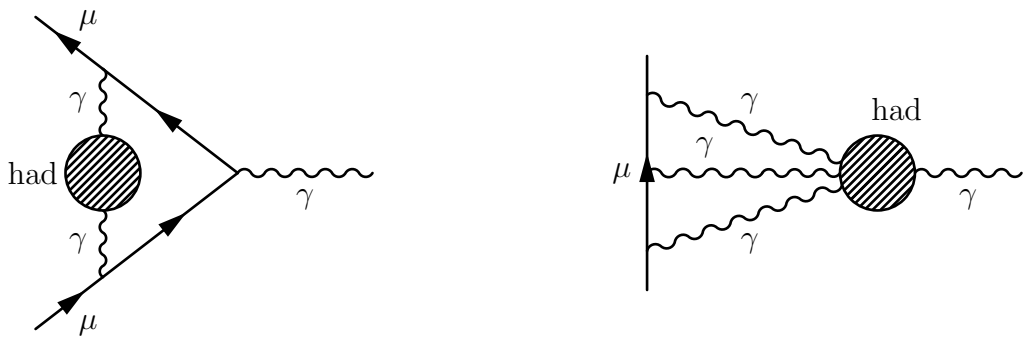
$$a_\mu^{\mathrm{QED}}(a_e) = 116\ 584\ 718.833(103) \cdot 10^{-11}. \tag{2.22}$$

In each case, the uncertainty on a_μ^{QED} is dominated by an estimate of the size of the $\mathcal{O}(\alpha^6)$ term, which contributes $0.100 \cdot 10^{-11}$ to the total uncertainties [16]. The remaining uncertainty contributions (fine structure constant, numerical integration, and the lepton masses) all contribute about an order of magnitude less. Considering the largest difference between these three a_μ^{QED} results is on the same scale as the uncertainty due to the $\mathcal{O}(\alpha^6)$ term, a combined value for a_μ^{QED} can be taken as [16]

$$a_\mu^{\mathrm{QED}} = 116\ 584\ 718.8(2) \cdot 10^{-11}. \tag{2.23}$$

2.2.2 Quantum chromodynamic contribution

The hadronic, or quantum chromodynamic (QCD) sector contributions are the largest source



(a) Hadronic vacuum polarization (b) Hadronic light-by-light scattering

Figure 2.2: Hadronic sector Feynman diagrams: Feynman diagrams illustrating the two mechanisms in the hadronic sector

of uncertainty in a_μ^{SM}, as they cannot be calculated perturbatively but must rely on other means such as data-driven dispersive methods and lattice-QCD (l-QCD) calculations. a_μ^{had} includes all interactions involving only quarks and strongly interacting photons[2], and is split again into two terms corresponding to hadronic vacuum polarization (HVP) and hadronic light-by-light scattering (HLbL) [20],

$$a_\mu^{\mathrm{had}} = a_\mu^{\mathrm{HVP}} + a_\mu^{\mathrm{HLbL}}. \tag{2.24}$$

Hadronic vacuum polarization

The dominant contribution to a_μ^{had} is the leading-order (LO) HVP term at $\mathcal{O}(\alpha^2)$, illustrated in fig. 2.2a. For a data-driven approach, the loop integral containing the HVP insertion into the photon propagator can be expressed equivalently as a dispersion integral [20, 16],

$$a_\mu^{\mathrm{HVP,LO}} = \frac{\alpha^2}{3\pi^2} \int_{M_\pi^2}^{\infty} \frac{K(s)}{s} R(s)\ \mathrm{d}s, \tag{2.25}$$

where s is the center-of-mass (COM) energy squared, M_π is the mass of the pion π^0, $R(s)$ is the hadronic R-ratio, obtained from electron-positron annihilation experiments,

$$R(s) = \frac{\sigma^0(e^+e^- \to (\gamma+)\mathrm{hadrons})}{\sigma_{pt}}, \tag{2.26}$$

[2]Mixed hadronic and electroweak interactions are included in the electroweak term, mixed hadronic and leptonic interactions contribute less than 10^{-11} to a_μ^{had} [97].

and

$$K(s) = \frac{x^2}{2}\big(2 - x^2\big) + \frac{(1+x^2)(1+x)^2}{x^2}\Big(\log(1+x) - x + \frac{x^2}{2}\Big) + \frac{1+x}{1-x}x^2\log(x), \quad (2.27)$$

with $x = \frac{1-\beta_\mu}{1+\beta_\mu}$, $\beta_\mu = \sqrt{1 - \frac{4m_\mu^2}{s}}$. $R(s)$ is the ratio of the bare cross-section σ^0, i.e. excluding vacuum polarization effects, to the lowest order point-like muon pair cross-section $\sigma_{pt} = \frac{4\pi\alpha^2}{3s}$. Vacuum polarization insertions are taken into account in the higher-order HVP terms.

Due to the $\frac{K(s)}{s}$ term, low energies are weighted most strongly, and in this region ($\sqrt{s} <\sim 2\,\mathrm{GeV}$) all hadronic channels must be evaluated and summed explicitly [20]. The largest contribution to the cross-section is the two-pion channel $e^+e^- \to \pi^+\pi^-$ at more than 70%. Fig. 2.3 shows the values of the two-pion channel contribution $a_\mu^{\pi^+\pi^-,\mathrm{LO}}$ to $a_\mu^{\mathrm{HVP,LO}}$ determined by different experiments in the energy range $0.6\,\mathrm{GeV} < \sqrt{s} < 0.88\,\mathrm{GeV}$ (a range covered by all experiments). Notably, this figure includes two results that were not yet available for the Theory Initiative's publication in 2020: the 2021 result of the SND collaboration ("SND2k") [4], and the 2024 result of the CMD-3 collaboration [99], which also has one of the smallest uncertainties. In their 2020 review, the Theory Initiative based their recommendation of $a_\mu^{\mathrm{HVP,LO}}$ entirely on these data-driven results, due to the significantly smaller uncertainties compared to other methods. While the SND2k result is consistent with previous measurements, the most recent CMD-3 result disagrees with all previous experiments, even CMD-2, to such a degree that forming a combined or averaged result is no longer appropriate [16]. It should also be noted that even within the results that contributed to the Theory Initiative's combined value in 2020, the KLOE [18] and BaBar [111] results in particular are in significant tension.

Alternatively, the hadronic cross-section in eq. 2.26 can be related to spectral functions obtained from hadronic τ decay via simple kinematic factors, assuming conserved vector current (CVC) i.e. isospin invariance [12, 16]. For example, the dominant two-pion channel $e^+e^- \to \pi^+\pi^-$ can be correlated with the dominant τ decay mode $\tau^- \to \pi^-\pi^0\nu_\tau$, since π^+, π^-, and π^0 form an isospin triplet. Spectral functions and branching ratios for this decay channel have been obtained by four different experiments and are in excellent agreement [64, 16]. However, isospin-breaking (IB) corrections have to be applied, which are the largest source of uncertainty in this approach [115, 16]. Efforts to improve the evaluation of these corrections are ongoing.

A full prediction of $a_\mu^{\mathrm{HVP,LO}}$ from τ decay data is not yet possible due to insufficient knowledge or control of the IB corrections for other channels [16], some of which might soon be accessible via l-QCD calculations [46].

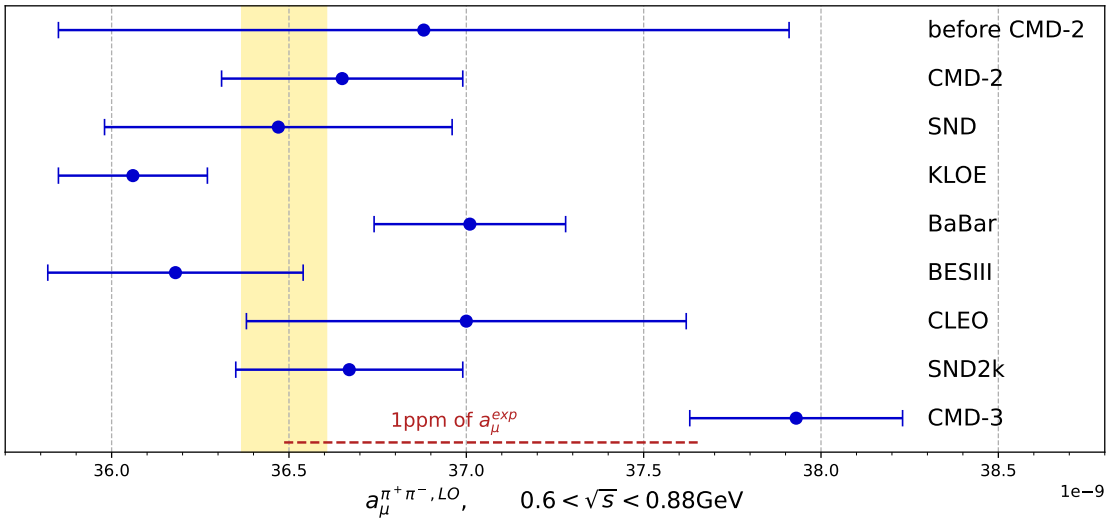


Figure 2.3: Results for $a_\mu^{\pi^+\pi^-,\mathrm{LO}}$ in the energy region $0.6\,\mathrm{GeV} < \sqrt{s} < 0.88\,\mathrm{GeV}$, which accounts for more than 50% of $a_\mu^{\mathrm{HVP,LO}}$. The yellow area shows a simple weighted average of all experiments before (visually above) CMD-3, which illustrates significant tension with the CMD-3 result (bottom). Also shown is 1 ppm of the current experimental world average [7] for reference (dashed red line). Reproduced from data in [99].

Next-to-leading-order (NLO) [102] and next-to-next-to-leading-order (NNLO) [109] contributions have also been calculated with dispersion relations. The Theory Initiative recommends the values

$$a_\mu^{\mathrm{HVP,NLO}} = -9.96(13) \times 10^{-10}, \tag{2.28}$$

$$a_\mu^{\mathrm{HVP,NNLO}} = 1.24(1) \times 10^{-10}, \tag{2.29}$$

where the tension between CMD-3 and previous experiments is taken into account through an increased uncertainty [16].

Lattice-QCD is a non-perturbative approach to QCD that avoids the problem of the strong force's non-linearity and the increase of the coupling constant at low energies, by discretizing spacetime into a grid of points (the lattice) with fixed distances from each other [143]. As such, path integrals that cannot be solved in continuous spacetime become well-defined and can be evaluated numerically, allowing for a calculation of $a_\mu^{\mathrm{HVP,LO}}$ from first principles. $a_\mu^{\mathrm{HVP,LO}}$ can be expressed in terms of a known function $\tilde{K}(t)$ and the 2-point correlation

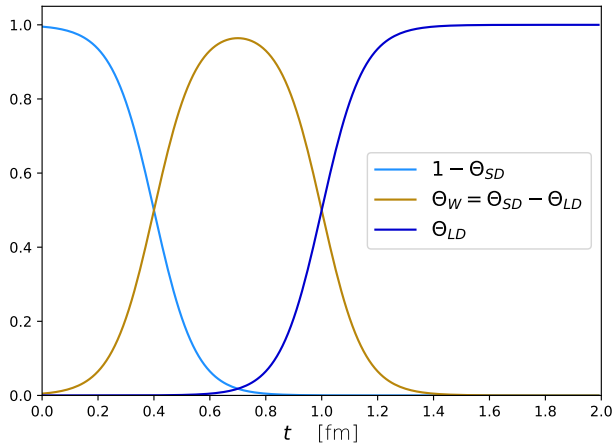


Figure 2.4: Illustration of HVP window weights: Illustration of the three window weights for short distance (SD), intermediate window (W) and long distance (LD) in eqs. 2.35-2.37, where $\Theta_{SD} = \Theta(t, t_0, \Delta)$, $\Theta_{LD} = \Theta(t, t_1, \Delta)$, and $t_0 = 0.4\,\text{fm}$, $t_1 = 1.0\,\text{fm}$, $\Delta = 0.15\,\text{fm}$.

function $G(t)$ [124] as

$$a_\mu^{\text{HVP,LO}} = \left(\frac{\alpha}{\pi}\right)^2 \int_0^\infty \mathrm{d}t\; G(t)\tilde{K}(t), \tag{2.30}$$

where t is Euclidean time, and

$$\tilde{K}(t) = 8\pi^2 \int_0^\infty \mathrm{d}\omega\; \frac{1}{\omega}\; f(\omega) \left(\omega^2 t^2 - 4\sin^2(\frac{\omega t}{2})\right), \tag{2.31}$$

$$f(\omega^2) = \frac{1}{\omega\sqrt{\omega^2 + 4m_\mu^2}} - \frac{1}{m_\mu^2} + \frac{\omega}{2m_\mu^4}\left(\sqrt{\omega^2 + 4m_\mu^2} - \omega\right). \tag{2.32}$$

The main effort lies in calculating the correlation function $G(t)$ for different quark flavors at different values of lattice spacing a and extrapolating to the continuum limit ($a \to 0$). Additionally, it is helpful to perform the calculations in an isospin-symmetric scheme, and treat the isospin-breaking component that arises due to the mass and charge difference of the up- and down-quarks as a correction $\delta a_\mu^{\text{HVP,LO}}(\text{IB})$ [16]. The total HVP contribution can thus be written as a sum of individual quark flavor terms[3],

$$\begin{aligned} a_\mu^{\text{HVP,LO}} = {} & a_\mu^{\text{HVP,LO}}(l) + a_\mu^{\text{HVP,LO}}(s) + a_\mu^{\text{HVP,LO}}(c) + a_\mu^{\text{HVP,LO}}(b) \\ & + a_\mu^{\text{HVP,LO}}(\text{disc}) + \delta a_\mu^{\text{HVP,LO}}(\text{IB}). \end{aligned} \tag{2.33}$$

The key contribution is the connected, isospin-symmetric light-quark contribution $a_\mu^{\text{HVP,LO}}(l)$, $l \in (u, d)$, which makes up about 90% of $a_\mu^{\text{HVP,LO}}$ [20], followed by the strange- ($a_\mu^{\text{HVP,LO}}(s)$)

[3]The top-quark term is omitted, because its contribution is negligible at the currently accessible precision of $a_\mu^{\text{HVP,LO}}$ [107].

and charm-quark ($a_\mu^{\mathrm{HVP,LO}}(c)$) contributions. The disconnected contribution $a_\mu^{\mathrm{HVP,LO}}(\mathrm{disc})$ accounts for disconnected quark loops, i.e. sea quarks of all flavors.

Because calculations of different energy regimes have different systematic effects, a window approach [36, 16] can be useful, which splits the integral of eq. 2.30 into three parts - short distance (SD), intermediate window (W), and long distance (LD)[4]:

$$a_\mu^{\mathrm{HVP,LO}} = a_\mu^{\mathrm{SD}} + a_\mu^{\mathrm{W}} + a_\mu^{\mathrm{LD}}, \tag{2.34}$$

$$a_\mu^{\mathrm{SD}} = \left(\frac{\alpha}{\pi}\right)^2 \int_0^\infty \mathrm{d}t\ G(t)\tilde{K}(t)\ \big[1 - \Theta(t, t_0, \Delta)\big], \tag{2.35}$$

$$a_\mu^{\mathrm{W}} = \left(\frac{\alpha}{\pi}\right)^2 \int_0^\infty \mathrm{d}t\ G(t)\tilde{K}(t)\ \big[\Theta(t, t_0, \Delta) - \Theta(t, t_1, \Delta)\big], \tag{2.36}$$

$$a_\mu^{\mathrm{LD}} = \left(\frac{\alpha}{\pi}\right)^2 \int_0^\infty \mathrm{d}t\ G(t)\tilde{K}(t)\ \Theta(t, t_1, \Delta). \tag{2.37}$$

The function

$$\Theta(t, t', \Delta) = \frac{1}{2}\left[1 + \tanh(\frac{t - t'}{\Delta})\right] \tag{2.38}$$

is utilized to define a smooth window with width parameter Δ as illustrated in fig. 2.4. A common choice for t_0, t_1, and Δ are 0.4 fm, 1.0 fm, and 0.15 fm respectively. Within each window, eq. 2.33 still holds. The intermediate window is designed to be less sensitive to certain systematic effects like discretization errors (SD) or signal-to-noise problems (LD) and lends itself to comparison between independent lattice calculation efforts [16]. It also allows for direct comparisons between lattice and data-driven methods [44, 55, 32, 63] (see fig. 2.6), as the R-ratio $R(s)$ (eq. 2.26) can be translated into a value for $G(t)$ via [33, 20]

$$G(t) = \frac{1}{12\pi^2} \int_0^\infty \mathrm{d}(\sqrt{s}) R(s) s e^{-\sqrt{s}t}. \tag{2.39}$$

Additionally, the window approach can also be combined with the dispersive approach in the SD and/or LD regions to obtain a hybrid result [36, 41].

At the time of the previous Theory Initiative review [20], l-QCD uncertainties were too large to compete with the data-driven dispersive approach, and were not included in the combined recommended value for $a_\mu^{\mathrm{HVP,LO}}$. In the past few years however, several groups

[4] Euclidean time t and COM energy $\sqrt{s}$ are related through a Laplace transform and thus don't exhibit a simple analytic correlation. However, the three windows weights (SD, W, LD) roughly map to inversely proportional energy weights - The SD window emphasizes larger COM energies, the LD window focuses on low-energy effects, and the intermediate window weight peaks around a mid-level energy of $\sqrt{s} \sim 1.5\,\mathrm{GeV}$ [55].

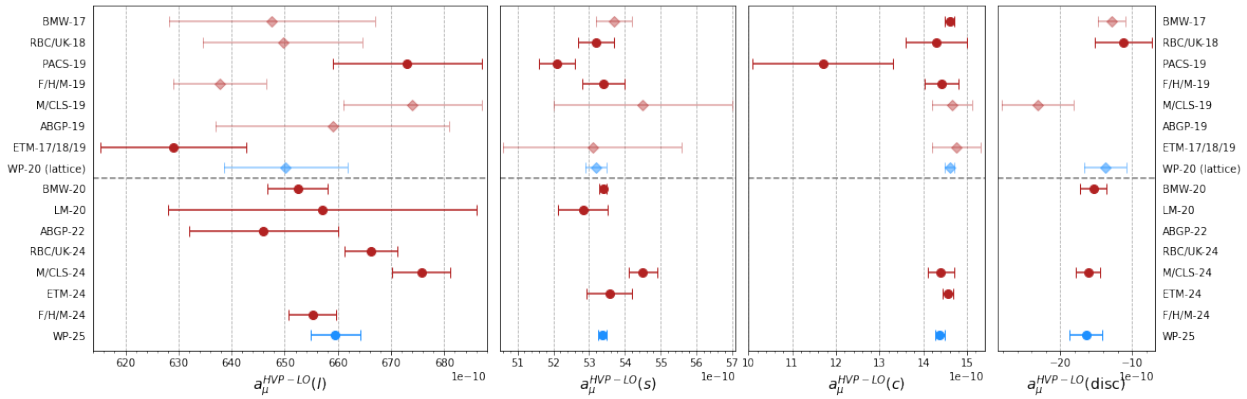


Figure 2.5: Quark flavor contributions to $a_\mu^{\mathrm{HVP,LO}}$: Individual isospin-symmetric connected light- (l), strange- (s) and charm-quark (c) flavor contributions to $a_\mu^{\mathrm{HVP,LO}}$, as well as the isospin-symmetric disconnected (disc) contribution, calculated in l-QCD, are shown in red. The Theory Initiative's combined lattice values from 2020 (WP-20 [20]) and 2025 (WP-25 [16]) are shown in light blue. Data points above the grey dashed line entered the WP-20 average, values below were published later. Semi-transparent points with diamond markers have been superseded with more recent results by the same collaboration, and did not enter the WP-25 average.

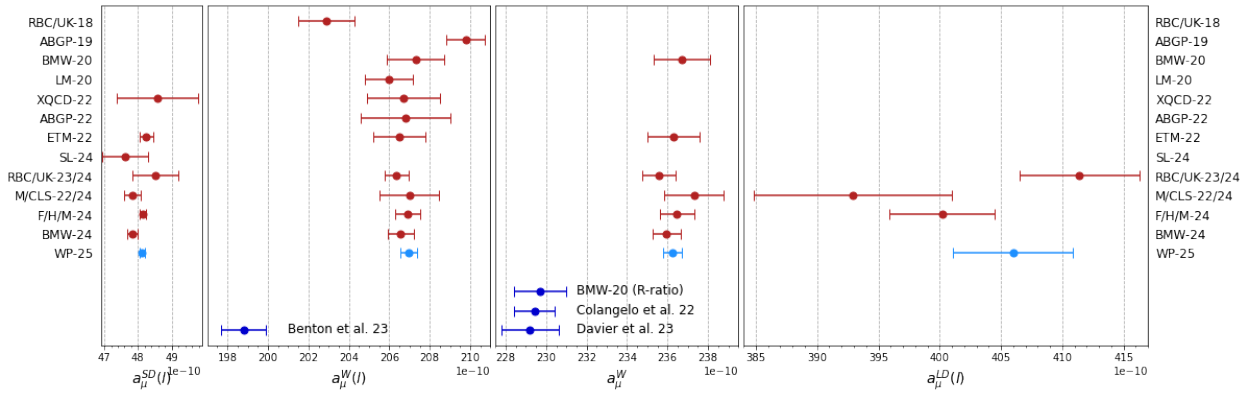


Figure 2.6: Window contribution to $a_\mu^{\mathrm{HVP,LO}}$: Light-quark contributions (connected, isospin-symmetric) to the SD, intermediate, and LD window, as well as total intermediate window contribution to $a_\mu^{\mathrm{HVP,LO}}$. Red data points show results of lattice-QCD calculations; light blue data points indicate the averaged values from the Theory Iniative's 2025 review (WP-25 [16]); dark blue data points show intermediate window contribution values derived from R-ratio data: BMW-20 [44] (performed for comparison as part of their l-QCD calculations), Colangelo et al. [55], Benton et al. [32], and Davier et al. [63]. For both the light-quark only and the total intermediate window contribution, there is a clear discrepancy between the lattice-QCD and the data-driven results.

have made significant progress in calculating new contributions (in particular LD window contributions by the Fermilab/HPQCD/MILC [30], Mainz/CLS [67], and RBC/UKQCD [38] collaborations) and greatly improving uncertainties of previously calculated values.

Fig. 2.5 shows results for the four largest terms in eq. 2.33: the light-, strange-, and charm-quark isospin-symmetric connected contributions $a_\mu^{\mathrm{HVP,LO}}(l), a_\mu^{\mathrm{HVP,LO}}(s)$, and $a_\mu^{\mathrm{HVP,LO}}(c)$ respectively, as well as the isospin-symmetric disconnected contribution $a_\mu^{\mathrm{HVP,LO}}(\mathrm{disc})$. Six groups have released improved values for some or all of their previously calculated values since the 2020 Theory Initiative review. The BMW collaboration has updated their 2017 values (BMW-17 [43]) in 2020 (BMW-20 [44]) for the light- and strange-quark and disconnected contributions, maintaining the charm-quark contribution; the RBC/UKQCD group released a new value for the light-quark contribution in 2024 (RBC/UK-24 [38]), while keeping their results from 2018 (RBC/UK-18 [36]) for the other three terms; the Fermilab/HPQCD/MILC collaborations similarly updated the light-quark term in 2024 (F/H/M-24 [30]) from their previous value in 2019 (F/H/M-19 [65]), and retain the results for strange- and charm-quark contributions; all four terms were updated in 2024 by the Mainz/CLS group (M/CLS-24 [67]) from their 2019 values (M/CLS-19 [78]); Aubin et al., who have only published a light-quark result in 2019 (ABGP-19 [24]), have released an update in 2022 (ABGP-22 [23]); and the ETM collaboration published their last light-quark result in 2019 (ETM-17/18/19 [81, 83, 82, 84]), but released new values for the strange- and charm-quark contributions in 2024 (ETM-24 [15]). Fig. 2.5 shows the significant uncertainty improvements for these updated values compared to the previous ones, in particular for the dominant light-quark contribution. It should also be noted that while the total $a_\mu^{\mathrm{HVP,LO}}$ (as well as the total window contributions in eq. 2.34) is independent if the specific isospin-symmetric scheme chosen, the split into isospin-symmetric flavor terms and $\delta a_\mu^{\mathrm{HVP,LO}}(\mathrm{IB})$ depends on the exact renormalization inputs and varies among groups. Typically however, the difference arising from the use of different schemes is smaller than the uncertainties on the respective terms [16].

In addition to the total quark-flavor contributions, many groups have also released results for flavor contributions in the window approach since 2020. Fig. 2.6 shows a summary of available values for the light-quark contribution in all three windows, as well as an overview of the total intermediate window contribution from both l-QCD and the dispersive method according to eq. 2.39. Almost all groups who published total flavor contributions also published flavor contributions for some or all windows - Aubin et. al [24, 23], the BMW collaboration [44, 41], the ETM collaboration [13], the Fermilab/HPQCD/MILC group [29, 30], Lehner and Meyer (LM-20) [113], the Mainz/CLS collaboration [47, 108, 67], and the RBC/UKQCD group [37, 38] - as well as two groups who only calculated window

quantities but not total flavor contributions: the χ-QCD collaboration (XQCD-22) [142], and Spiegel and Lehner (SL-24) [134]. In their 2020 publication, the BMW group also determined a value based on R-ratio data for the total intermediate window contribution. Similar results have been released by Colangelo et al. in 2022 [55] and Davier et al. in 2023 [63] for the total, and by Benton et al. in 2023 [32] for the light-quark connected intermediate window contribution. All four of these data-driven results are in significant tension with the ones obtained from l-QCD.

Both the CMD-3 results, as well as the disagreement of l-QCD calculations with the dispersive values (while being very consistent among l-QCD results) calls the reliability of data-driven determinations of $a_\mu^{\mathrm{HVP,LO}}$ into question. Given the increased number and precision of l-QCD results, the Theory Initiative's 2025 averaged value is based solely on l-QCD calculations and does not include data-driven results. After carefully correcting for scheme-dependent differences, averaging all available and current results for each term in eq. 2.33 within each window, and summing it all up, they arrive at [16]

$$a_\mu^{\mathrm{HVP,LO}} = 713.2(6.1) \times 10^{-10}. \tag{2.40}$$

This procedure of determining the average includes the largest number of independent results. As a cross-check, averages and sums of results for the total quark-flavor terms (not limited to a window, cf. fig. 2.5), and an average of directly reported $a_\mu^{\mathrm{HVP,LO}}$ values were also performed, which all agree within their uncertainties with the above value [16].

Hadronic light-by-light scattering

Similarly to the hadronic vacuum polarization contribution, the hadronic light-by-light scattering contribution cannot be calculated perturbatively and must be evaluated either through a dispersive approach, or lattice-QCD calculations. It starts to contribute at $\mathcal{O}(\alpha^3)$ and is illustrated in fig. 2.2b. Unlike for HVP however, the individual contributions to HLbL are not unique in those calculations, but rather model-dependent [20]. HLbL requires the evaluation of a four-point correlator function instead the HVP two-point function $G(t)$ (cf. eq. 2.30) and thus contains significantly more intermediate states (e.g. exchange of single pseudoscalar mesons like π^0, η, or η', also called pole diagrams; loops with charged pions and kaons; or quark loops at higher momenta [20]). It is assumed (and so far confirmed by all calculations) that the size of the individual contributions scales inversely with the corresponding energy threshold (i.e. particle mass), and the majority of $a_\mu^{\mathrm{HLbL,LO}}$ is expected to come from states below 1.5 GeV, with the π^0-pole as the dominant contribution $a_\mu^{\mathrm{HLbL},\pi^0\text{-pole}}$.

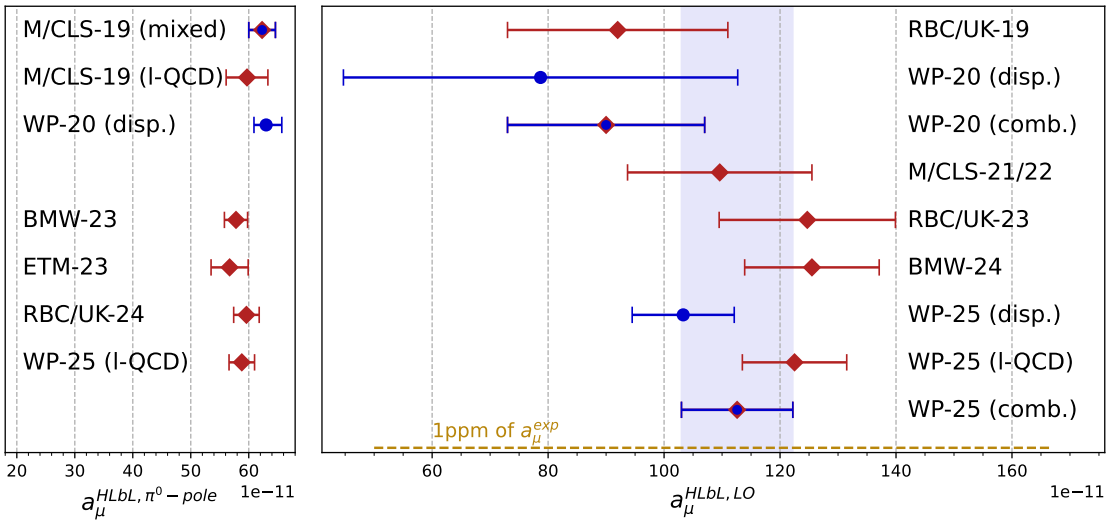


Figure 2.7: $a_\mu^{\mathrm{HLbL},\pi^0\text{-pole}}$ and $a_\mu^{\mathrm{HLbL,LO}}$. Summary of current values for $a_\mu^{\mathrm{HLbL},\pi^0\text{-pole}}$ (left) and $a_\mu^{\mathrm{HLbL,LO}}$ (right). Blue circles show data-driven results, red diamonds show l-QCD calculation results, and red diamonds overlayed with blue circles indicate 'mixed' results. The Mainz/CLS collaboration in 2019 [80] reported a result which uses a form factor normalization obtained from experimental data, denoted here as "M/CLS-19 (mixed)", in addition to their pure lattice result. Likewise, the combined values from both the 2020 [20] and the 2025 Theory Initiative review [16] average dispersive and lattice results and are thus 'mixed' results. The blue transparent band indicates the most recent combined value, which all previous results agree with within errors. Also shown is 1 ppm of the current experimental world average [5] for reference (dashed brown line).

The Theory Initiative's most recent recommended values for the LO HLbL contribution [16, 20, 120, 119, 58, 96, 80, 35, 57, 56, 39],

$$a_\mu^{\mathrm{HLbL,LO}} = 112.6(9.6) \times 10^{-11}, \tag{2.41}$$

is based on a combination of dispersive evaluations of the intermediate states - with transition form factors of the pion, η, and η', pair production measurements of heavier pesudoscalar mesons [20], meson decay data, and even hadronic cross-section data as experimental input [16] - and the average of four l-QCD results produced by the RBC/UKQCD collaboration in 2019 (RBC/UK-19 [39]) and 2023 (RBC/UK-23 [40]), the Mainz/CLS collaboration in 2021/2022 (M/CLS-21/22 [48, 49], and the BMW collaboration in 2024 (BMW-24 [75]). Fig. 2.7 shows a comparison of the available values for the total LO HLbL term, as well as the dominant pion-pole contribution. On the data-driven side, the dispersive result for the pion-pole has not changed since the Theory Initiative's 2020 publication [20], but major advancements have been made e.g. in calculating other, heavier-mass pseudoscalar contributions, leading to a new value for $a_\mu^{\mathrm{HLbL,LO}}$. Lattice-QCD calculations typically don't rely on calculating exclusive state contributions, but rather quark-flavor terms similar to the HVP contributions. Nevertheless some groups did calculate the pion-pole contribution explicitly, like the Mainz/CLS collaboration in 2019 (M/CLS-19 [80]), the BMW collaboration in 2023 (BMW-23 [79]), the ETM collaboration also in 2023 (ETM-23 [14]), and the RBC/UKQCD collaboration in 2024 (RBC/UK-24 [114]). Good agreement is observed between data-driven and l-QCD results across the board [16], as can be seen in fig. 2.7.

The NLO term is determined purely in the dispersive framework as [56, 16]

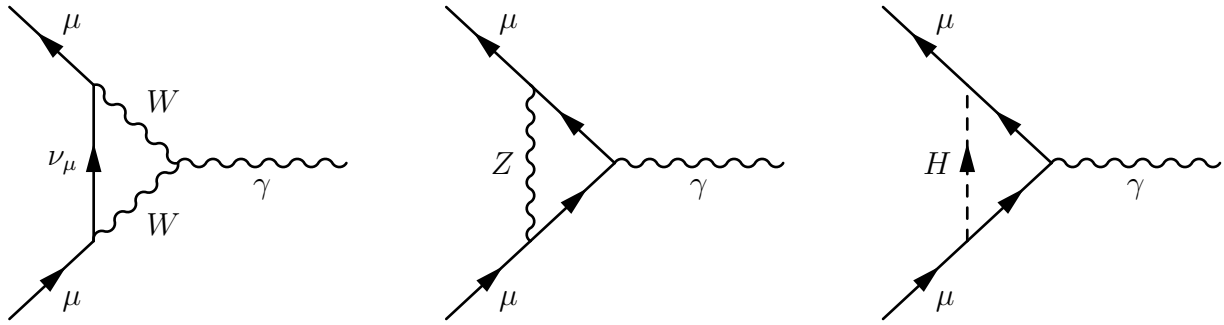
$$a_\mu^{\mathrm{HLbL,NLO}} = 2.8(6) \times 10^{-11}, \tag{2.42}$$

such that the total HLbL contribution becomes [16]

$$a_\mu^{\mathrm{HLbL}} = 115.5(9.9) \times 10^{-11}. \tag{2.43}$$

2.2.3 Electroweak contribution

Electroweak corrections include all interactions with at least one W, Z or Higgs boson, i.e. all interactions that are not purely QED or QCD contributions. One-loop diagrams are shown in fig. 2.8. The EW contribution is the smallest contribution, due to a suppression factor of $\frac{m_\mu^2}{M_W^2} \approx 10^{-6}$, where m_μ denotes the muon mass and M_W the W-boson mass [16].



(a) 1-loop electroweak vertex function with charged vector boson (W^+, W^- depending on charge of muon)

(b) 1-loop electroweak vertex function with neutral vector boson (Z^0)

(c) 1-loop electroweak vertex function with Higgs boson

Figure 2.8: Electroweak sector Feynman diagrams: Feynman diagrams of the leading order contributions to a_μ^{EW}

Two-loop diagrams with heavy particles (like the top-quark or Higgs-boson) and a photon lead to logarithmically enhanced corrections, which are large ($\sim 20\%$ from enhancement factors of the form $\log(\frac{M_Z^2}{m_\mu^2})$, with the Z-boson mass M_Z) compared to the one-loop contribution [20]. Additionally, non-perturbative hadronic corrections become important at the two-loop order as well. These can be evaluated with e.g. dispersion relations similar to the HVP and HLbL terms.

The EW contribution has been evaluated fully up to two-loop level, and leading-order three-loop effects have been estimated using effective field theory and renormalization group methods, which results in [61, 85]

$$a_\mu^{\mathrm{EW}} = 154.4(4) \cdot 10^{-11}. \tag{2.44}$$

The uncertainty is dominated by the hadronic terms from two-loop diagrams, and while it is an order of magnitude larger than the QED uncertainty, it is still insignificant compared to the purely hadronic sector as well as the experimental precision.

2.3 Current status of theory

Table 2.1 summarizes the currently published individual contributions to a_μ^{SM} that have been discussed in the previous sections, as well as the total value reported by the Theory Initiative

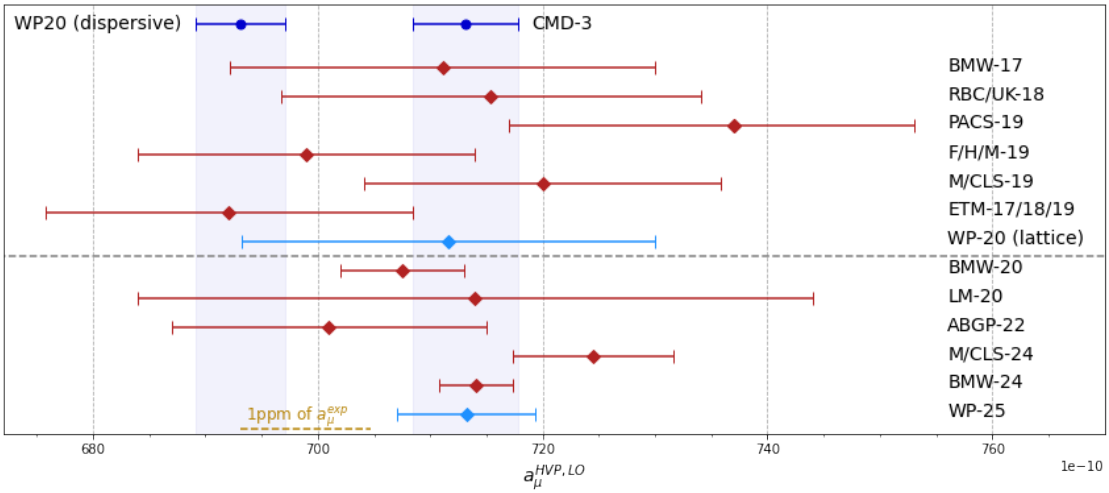


Figure 2.9: $a_\mu^{\mathrm{HVP,LO}}$ from lattice-QCD: Summary of lattice-QCD results for $a_\mu^{\mathrm{HVP,LO}}$ (red diamonds), compared to the Theory Initiative's 2020 data-driven recommended value WP20 [20] (blue circle), and the relatively new CMD-3 result [99]. Data points above the dashed horizontal line were available at the time of WP-20 [20], data points below were only available for the 2025 review WP-25 [16]. The combined values based on l-QCD results determined by the Theory Initiative are shown in light blue for both 2020 and 2025. Note that the BMW-24 result [41] has not been published yet and does not enter the WP-25 average [16]. Also shown is 1 ppm of the current experimental world average [7] for reference (dashed brown line).

for a_μ^{SM} [16],

$$a_\mu^{\mathrm{SM}} = 116\ 592\ 033(62) \cdot 10^{-11}. \tag{2.45}$$

The hadronic contributions are responsible for the majority of the uncertainty, in particular the leading-order HVP term, which is also the main reason for the significant change of the recommended SM value for a_μ from 2020 to 2025. A summary of l-QCD results for the total LO HVP contribution to a_μ^{SM} (including the l-QCD averages of the Theory Initiative from their 2020 [20] and 2025 [16] publications) is shown in fig. 2.9 vis-à-vis the dispersive average determined in 2020 [20] and the recent CMD-3 result [99][5]. While the 2020 average of available lattice results was still technically in agreement with the dispersive average due to the former's large uncertainty, the 2025 l-QCD average is now 2.8σ larger than the dispersive average from 2020 but shows very good agreement with the CMD-3 result. Further improvements in the precision of l-QCD results is expected in the future, as the focus of the lattice community shifts towards LD window contributions and the IB corrections [16].

[5]For the data point shown here, the two-pion contribution in the WP20 dispersive average was replaced with the CMD-3 result.

Nevertheless, it will be crucial to also investigate and understand the discrepancies among dispersive results. A 2023 measurement of higher order radiative effects by BABAR revealed that one of the Monte Carlo simulators used in some of the previous experiments, including KLOE and BESIII, does not accurately predict these higher order effects [112, 64, 66], which suggests the existence of systematic effects not included in the published uncertainties. The recently formed RadioMonteCarLow 2 working group has begun a dedicated effort to evaluate and improve current event generators for hadron and lepton production at low-energy e^+e^- colliders [17], which will be especially important to better understand and control radiative corrections [16]. Additionally, several major collaborations are expected to release updated results from larger datasets in the near future - CMD-3, SND, BaBar, BESIII, KLOE, and Belle-II [16].
Belle-II is also preparing to perform a new measurement of hadronic τ-decay with significantly larger statistics than any previous measurements [16]. While a combination of τ-decay data for the dominant two-pion channel and e^+e^- data for the other channels is in principle already possible (and might seem like a way to avoid the discrepancies in the two-oion channel), the τ-data based prediction of $a_\mu^{\pi^+\pi^-,LO}$ does at present not achieve the required precision; to yield a competitive result, IB corrections have to be well understood and included. Efforts for this are ongoing in both l-QCD and dispersive approaches. Furthermore, new measurements of other channels also show tensions between e^+e^- annihilation experiments (albeit on a significantly smaller scale than in the two-pion channel), which would call the reliability of such a combined result nonetheless into doubt.

With the updated SM value of a_μ from eq. 2.45 and the current experimental world average $a_\mu^{\text{exp}} = 116\,592\,071.5(14.5) \cdot 10^{-11}$ [7] (cf. table 1.1), theory and experiment are only separated by 0.6σ, which indicates agreement. Fig. 2.10 shows a comparison of a_μ^{exp} with the current theory value a_μ^{SM} (WP 2025) and its predecessor (WP 2020), as well as select individual results. There would be a 5.8σ discrepancy between theory and experiment if the WP 2020 value were still in effect.

	result (uncertainty) [$\times 10^{-11}$]	**uncertainty** [ppb of a_μ^{SM}]
a_μ^{QED}	116 584 718.8 (2)	1.7
a_μ^{HVP}	7 045 (61)	523
a_μ^{HLbL}	115.5 (9.9)	84.9
a_μ^{EW}	154.4 (4)	3.4
total a_μ^{SM}	116 592 033 (62)	532

Table 2.1: Summary of individual contributions to a_μ^{SM}, with uncertainties, reviewed and reported by [16].

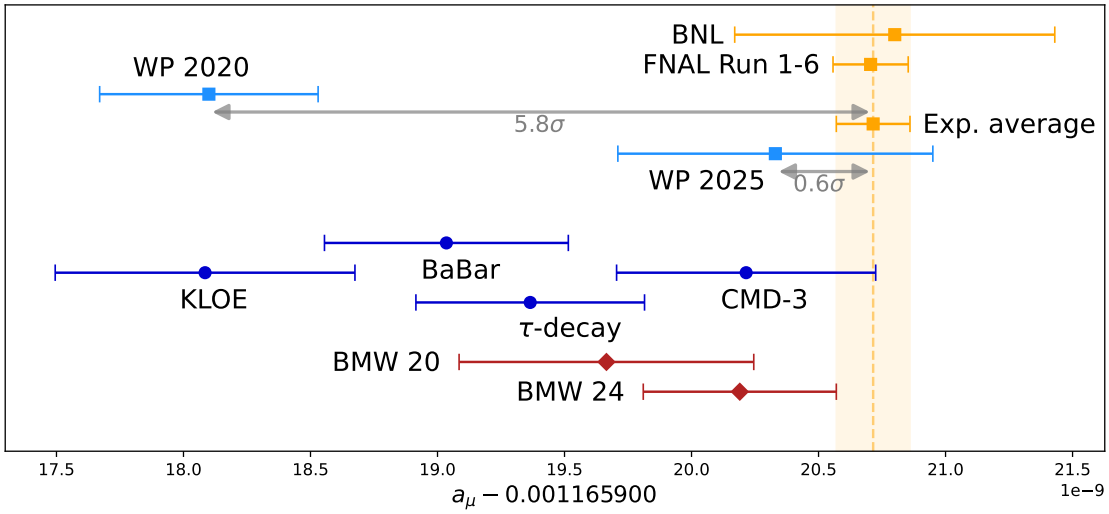


Figure 2.10: Comparison of a_μ^{exp} with select theory predictions, namely the currently accepted theory value a_μ^{SM} (WP 2025 [16]), the previous theory value from 2020 (WP 2020 [20]), as well as theory predictions using $a_\mu^{\mathrm{HVP,LO}}$ based on select individual results. Orange squares denote experimental results (BNL [31], FNAL Run 1-6 [7], and combined experimental world average); blue circles are data-driven theory predictions (BaBar [111], KLOE [18] and CMD-3 [99] from e^+e^- cross-sections, and a τ-decay data result [64]); and red diamonds mark lattice-QCD results (BMW 2020 [44], BMW 2024 [41]). BaBar, KLOE, CMD-3, and τ-decay predictions are calculated based on the respective experimental result in its available energy range, combined with an average of other available results outside that range to get a prediction for $a_\mu^{\mathrm{HVP,LO}}$. Values for a_μ^{SM} based on these four results have been taken from [64].

CHAPTER 3

Experimental Overview

The Fermilab Muon g-2 experiment measures the muon's anomalous magnetic moment a_μ (cf. eq. 2.13) by observing the decay positrons of positively charged muons orbiting in a storage ring. This chapter will provide details of the measurement principle (section 3.2.1), the muon production and delivery (section 3.1.1) as well as the physical components of the experimental apparatus: the magnetic storage ring including the electrostatic quadrupoles needed for focusing (section 3.1.2), the detectors that record the muon decay products (section 3.2.2), and the magnetic field measurement system (section 3.2.3).

3.1 Experimental setup

3.1.1 Muon production and delivery

Fig. 3.1 shows a schematic of the Fermilab Muon Campus, which is comprised of all facilities between (and including) the proton storage ring (the "Recycler") and the Muon g-2 storage ring. Pulses of 10^{12} protons ("bunches") with energy of 8.89 GeV and separated by at least 10 ms are created in the Recycler, guided through the P1, P2, and M1 beamlines, and hit the Inconel production target inside the AP0 target hall. This produces a secondary beam of pions, protons, positrons, and deuterons, with an energy of 3.11 GeV $\pm \sim 10\%$. [76]
The secondary beam travels through the M2 and M3 beamlines to the delivery ring, during which time 70% of the pions are expected to decay. The dominant decay channel for pions is $\pi^+ \to \mu^+ + \nu_\mu$ (branching ratio $> 99.98\%$ [125]), depicted in fig. 3.2. Since it is mediated by the parity-violating weak force, the neutrino as an almost massless particle will be left-handed. The pion is a spin 0 particle and angular momentum conservation thus requires the μ^+ to also be left-handed (in the center-of-mass frame). By capturing the highest energy muons, which are boosted along the pion's trajectory, 95% polarization of these muons anti-parallel to their momenta is achieved [76]. The secondary beam is then sent into the "Delivery Ring", where it is stored for four turns. Only protons and muons are successfully

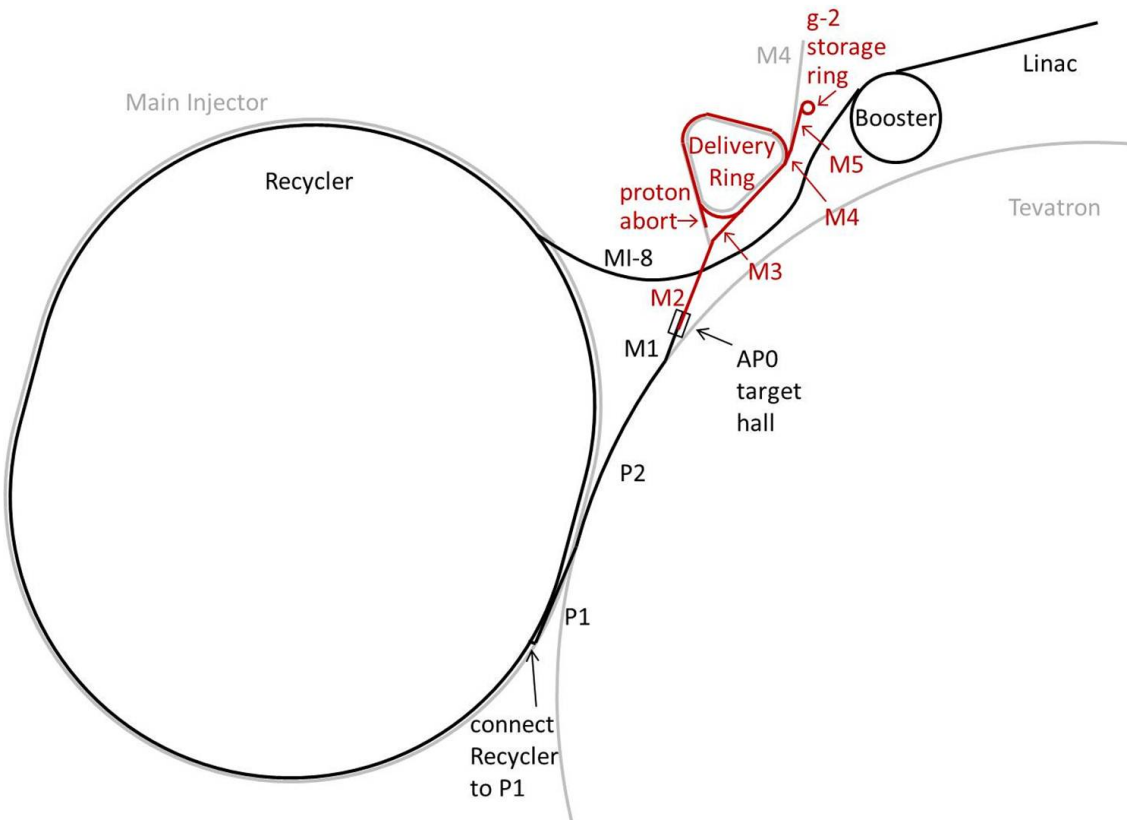


Figure 3.1: Schematic of the Fermilab Muon Campus, reproduced from [88].

stored, and the protons will separate from the muons by time of flight and are kicked off before the muons are extracted into the M5 beamline [76]. A wedge absorber system in the M5 beamline further reduces the momentum spread of the muons, such that at the entrance to the Muon g-2 storage ring, the muon beam has a central momentum of $p_0 = 3.094\,\text{GeV}$ with a $\pm 1.6\%$ spread [9], and a contamination fraction of less than 2% [110].

3.1.2 Muon storage

The Muon g-2 magnetic storage ring consists of a toroidal vacuum chamber with major radius $R_0 = 7.112\,\text{m}$, surrounded by a C-shaped superconducting electromagnet producing a nominal dipole field of $B_0 = 1.4513\,\text{T}$ in the vertical (y) direction [88]. The opening of the magnet is towards the inside of the ring, where the detectors are located. Fig. 3.4 shows a cross-sectional view of this setup. Fig. 3.5 depicts a bird-eye view of the storage ring, indicating many of the experiment's subsystems.

The muons are delivered into the storage ring through a superconducting inflector magnet

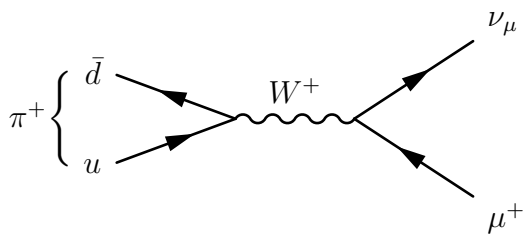


Figure 3.2: Dominant pion decay channel

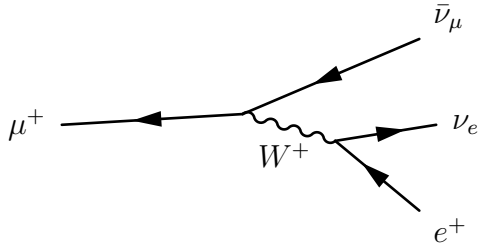


Figure 3.3: Muon decay

that passes through one of the ring magnet's iron yokes. To ensure the incoming muons can be efficiently transported into the storage volume without being deflected by the ring's magnetic field, the inflector magnet is designed to almost perfectly cancel the storage ring's fringe field inside the yoke and provide a field-free path for the muons [88]. To avoid interference with any muons currently circulating in the ring, the injection point is offset from the center of the storage volume by 77 mm radially outward, and delivers the muons onto a path tangential to the central orbit. The injected muons will intersect the central orbit after about a quarter turn, at which point there are three "kickers", pulsed electromagnets, located (K1-K3 in fig. 3.5). At the intersection with the central orbit, the injected muons have a small momentum component radially inwards. These kickers produce a vertical magnetic field for a very brief amount of time, which causes a momentum shift radially outwards that compensates the inward component and "kicks" the muons onto the central orbit [9].

The active storage volume of the Muon g-2 storage ring has a circular cross-section with minor radius $r_0 = 45\,\text{mm}$, and stores muons within $\pm 0.15\%$ of the ideal momentum p_0 [9]. Five ring-shaped collimators made of copper are placed along the ring (indicated with "C" in fig. 3.5) to slow down (and quickly lose) muons outside of the storage radius r_0. In order to create a highly uniform magnetic field inside the storage volume, both active and passive shimming techniques are employed. Passive shimming, during which pieces of various sizes and materials are precisely adjusted, was performed initially during the assembly of the experiment, before data was taken. Adjustable components, which are indicated in fig. 3.4, are the top hats (long steel pieces on top of/below the yokes), pole pieces (iron segments directly above and below the storage volume), steel wedges between the yokes and pole pieces, as well as "edge shims" (iron foil laminations on the pole faces). With passive shimming alone, the azimuthal variations of the dipole field had been reduced to 25 ppm [126]. Active shimming is an ongoing process during data-taking, which uses the "surface coils" (two sets of 100 concentric coils, located on each of the pole faces; also shown

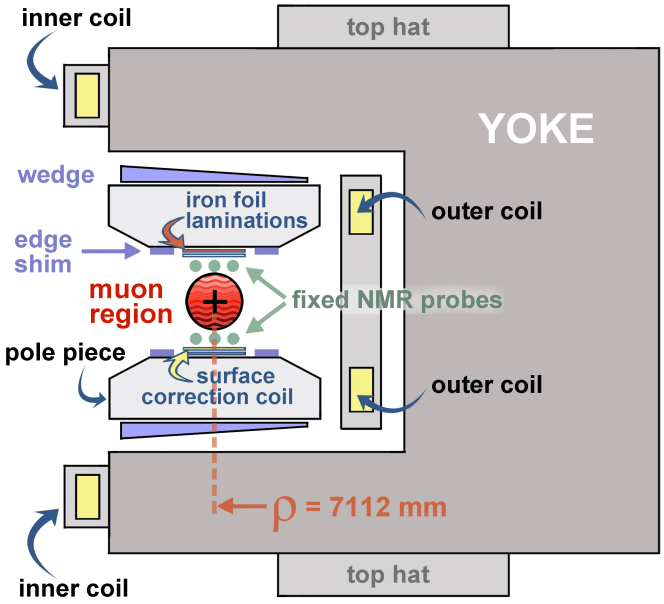


Figure 3.4: Cross-section of the Muon g-2 magnetic storage ring, reproduced from [10]. C-shaped iron yokes are magnetized by superconducting electromagnetic coils to produce the vertical magnetic field B_0. The center of the storage ring is to the left.

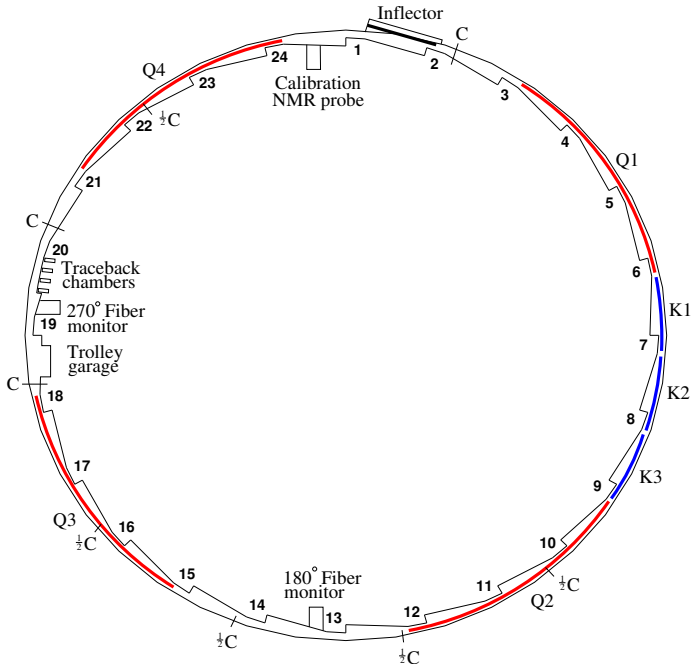


Figure 3.5: Top-down schematic view of the Muon g-2 storage ring, reproduced from [88]. Indicated are, among others, the electrostatic quadrupoles in red (labeled "Q"), the kicker magnets in blue (labeled "K"), and the locations of the traceback chambers behind the fiber monitors at about 180° and 270° from the inflector.

in fig 3.4) to counteract any remaining spatial field gradients and inhomogeneities, and a power supply feedback loop that adjusts the current of the main magnet in response to field drifts over time. Active and passive shimming together resulted in a ~ 14 ppm spatial homogeneity across the storage volume, and less than 1 ppm when azimuthally averaged [10].

Once the muons have been injected into the ring, the magnetic field inherently offers horizontal focusing. Vertical focusing is provided by four electrostatic quadrupoles (ESQs) placed symmetrically around the ring as seen in fig. 3.5. Each quadrupole section consist of a short (13°) and a long (26°) segment, with a high voltage (HV) electrode plate on all four sides of the ring (top and bottom at positive voltage, inside and outside at negative voltage) [88]. In Run-4/5/6, the typical ESQ voltage was 18.2 kV. The ESQ plates stay charged for the entire measurement period for one muon bunch, called a "fill", which lasts about 700 µs. In addition to vertical focusing, the ESQs are also used for "scraping" the beam in the first ~ 20 µs, a procedure during which the ESQ plates are asymmetrically and the beam shifted (vertically and horizontally) off-center. This brings the edges of the beam outside of the active storage volume and "scrapes" the highest- and lowest-energy muons away by scattering on the collimators [88]. The ESQ plate voltages are brought to their nominal setpoints and the beam returns to its equilibrium position within 30 µs after injection.

3.2 Methodology

3.2.1 Measurement principle

Even though muons are quantum particles, it is helpful to treat the polarization of the ensemble of muons in a semi-classical way, combining equations 2.5 and 2.8 from section 2.1. Since the vast majority of muons enter the storage ring with their spin anti-parallel to their momentum, both of which are perpendicular to the (vertical) magnetic field, one can observe Larmor precession of the polarization vector. In the ideal (i.e. all muons are on the central orbit, no vertical focusing is necessary), non-relativistic case, the spin precession frequency $\vec{\omega}_s$ and cyclotron frequency $\vec{\omega}_c$ are

$$\vec{\omega}_s = -\frac{gq}{2m}\vec{B}, \tag{3.1}$$

$$\vec{\omega}_c = -\frac{q}{m}\vec{B}, \tag{3.2}$$

and the "anomalous frequency", defined as the rate of dephasing between polarization and momentum vector, can be written as

$$\vec{\omega}_a \equiv \vec{\omega}_s - \vec{\omega}_c = -\frac{q}{m}\left(\frac{g-2}{2}\right)\vec{B} = -a_\mu \frac{q}{m}\vec{B}. \tag{3.3}$$

This illustrates the basic measurement principle of the experiment: Because the anomalous frequency is directly proportional to a_μ, the magnetic field, and natural constants, we can extract the muon's anomalous magnetic moment from precise measurements of $\vec{B}$ and $\vec{\omega}_a$.
In practice, with the addition of the electrostatic quadrupole fields and including relativistic effects, spin precession becomes a lot more complicated. The Thomas-Bargmann-Michel-Telegdi (Thomas-BMT) equation describes spin precession of the muon spin $\vec{s}$ in the lab frame [100],

$$\frac{\mathrm{d}\vec{s}}{\mathrm{d}t} = \frac{q}{m}\vec{s} \times \left[(\frac{g}{2} - 1 + \frac{1}{\gamma})\vec{B} - (\frac{g}{2} - 1)\frac{\gamma}{\gamma+1}(\vec{\beta}\cdot\vec{B})\vec{\beta} - (\frac{g}{2} - \frac{\gamma}{\gamma+1})\vec{\beta}\times\vec{E}\right]. \tag{3.4}$$

Included here are the full relativistic magnetic and electric fields ($\vec{B}$ and $\vec{E}$ respectively, with the relativistic Lorentz factor $\gamma = \frac{1}{\sqrt{1-\beta^2}}$ and fractional velocity $\beta = \frac{v}{c}$ where v is velocity and c is speed of light, which leads to the $(\vec{\beta}\cdot\vec{B})\vec{\beta}$ and $\vec{\beta}\times\vec{E}$ terms, as well as a correction due to the Thomas precession, $\vec{\omega}_T = \frac{\gamma^2}{\gamma+1}(\frac{\mathrm{d}\vec{\beta}}{\mathrm{d}t}\times\vec{\beta})$, which gives rise to the $(\frac{1}{\gamma}-1)\vec{B}$ term. Then the anomalous frequency becomes

$$\vec{\omega}_a = -\frac{q}{m}\left[a_\mu\vec{B} - a_\mu\frac{\gamma}{\gamma+1}(\vec{\beta}\cdot\vec{B})\vec{\beta} - (a_\mu - \frac{1}{\gamma^2-1})\vec{\beta}\times\vec{E}\right]. \tag{3.5}$$

The muon momentum $p_0 = 3.094\,\mathrm{GeV}$ ("magic momentum"), corresponding to $\gamma \simeq 29.3$, has been chosen so that the motional magnetic field term $(\vec{\beta}\times\vec{E})$ will vanish, as a_μ and $\frac{1}{\gamma^2-1}$ cancel each other [88]. A first-order perturbative correction, the "electric field correction", is applied due to the muon beam's finite momentum spread around p_0.
The $\vec{\beta}\cdot\vec{B}$ term only contributes for muons that have a momentum component parallel to the magnetic field. The electrostatic vertical focusing introduces vertical betatron motion for muons above or below the central orbit, but since the vertical velocities are small compared to the azimuthal relativistic motion, it is equally treated as a perturbative effect, and a "pitch correction" is applied to the measurement result.
Additionally, so-called coherent betatron oscillations (CBOs) - betatron motion of the beam centroid, i.e. many muons moving in phase - and muon losses can be reduced by applying radio frequency (RF) electric fields to the ESQ plates for the first 6 µs after injection [103, 7].

This was implemented during early Run-5 for the radial CBO, and midway through Run-5 for the vertical CBO as well. Analysis datasets for Run-4/5/6 have been defined to align with the changes in RF configuration ("RF datasets") and are defined in appendix A, table A.2.

While the magnetic field measurement will be discussed in more detail in section 3.2.3 and chapter 4, it is relevant to note that the magnetometers are pulsed nuclear magnetic resonance (NMR) probes. Thus the physical quantity measured is not the magnetic field $\vec{B}$ itself, but the Larmor precession frequency of the (shielded) proton samples via observation of free induction decay (FID) inside the NMR probes. The proton spin precession frequency can be defined analogously to eq. 2.8,

$$\omega_p = \frac{g_p q_p}{2m_p} B = \frac{2\mu_p}{\hbar} B, \tag{3.6}$$

where g_p is the proton g-factor, $q_p = e$ is the proton charge, m_p is the proton mass, and $\mu_p = \frac{g_p q_p \hbar}{4m_p}$ is the proton magnetic moment defined as in eq. 2.6 with the proton spin $S = \frac{\hbar}{2}$. Furthermore, as ω_a is a statistical average of all detected muons, but the magnetic field is measured at discrete spatial positions, a procedure called "muon weighting" is employed to find the average magnetic field encountered by the muons (explained in depth in chapter 5). Muon-weighted quantities are indicated by a tilde, e.g. the muon-weighted magnetic field $\tilde{B}$. Taking into account the muon-weighted magnetic field, and the use of *shielded* proton NMR probes (indicated by the primed quantities ω_p', μ_p'), the muon anomalous magnetic moment can be calculated as

$$a_\mu = \frac{\omega_a}{\tilde{B}} \frac{m_\mu}{e} = \frac{\omega_a}{\tilde{\omega}_p'} \frac{2\mu_p'}{\hbar} \frac{m_\mu}{e} = \frac{\omega_a}{\tilde{\omega}_p'(T_r)} \frac{\mu_p'(T_r)}{\mu_e(H)} \frac{\mu_e(H)}{\mu_e} \frac{m_\mu}{m_e} \frac{g_e}{2}, \tag{3.7}$$

with the electron magnetic moment $\mu_e = g_e \frac{e\hbar}{4m_e}$, electron mass m_e, electron g-factor g_e, and reference temperature $T_r = 34.7\,^\circ\mathrm{C}$ [6]. Ratios of magnetic moments are, in many cases, known with much higher precision than their absolute values or the muon's charge-to-mass ratio. Table 3.1 summarizes the uncertainties and sources for the inputs to eq. 3.7 that are not measured by the Fermilab Muon g-2 experiment.

3.2.2 Measurement of ω_a

Muons have a mean lifetime of $\tau_\mu \simeq 2.2\,\mu\mathrm{s}$ [125], which dilates to about 64 µs in the lab frame, and are too short-lived to be detected directly. However, they decay nearly exclusively in a

Quantity	Description	Uncertainty	Source
$\mu'_p(T_r)/\mu_e(H)$	magnetic moment ratio of proton in a spherical water sample at 34.7 °C and electron in a hydrogen atom	10.5 ppb	[129]
$\mu_e(H)/\mu_e$	magnetic moment ratio of e^- in a hydrogen atom and free e^- in vacuum	0	[101]
m_μ/m_e	muon mass to electron mass ratio	22 ppb	[121]
g_e	electron g-factor	$1.2 \cdot 10^{-13}$	[125]
Total	all external parameters	25 ppb	[6]

Table 3.1: External parameters used to calculate a_μ. Description and uncertainty of externally sources ratios and constants used to calculate a_μ.

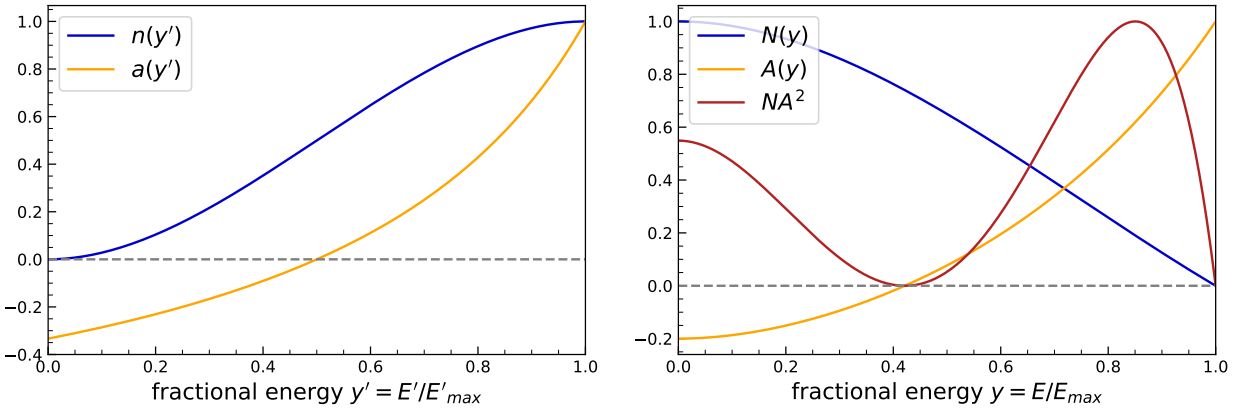
single channel [125] into a positron[1] and two neutrinos,

$$\mu^+ \to \bar{\nu}_\mu + \nu_e + e^+, \tag{3.8}$$

shown in fig. 3.3. The muon decay is, like the pion decay, a weak force interaction and requires that the muon anti-neutrino be right-handed, and the electron neutrino be left-handed. In the center-of-mass frame, the positron will have the biggest momentum when the neutrinos are emitted parallel to each other and opposite of the positron. In that case, the neutrinos have a combined spin of 0 and the positron carries the spin of the muon. Since the positron preferentially has right-handed helicity[2], the majority of high-energy positrons will be emitted along the muon's spin direction. Conversely, the positron will have the lowest momentum when the neutrinos are emitted back to back, carrying a total spin of 1, and the positron will preferentially be emitted opposite the muon's spin [88]. Hence the positron momentum direction is effectively correlated with the muon's spin at the time of decay. These relationships are illustrated in 3.6a, which shows the asymmetry (a) and relative number (n) of decay positrons depending on the normalized electron energy y' in the muon's rest frame.

[1]Since the Muon g-2 experiment measures exclusively positively charged muons, this work will use the corresponding terminology and refer to the decay products as positrons, electron neutrinos, and muon anti-neutrinos. Negatively charged muons of course produce electrons, electron anti-neutrinos, and muon neutrinos.

[2]Weakly mediated interactions enforce chirality, but massive leptons can have a helicity different from their chirality. The mass term in the Lagrangian is responsible for mixing between the left- and right-handed states, and this effect is more pronounced for heavier particles [128]. This is why in the pion decay, the μ^+ is able to have left-handed helicity, even though it is an anti-particle. In muon decay however, because the positron is about 200 times lighter than the muon, it is more likely to have right-handed helicity.



(a) Center-of-mass frame

(b) Lab frame

Figure 3.6: Positron number and asymmetry distributions, and figure-of-merit. Relative number (n,N) and asymmetry (a,A) distributions versus fractional energy (y',y) in the center-of-mass frame (left) and lab frame (right), reproduced from [88]. In the center-of-mass frame, the maximum energy $E'_{max} \approx m_\mu/2$, and in the lab frame the maximum energy $E_{max} \approx 3.1\,\text{GeV}$. The figure-of-merit NA^2, which determines the statistical uncertainty of ω_a, is also displayed. Distributions have been normalized for visibility.

These quantities can be understood with the equation for the differential probability in the center-of-mass frame of a positron emerging with energy fraction $y' = E'_e/E'_{e,max}$ of the maximum possible positron energy $E'_{e,max} = \frac{\sqrt{m_\mu^2+m_e^2}}{2m_\mu} \approx m_\mu/2$ and at angle ϑ with respect to the muon's spin direction [88],

$$\frac{\mathrm{d}^2P(y',\vartheta)}{\mathrm{d}y'\mathrm{d}\Omega} \propto n(y')[1 + a(y')\cos\vartheta]. \tag{3.9}$$

The relative number distribution $n(y')$ gives the overall likelihood of a positron with fractional energy y' being emitted, while the asymmetry $a(y')$ is a measure of the probability that a positron at energy y' is emitted along the muon spin direction. It is clear that since both n and a are large for large y', the majority of high-energy positrons will be emitted along the muon spin direction.

Only positrons with a significant momentum component parallel to the muon momentum are boosted enough to still have large fractional energy y in the lab frame. Fig. 3.6b shows the relative number density $N(y)$ and the asymmetry $A(y)$ in the lab frame as a function of lab frame fractional energy y. The maximum possible energy in the lab frame is equivalent to the muons with the highest momentum, $E_{max} = E_0 + 0.15\% \approx 3.1\,\text{GeV}$.

As the muon spin is precessing faster than the cyclotron frequency, the number of decay positrons above a certain threshold energy E_{thr} (at a fixed angle with respect to the muon's

momentum direction) will be modulated with ω_a. The expected rate is given by [88]

$$N(t, E_{thr}) = N_0(E_{thr}) \exp^{-t/\gamma\tau_\mu}[1 + A(E_{thr}) \cos(\omega_a t + \phi(E_{thr}))], \tag{3.10}$$

where N_0 is the positron number at the start of the measurement, A is the lab frame asymmetry, $\gamma\tau_\mu = 64.4\,\mu s$ is the muon lifetime in the lab frame, and ϕ is the initial muon spin phase. The angle between positron momentum and muon spin from eq. 3.9 is precisely the accumulated anomalous phase, $\vartheta = \omega_a t$. N_0, A and ϕ will depend on E_{thr}. This is the basic, five-parameter fit function with which ω_a can be determined from the positron decay spectrum [88]. A correction for muon losses, which are higher at early times and can skew the distribution, has to be made. Additionally, a "differential decay" correction, which accounts for a correlation between the muon lifetime and its momentum that leads to a time-dependence in the phase ϕ, and a "phase-acceptance" correction, which resolves time-dependencies of the phase ϕ due to transverse beam profile changes, have to be applied. The statistical error on the anomalous frequency when fitting with eq. 3.10, $\delta\omega_a/\omega_a$, is inversely proportional to $\sqrt{NA^2}$ [88]. When choosing E_{thr} it is therefore important to try and maximize this figure-of-merit (which is also depicted in fig. 3.6b), in order to minimize the statistical error.

Fig. 3.7 illustrates a typical measurement result. The "wiggle plot" shows the number of detected positrons (above E_{thr}, and binned per cyclotron period $T_c = 149.2\,ns$) over time. The number oscillates with the anomalous frequency, and an overall decrease over time is visible, which corresponds to the exponential decay with lifetime τ_μ. After fitting the wiggle plot with the five-parameter function, a fast Fourier transform (FFT) of the residuals shows no spike at the ω_a frequency value, indicating a successful fit. Other identifiable frequencies in the FFT are all well understood, like the CBO frequency.

The calorimeters

The positrons are measured by 24 electromagnetic calorimeters [88, 9], located evenly spaced on the inner wall of the vacuum chamber (see fig. 3.5). Due to the positron's lower mass compared to the muon, its momentum even for high-energy positrons is too low to stay on a stable orbit in the ring, and it will curl inwards into the detector (shown in fig. 3.8). The detector acceptance depends both on the positron energy (which has to be greater than the threshold energy of the detector) as well as on the angle of incidence.
Each calorimeter consists of a 6x9 array of $25 \times 25 \times 140\,mm^3$ lead fluoride (PbF_2) crystals. Positrons will shower and produce Cerenkov radiation proportional to the deposited energy

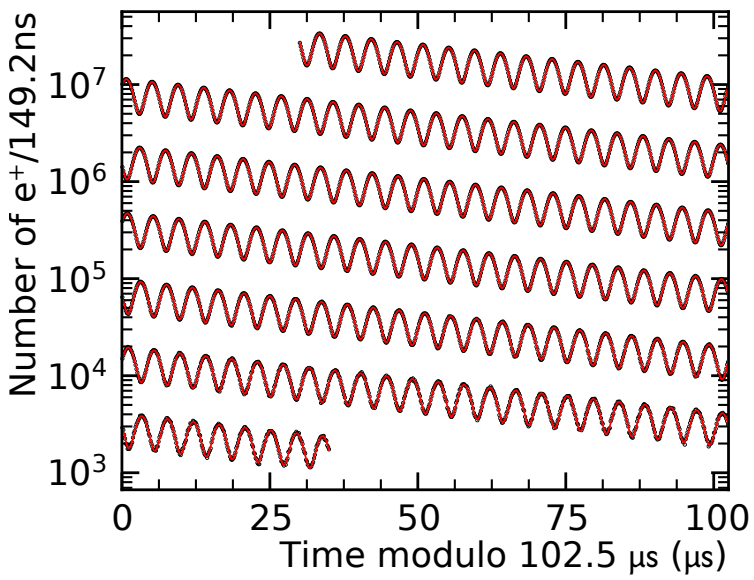


Figure 3.7: Wiggle plot: Representative result for a ω_a measurement, reproduced from [5], showing the "wiggle plot" of the Run-3a dataset, i.e. the number of detected positrons over time, binned per cyclotron period (149.2 ns) and asymmetry-weighted. The oscillation in the wiggle plot is precisely ω_a, with overall decreasing numbers due to the exponential decay of the muon beam.

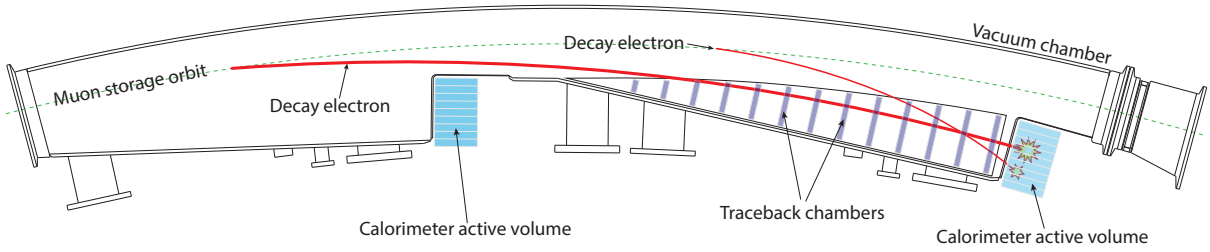


Figure 3.8: Schematic view of a decay positron hitting the detectors, reproduced from [88]. A positron's trajectory will be bent more strongly by the magnetic field in the storage ring due to its lower momentum compared to a muon, causing it to curve inward into the detector.

in the crystal, which is read out by silicon photo-multipliers (SiPMs). The calorimeters measure both the energy of the incident positrons as well as their arrival times in order to extract a frequency. Limited spatial information can be gained, since the crystal in which the hit occurs is also registered. Calibration of both gain and time are done via regular laser pulses with fixed intensity and wavelength. The front face of the calorimeters (facing the storage ring) also has another layer of scintillator detectors to account for muon losses. Hits in the front scintillators that aren't followed by showers in the crystal indicate a lost muon instead of a decay positron. The number of detected positron hits in the calorimeters is also called "CTAG" (**C**alorimeter **TAG**ged) count, or CTAGs for short.

The tracking detectors

The Fermilab Muon g-2 experiment also uses two traceback chambers [88, 9] to reconstruct the transverse muon beam profile as a function of time, and to reduce systematic uncertainties of the calorimeters due to pile-up by identifying times with multiple decay vertices close together. These traceback chambers, also called straw tracker detectors, are placed in front of a calorimeter (as shown in fig. 3.8) at about 180° and 270° from the inflector (cf. fig. 3.5). Each straw tracker consists of eight modules with two double layers ("planes") of straw chambers. The straws of these two planes are angled at 7.5° from the vertical opposite of each other, as seen in fig. 3.9. This design permits to collect some height information without the need for a dedicated layer of horizontal straws. Each straw has a fiducial length of 8.5 cm and a diameter of 5 mm, containing a 25 µm gold-plated tungsten wire at 1.5 kV and filled with a drift gas mixture of 50:50 Ar:C_2H_6 [136].
By connecting the hits in the different layers and modules, a tracking algorithm can determine the positron's momentum and extrapolate its trajectory back into the ring[3]. The muon decay vertex position is reconstructed by finding the point of radial tangency of the extrapolated positron trajectory to the central orbit. Through the two tracker stations, the transverse beam profile (i.e. a histogram of radial and vertical positions of muon decay vertices) can be obtained at an azimuthal point close to the tracker locations. Knowledge about the transverse beam profile is necessary both for systematic corrections to ω_a (e.g. the phase-acceptance correction), and as input to the muon weighting procedure, which will be described in detail in chapter 5.

The detector data acquisition (DAQ) system collects data for both calorimeters and trackers in time intervals of about 9 s ("subruns") and groups them into larger units spanning time periods on the order of one hour ("runs", not to be confused with the large scale data taking

[3]The extrapolation algorithm is discussed in more detail in chapter 6.

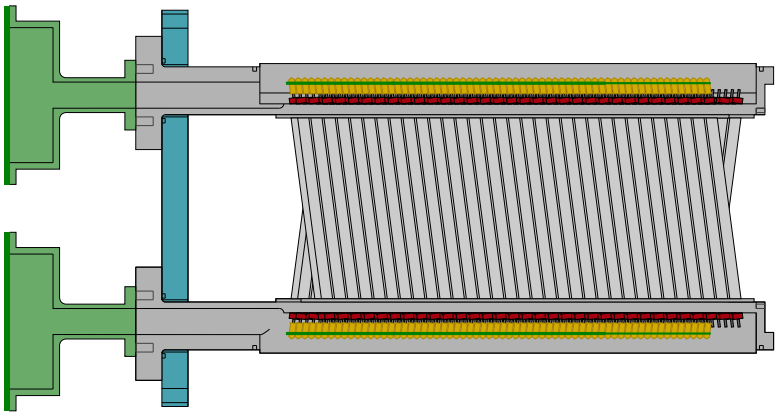


Figure 3.9: Side view of a straw tracker module. Schematic diagram of a straw tracker module, viewed from the side, reproduced from [88]. The closest two layers of straws are tilted 7.5° to the left, and slightly offset from each other, while the two layers towards the back are tilted 7.5° to the right.

periods of several months, Runs 1-6). While calorimeter and tracker data is stored separately, the grouping into subruns and runs is identical in structure and naming.

3.2.3 Measurement of ω_p

In order to measure a_μ, it is not only necessary to determine ω_a, but also to measure the magnetic field as precisely as possible. Because the NMR probes can't occupy any space inside the muon storage volume while beam is present (otherwise they are obstructing the muon beam's path), two types of probes are employed [10]: 378 stationary "fixed probes" at 72 azimuthal positions ("stations") along the ring just above and below the storage volume (indicated in fig. 3.4), that continuously monitor the magnetic field over time (as well as provide real-time feedback to the storage ring magnet's power supply for active shimming); and a "trolley" with 17 probes that can be moved inside the vacuum chambers along rails to create a high-resolution spatial map ($\sim$ 9000 azimuthal points) of the field inside the storage volume, which is done about every three days while beam is turned off [10]. Note that the field DAQ system is independent of the detector DAQ system, and field runs are not synchronized with the detector DAQ [6]. A cross-sectional view of the trolley probes and fixed probes in the storage ring is shown in fig. 3.10 and indicates the dimensions of the trolley and distances between probes.

As mentioned previously (eq. 3.7), the value that is needed to calculate a_μ is the shielded proton spin precession frequency at a reference temperature of $T_r = 34.7\,^\circ\mathrm{C}$ [6]. Both the trolley probes and the fixed probes use petroleum jelly samples, and have to be temperature-

corrected as well as (individually) carefully calibrated with a pure water sample [10]. Several other effects have to be taken into account and adjusted for during the analysis of the FID frequencies, such as [10, 8, 51]

- Magnetic field distortion inside the trolley due to material effects
- Motion effects in the trolley probes (mainly from Eddy currents in the trolley shell)
- Position variations of the trolley rails
- Electrical ground loops created by the trolley cable making contact with the kicker plates (affecting both the trolley and the fixed probes)
- Magnetic field distortions at the fixed probe locations due to the trolley material and Eddy currents (so-called "trolley footprint")

During each trolley run, the trolley probes and fixed probes are synchronized, since they are measuring the magnetic field at the same time and spatially close (after aforementioned corrections). This subsequently allows for the spatial trolley maps to be extrapolated in time, using the information of the fixed probes to correct for field drifts, in order to determine the field while the muon beam is present. A pair of trolley runs together with the data taken in between (both fixed probe data and ω_a data from the calorimeters and trackers) defines a "trolley pair dataset". Typically, extrapolating forward in time from one trolley run and backward from the following will not give exactly the same result. The discrepancy is called the "sync offset", and leads to another correction during the magnetic field analysis. The temporally extrapolated spatial field maps, called virtual trolley runs (VTRs), are then used together with information from the straw tracking detectors for the muon weighting of the magnetic field (described in more detail in chapter 5).

While efforts are made to keep conditions during trolley measurements the same as during data taking, this isn't entirely possible. Both the kickers and the ESQs have to be turned off, as they would interfere with the NMR measurements. Additionally, to make the movement of the trolley in the storage ring possible, two physical elements have to be brought into different configurations: The trolley is inserted into the ring on the "garage rails", which are not inside the ring during data taking; and the beam collimators, which help to eliminate muons outside the active storage volume, have to be retracted to allow the trolley to pass. In particular the magnetic kick during muon injection, and the ESQ pulses, produce transient perturbations to the magnetic field that cannot be measured with the trolley. Both of these effects are evaluated independently, and a correction is applied to

the muon weighted magnetic field.

As described in section 3.2.1, the measured frequency of the FID of the trolley and fixed probes is related to the magnetic field via fixed constants, thus the two can be used interchangeably to describe the field strength. The proton precession frequency at the nominal field strength $B_0 = 1.4513\,\mathrm{T}$ is $f_0 \simeq 61.79\,\mathrm{MHz}$. In the following chapters, the magnetic field strength is typically reported as a frequency $\frac{\omega_p}{2\pi}$ either directly in Hz, or relative to the reference frequency f_0 (in e.g. ppm or ppb).

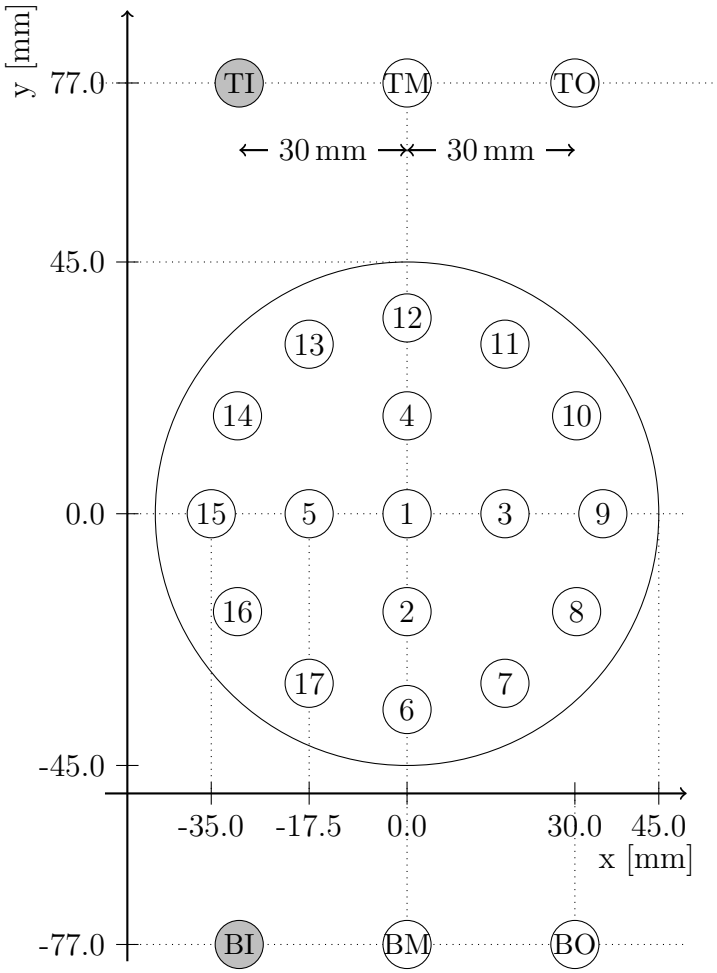


Figure 3.10: Cross-sectional view of trolley and fixed probe locations: Cross-sectional view of trolley probe (1-15) and fixed probe (TI/TM/TO/BI/BM/BO) locations, reproduced from [10]. The central probe of the trolley coincides with the central orbit of the muon storage ring. The fixed probes are located outside of the storage volume, 77 mm above and below the central orbit. Most fixed probe stations have 6 probes, 3 each on top (T) and bottom (B), at the inner (I), middle (M), and outer (O) positions. A few stations only have 4 fixed probes, 2 each on top and bottom in middle and outer positions. y indicates the vertical direction, x the radial direction.

CHAPTER 4

Spatial Magnetic Field Fitting

As was just discussed in section 3.2.3, any true magnetic field measurement *inside* the storage volume can not coincide with the muon beam. The magnetic field monitoring during ω_a measurements is provided by the fixed probes located outside the storage volume, which depends on a correlation/synchronization between trolley and fixed probe data. Additionally, it is also necessary to be able to spatially correlate the magnetic field with the beam shape for the muon weighting, detailed in chapter 5. Both of these tasks require a method to parameterize and fit the magnetic field that can provide reliable field value estimates at arbitrary positions within the storage ring.

Section 3.1.2 described the efforts to make the magnetic field inside the muon storage volume as close to a perfectly uniform vertical field as possible. Dedicated measurements using Hall probes were performed before the start of data-taking in Run-1 to determine the non-vertical components of the magnetic field [126]. Due to the physical requirements of the probes and their mount, these measurements could only be performed before the installation of the vacuum chambers, and therefore only passive shimming was present. The results of over 100 measured azimuthal locations showed that the magnitude of the radial field B_R was less than 300 ppm compared to B_0, and the longitudinal field B_ϕ varied less than 100 ppm in magnitude. The average azimuthal values were $\langle B_R/B_0\rangle_\phi = (68.5\pm68.9)$ ppm and $\langle B_\phi/B_0\rangle_\phi = (0.14 \pm 36.90)$ ppm [126]. Both non-vertical components were reduced further by applying active shimming ($\langle B_R/|\vec{B}|\rangle_\phi < 40$ ppm during Run-1 [10]), and in Run-4/5/6 the azimuthally averaged radial magnetic field was able to be reduced to less than 1 ppm [90].

The NMR probes used in both the trolley and the fixed probe stations on the other hand only measure the magnitude of the magnetic field, $|\vec{B}| = \sqrt{B_R^2 + B_y^2 + B_\phi^2}$. However, since the vertical field is much larger than the other components, the NMR measurements can be taken as a measurement of B_y to first order. Given the size of the non-vertical components with active shimming, the difference between the magnitude of the field and the vertical

component is on the order of 10 ppb (compared to B_0) [10],

$$|\vec{B}| - B_y \approx \frac{B_x^2 + B_z^2}{2B_y} \approx \mathcal{O}(10^{-8} B_0). \tag{4.1}$$

This justifies the approximation $|\vec{B}| \approx B_y$ which is used in the analysis and fitting procedure of the magnetic field. This chapter will first introduce in section 4.1 the 2-dimensional multipole parameterization of the magnetic field that has already been used in the Brookhaven experiment [31], and continues to be used in the standard magnetic field analysis at the Fermilab Muon g-2 experiment [2, 5, 7, 10, 140]. Then in section 4.3, a 3-dimensional parameterization in toroidal coordinates will be introduced, which is used to validate the use of the 2D multipoles and estimate the uncertainty introduced by the above approximation in section 4.5.

4.1 2D parameterization

Given that the main component of the muon momentum is along the azimuthal direction, the radial magnetic field perturbation (being perpendicular to the largest momentum component) will have a more noticeable effect on the beam than the longitudinal field (which only acts on the "perturbations" of the momentum). A natural choice of coordinate system is then to neglect the azimuthal component of the magnetic field, and only look at the transverse components in a 2-dimensional cross-section of the storage ring. The active storage volume is a source free region, and therefore the scalar magnetic potential Φ inside must fulfill Laplace's equation, i.e. [100]

$$\nabla^2 \Phi = \frac{\partial^2 \Phi}{\partial x^2} + \frac{\partial^2 \Phi}{\partial y^2} = \frac{1}{r}\frac{\partial}{\partial r}\left(r \frac{\partial \Phi}{\partial r}\right) + \frac{1}{r^2}\frac{\partial^2 \Phi}{\partial \theta^2} = 0, \tag{4.2}$$

where x and y are the usual 2D cartesian coordinates, and $r = \sqrt{x^2 + y^2}$ and $\tan\theta = y/x$ are 2D polar coordinates. In the Muon g-2 storage ring cross-section, $x = R - R_0$ is the radial offset from the nominal muon orbit R_0, and y is the vertical position[1].

The solution to the source-free Laplace equation in polar coordinates[2] for a region

[1] A visualization of the different coordinate systems used in the Muon g-2 experiment is provided in appendix A.

[2] A detailed derivation can be found in appendix B.1.

including the origin is [100] (with free parameters a_0, a_n, and b_n)

$$\Phi(r,\theta) = a_0 + \sum_{n=1}^{\infty} r^n \Big(a_n \sin(n\theta) + b_n \cos(n\theta) \Big), \tag{4.3}$$

and the magnetic field components can be calculated as the gradients of the scalar potential, $\vec{B} = -\nabla\Phi$. This solution is equivalent to a 3D solution in cylindrical coordinates with no z-dependence. It is helpful to write these in terms of the normalized radius r/r_0, where $r_0 = 45$ mm is the outer edge of the storage volume, and let the normalization terms be absorbed by the free parameters. Thus the radial (x) and vertical (y) magnetic field components in terms of the cross-sectional radius r and polar angle θ are

$$B_y(r,\theta) = \tilde{a}_0 + \sum_{n=1}^{\infty} \left(\frac{r}{r_0}\right)^n \Big(\tilde{a}_n \cos(n\theta) + \tilde{b}_n \sin(n\theta) \Big), \tag{4.4}$$

$$B_x(r,\theta) = -\tilde{b}_0 + \sum_{n=1}^{\infty} \left(\frac{r}{r_0}\right)^n \Big(\tilde{a}_n \sin(n\theta) - \tilde{b}_n \cos(n\theta) \Big). \tag{4.5}$$

The $\tilde{a}_n$ and $\tilde{b}_n$ parameters are called the "multipole strengths", with $\tilde{a}$ denoting the "normal" multipoles and $\tilde{b}$ the "skew" multipoles. $\tilde{a}_0$ is the "normal dipole", i.e. the constant vertical field and thus by far the biggest component. The "skew dipole" $\tilde{b}_0$ would correspond to a constant radial field and is assumed to be zero in the analysis, since the NMR probes as scalar magnetometers are unable to measure the radial component. By renaming the $\tilde{a}_n$ and $\tilde{b}_n$ parameters to independently numbered m_i and assigning corresponding base functions $f_i(r,\theta) = f_i(x,y)$ (see tab. 4.1), B in a 2D slice can be written as a sum of single terms as

$$B(r,\theta) \approx B_y(r,\theta) = \sum_{i=1}^{\infty} m_i f_i(r,\theta), \tag{4.6}$$

where the base functions are defined as[3] [8]

$$f_i(r,\theta) = \begin{cases} 1 & \text{for } i = 1, \\ \left(\frac{r}{r_0}\right)^{\frac{i}{2}} \cos(\frac{i}{2}\theta) & \text{for } i > 1,\ i \text{ even}, \\ \left(\frac{r}{r_0}\right)^{\frac{i-1}{2}} \sin(\frac{i-1}{2}\theta) & \text{for } i > 1,\ i \text{ odd}. \end{cases} \tag{4.7}$$

[3]For historical reasons, m_4 and m_5 as well as m_6 and m_7 are switched, i.e. m_4 and m_6 correspond to the skew sextupole and octupole, and m_5 and m_7 to the normal sextupole and octupole. Additionally, if B_x is calculated, m_0 is used to indicate the skew dipole strength.

Within a cross-sectional slice, only 17 data points at the 17 trolley probe locations are available, thus the infinite series has to be truncated. In Run-4/5/6, terms up to $n = 6$ are included [51], which corresponds to 13 multipoles. This choice was determined as optimal in a study of Run-2/3 data residuals [133] and results in a set of over-constrained equations to allow for a proper fit, while also minimizing the truncation error. It should be noted that due to the geometric arrangement of the trolley probes, the skew 14-pole can not actually be measured. Each of the trolley probes sits at a node of the skew 14-pole, and the fit will always return $\tilde{b}_6 = 0$ [140]. Table 4.1 shows all included multipoles in both polar and cartesian coordinates, and fig. 4.1 shows each non-zero multipole component in B_y individually with a normalized strength $m_i = 1$.

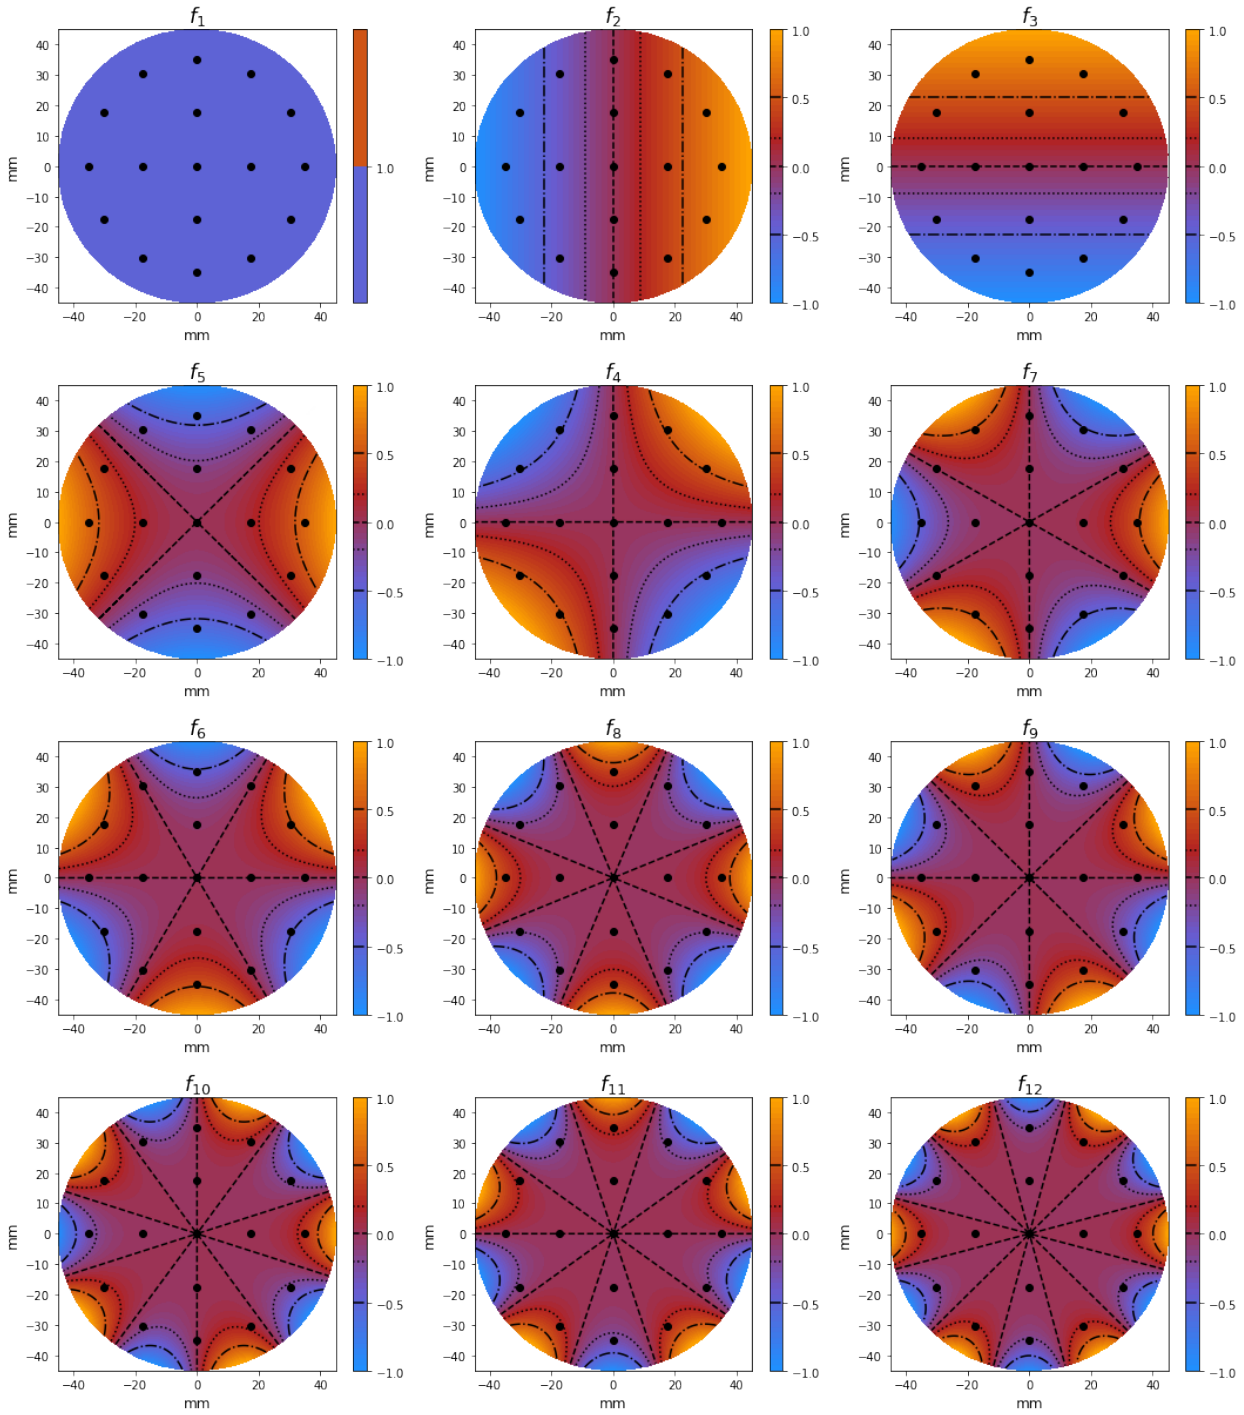


Figure 4.1: 2D multipole base functions. Individual 2D multipole base functions within the muon storage volume. The gradient shows the strength of B_y, with the dots indicating the position of the trolley probes.

Multipole moment	**n**	$\mathbf{m_i}$	$\mathbf{f_i(r,\theta)}$	$\mathbf{f_i(x,y)}$
Dipole (Dip)	0	$m_1 = \tilde{a}_0$	1.0	1.0
Normal quadrupole (NQ)	1	$m_2 = \tilde{a}_1$	$\frac{r}{r_0}\cos(\theta)$	x/r_0
Skew quadrupole (SQ)	1	$m_3 = \tilde{b}_1$	$\frac{r}{r_0}\sin(\theta)$	y/r_0
Skew sextupole (S6)	2	$m_4 = \tilde{b}_2$	$\left(\frac{r}{r_0}\right)^2\sin(2\theta)$	$(2xy)/r_0^2$
Normal sextupole (N6)	2	$m_5 = \tilde{a}_2$	$\left(\frac{r}{r_0}\right)^2\cos(2\theta)$	$(x^2-y^2)/r_0^2$
Skew octupole (S8)	3	$m_6 = \tilde{b}_3$	$\left(\frac{r}{r_0}\right)^3\sin(3\theta)$	$(3x^2y-y^3)/r_0^3$
Normal octupole (N8)	3	$m_7 = \tilde{a}_3$	$\left(\frac{r}{r_0}\right)^3\cos(3\theta)$	$(x^3-3xy^2)/r_0^3$
Normal decapole (N10)	4	$m_8 = \tilde{a}_4$	$\left(\frac{r}{r_0}\right)^4\cos(4\theta)$	$(x^4-6x^2y^2+y^4)/r_0^4$
Skew decapole (S10)	4	$m_9 = \tilde{b}_4$	$\left(\frac{r}{r_0}\right)^4\sin(4\theta)$	$(x^3y-xy^3)/r_0^4$
Normal 12-pole (N12)	5	$m_{10} = \tilde{a}_5$	$\left(\frac{r}{r_0}\right)^5\cos(5\theta)$	$\frac{x^5-10x^3y^2+5xy^4}{r_0^5}$
Skew 12-pole (S12)	5	$m_{11} = \tilde{b}_5$	$\left(\frac{r}{r_0}\right)^5\sin(5\theta)$	$\frac{5x^4y-10x^2y^3+y^5}{r_0^5}$
Normal 14-pole (N14)	6	$m_{12} = \tilde{a}_6$	$\left(\frac{r}{r_0}\right)^6\cos(6\theta)$	$\frac{x^6-15x^4y^2+15x^2y^4-y^6}{r_0^6}$
Skew 14-pole (S14)	6	$m_{13} = \tilde{b}_6$	$\left(\frac{r}{r_0}\right)^6\sin(6\theta)$	$\frac{6x^5y-20x^3y^3+6xy^5}{r_0^6}$
Total	≤ 6	$\lvert\vec{B}\rvert \approx B_y \simeq \sum_{i=1}^{13} m_i f_i$		

Table 4.1: List of 2D multipole moments included in the field analysis, based on eq. 4.4. The parameters $\tilde{a}_n$ and $\tilde{b}_n$ can alternatively be referred to via m_i with $i = 1, .., 13$ with their respective functions f_i such that the truncated series is a sum of 13 multipoles (instead of $\tilde{a}_0$ plus a sum of 6 terms).

4.2 Linear regression fit

Eq. 4.6 shows that in a given cross-sectional slice of the ring, the magnetic field strength B can be approximated as a linear combination of the base functions $f_i(r, \theta)$. Therefore, the set of magnetic field measurements of the 17 trolley probes can be expressed in matrix form as an overconstrained set of 17 linear equations,

$$\mathbf{b} = \mathbf{F} \cdot \mathbf{m}, \tag{4.8}$$

where $\mathbf{b} = \big(B(x_1, y_1), ..., B(x_{17}, y_{17})\big)^T$ is the vector of magnetic field values at the 17 probe positions, $\mathbf{m} = (m_1, ..., m_{13})^T$ is the vector of (yet unknown) multipole strengths, and the 17×13 matrix $\mathbf{F}$ is constructed from the base functions as

$$F_{ij} = f_j(x_i, y_i) = f_j(r_i, \theta_i). \tag{4.9}$$

Finding the vector $\mathbf{m}$ is then a linear regression problem which can be solved using the least squares method, i.e. finding $\mathbf{m}$ that minimizes the sum of squared residuals

$$S = \sum_{i=1}^{17} \big(\Delta B(x_i, y_i)\big)^2 = \sum_{i=1}^{17} \Big(B(x_i, y_i) - \sum_{j=1}^{13} m_j f_j(x_i, y_i)\Big)^2 = \|\mathbf{b} - \mathbf{F} \cdot \mathbf{m}\|^2, \tag{4.10}$$

where $\|\mathbf{x}\|$ denotes the Euclidean 2-norm, or the spectral norm for a matrix. The solution that minimizes $\|\mathbf{b}-\mathbf{F}\cdot\mathbf{m}\|$ (and thus also $\|\mathbf{b}-\mathbf{F}\cdot\mathbf{m}\|^2$) can be found using the pseudo-inverse of $\mathbf{F}$, $\mathbf{F}^+$ [130], such that

$$\mathbf{m} = \mathbf{F}^+ \cdot \mathbf{b}. \tag{4.11}$$

In contrast to the inverse of a matrix, which is only defined for square non-singular matrices, the pseudo-inverse can be calculated for any matrix through singular value decomposition (SVD) [86]. Any $m \times n$ matrix $\mathbf{A}$, $m \geq n$, can be written as

$$\mathbf{A} = \mathbf{U} \cdot \mathbf{D} \cdot \mathbf{V}^T, \tag{4.12}$$

(its SVD,) where $\mathbf{U}$ is a $m \times n$ column-orthogonal matrix ($\mathbf{U}^T \cdot \mathbf{U} = \mathbf{I}_n$), $\mathbf{D} = \mathrm{diag}(d_1, ..., d_n)$ is a $n \times n$ diagonal matrix with elements $d_1 \geq ... \geq d_n \geq 0$ (singular values[4]), and $\mathbf{V}$ is a $n \times n$ orthogonal matrix ($\mathbf{V}^T \cdot \mathbf{V} = \mathbf{V} \cdot \mathbf{V}^T = \mathbf{I}_n$). The pseudo-inverse can then be defined as

$$\mathbf{A}^+ = \mathbf{V} \cdot \mathrm{diag}(d_1^+, ..., d_n^+) \cdot \mathbf{U}^T, \tag{4.13}$$

[4]The number on non-negative singular values of a matrix is equal to the rank of that matrix [86].

with

$$d_i^+ = \begin{cases} \frac{1}{d_i} & \text{if } d_i > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{4.14}$$

Standard SVD algorithms are backward stable, which means the computed SVD is the exact decomposition of the slightly perturbed matrix $\mathbf{A} + \mathbf{\Delta A}$ with

$$\|\mathbf{\Delta A}\| \lesssim \mathcal{O}(\epsilon_m)\|\mathbf{A}\| = \mathcal{O}(\epsilon_m) d_1, \tag{4.15}$$

and machine precision $\epsilon_m \approx 10^{-16}$ for double precision floating point numbers [19]. As mentioned above, m_{13} will always return zero and is therefore in practice excluded from the fit. The remaining 17×12 matrix $\mathbf{F}$ has full column rank and the difference between computed and true pseudo-inverse is bounded by [135]

$$\frac{\|(\mathbf{F} + \mathbf{\Delta F})^+ - \mathbf{F}^+\|}{\|\mathbf{F}^+\|} \leq \sqrt{2}\ \kappa(\mathbf{F}) \frac{\|\mathbf{\Delta F}\|}{\|\mathbf{F}\|} \lesssim \mathcal{O}(\epsilon_m)\ \kappa(\mathbf{F}), \tag{4.16}$$

where $\kappa(\mathbf{F}) = \|\mathbf{F}\| \cdot \|\mathbf{F}^+\| = d_1/d_n$ is the condition number of $\mathbf{F}$, in this case $\kappa(\mathbf{F}) \approx 5.9$. This means the numerical error in the calculated pseudo-inverse due to finite machine precision will still be $\mathcal{O}(\epsilon_m)$, and since the spectral norm is defined as the maximum amplification of any unit vector, $\|A\| = \max_{\|x\|=1}(\|Ax\|)$, this also limits the error in any of the fitted multipoles m_i. Explicitly taking the numerical roundoff error in $\mathbf{b}$ into account as well (which is also $\mathcal{O}(\epsilon_m)$ and distinct from other sources of error in the magnetic field measurements), the algorithmic error on the solution is [86]

$$\frac{\|\mathbf{\Delta m}\|}{\|\mathbf{m}\|} \leq \kappa(\mathbf{F}) \left(\frac{\|\mathbf{\Delta F}\|}{\|\mathbf{F}\|} + \frac{\|\mathbf{\Delta b}\|}{\|\mathbf{b}\|} \right) + \mathcal{O}(\epsilon_m^2) \lesssim \mathcal{O}(\epsilon_m)\kappa(\mathbf{F}). \tag{4.17}$$

The ultimate goal of this linear regression fit is to determine the multipole strengths m_i that are best able to reconstruct the measured field $\mathbf{b}$, i.e. minimize the residual $\mathbf{r} = \mathbf{b}^{\text{fit}} - \mathbf{b}$. Any residual between the unperturbed fit result $\mathbf{b}^{\text{fit}}$ and the input $\mathbf{b}$ is due to effects that the model cannot capture, like the contribution of radial/longitudinal fields. It is important to ensure that any error introduced by the fit itself $\mathbf{\Delta b}^{\text{fit}} = \mathbf{\Delta r}$ is significantly smaller than the desired sensitivity of the fit. In the Muon g-2 experiments, magnetic field differences smaller than a part per billion can generally be neglected due to the combined size of systematic errors (see [10, 7]), thus the fit should ideally not introduce any changes larger than that,

i.e. $|\Delta b_i^{\text{fit}}| \leq \|\mathbf{\Delta b}^{\text{fit}}\| \leq 10^{-9}|b_i| \leq 10^{-9}\|\mathbf{b}\|$. From eqs. 4.15, 4.16, and 4.11 follows

$$\begin{aligned}\frac{\|\Delta\mathbf{b}^{\text{fit}}\|}{\|\mathbf{b}\|} &= \frac{\|(\mathbf{F}+\mathbf{\Delta F})\cdot(\mathbf{m}+\mathbf{\Delta m})-\mathbf{F}\cdot\mathbf{m}\|}{\|\mathbf{b}\|} \approx \frac{\|\mathbf{\Delta F}\cdot\mathbf{m}+\mathbf{F}\cdot\mathbf{\Delta m}\|}{\|\mathbf{b}\|} + \mathcal{O}(\epsilon_m^2) \\ &\lesssim \frac{\|\mathbf{\Delta F}\|\cdot\|\mathbf{m}\|+\|\mathbf{F}\|\cdot\|\mathbf{\Delta m}\|}{\|\mathbf{b}\|} \lesssim \mathcal{O}(\epsilon_m)\ \kappa(\mathbf{F})\big(1+\kappa(\mathbf{F})\big) \ll 10^{-9}, \end{aligned} \tag{4.18}$$

which confirms that the fit algorithm does not introduce errors to any significant level.

4.3 3D toroidal parameterization

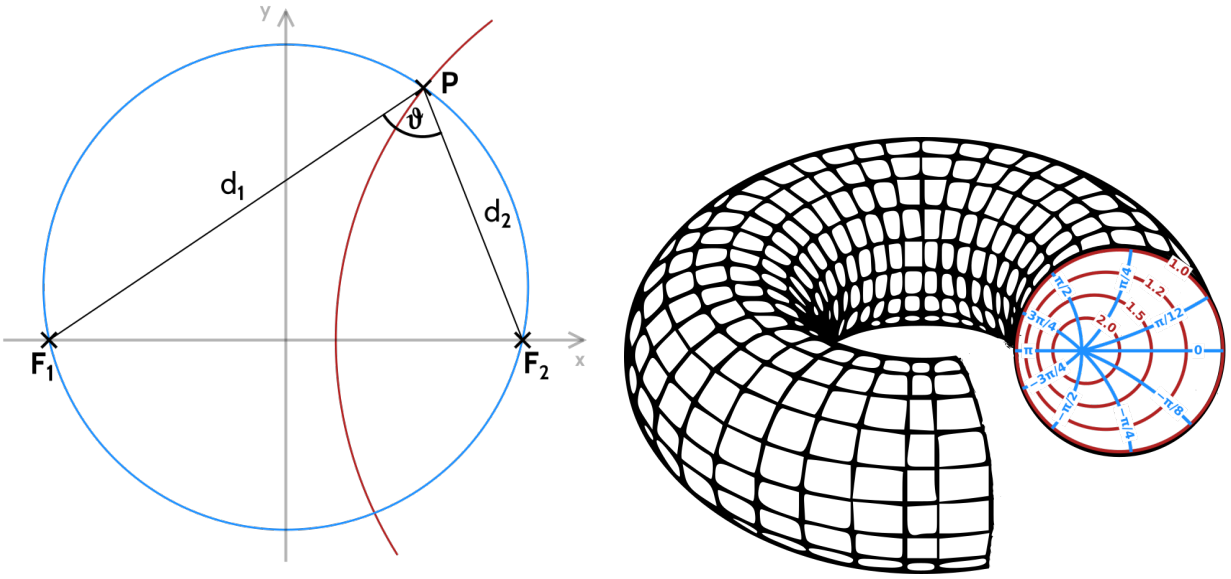
The 2-dimensional parameterization introduced above relies on the assumption that the longitudinal field (i.e. the field component perpendicular to each transverse 2D slice) is zero. In reality, this cannot be true, proven both by direct measurements of the azimuthal field component after passive shimming was completed (see section 3.1.2), as well as the existence of azimuthal variations of the field measured by the trolley[5].

Expanding the 2D parameterization to 3D is non-trivial, as it requires the definition of the storage ring's source-free region in the chosen 3-dimensional coordinate system to solve Laplace's equation analytically. While a cylindrical coordinate system could be a suitable choice in each slice if the cylinder length is taken as infinitesimally small, this cannot be generalized to describe the entire storage volume, due to the curvature of the storage ring. Neither spherical nor cartesian coordinates can be used, since no spherical or rectangular boundary can be defined that contains the storage volume but no magnetic field sources. There is, however, a lesser-known set of curvilinear coordinates ideally suited to this problem, called toroidal coordinates.

4.3.1 Introduction to toroidal coordinates

Toroidal coordinates are based on a 2D coordinate system called bipolar coordinates (with η, ϑ in the 2D plane), which are then rotated around the vertical (y) axis [105]. This gives the azimuthal angle ψ as the third coordinate. Bipolar coordinates relate a point P to two focal points F_1 and F_2 located at $\pm a$ on the x-axis (in 2D). The point is then defined by the angle it forms with the focal points, and the relative difference of the distances d_1 and d_2 to

[5] According to Maxwell's equations [91], in a source- and current-free region with no (or constant) electric field, $\nabla \times \vec{B} = 0$, and therefore in any given 2D slice of the storage ring with x pointing radially outward, y being the vertical direction, and z normal to the slice in the direction of the muons' trajectory, $\frac{\partial B_y}{\partial z} = \frac{\partial B_z}{\partial y}$. An azimuthal variation in the field measured by the trolley (which is almost entirely B_y) thus indicates a vertical variation of B_z, which means it can't be zero everywhere in the storage volume.



(a) Geometric definition of bipolar coordinates, cf. eq. 4.19. The blue circle is a line of constant ϑ, and the red circle segment is at constant η.

(b) Intersecting circles of constant η and constant ϑ. Torus cross-section with intersecting circles of constant η (red) and constant ϑ (blue), illustrating orthogonality at each intersection.

Figure 4.2: Illustration of bipolar coordinates.

the foci (see fig. 4.2a) [105],

$$\vartheta = \angle F_1 P F_2, \qquad \eta = \ln \frac{d_1}{d_2}. \tag{4.19}$$

Lines of constant ϑ form circles of different radii, centered on the y-axis, that always intersect the two foci (and therefore each other). Lines of constant η form non-concentric, non-intersecting circles of different radii around one of the focal points, with $\eta = 0$ ($d_1 = d_2$) being a straight line along the y-axis and $\eta \to \pm\infty$ approaching the foci. These two sets of circles are depicted in fig. 4.2b). Any circle in the first set (const. ϑ) will intersect any circle in the second set (const. η) at a right angle, and thus any given point with coordinates (η, ϑ) will lie on the intersection of a circle of constant η and a circle of constant ϑ, and the unit vectors (which are tangents to the circles) will be orthogonal.

In terms of the cylindrical coordinate system of the Muon g-2 ring with y being the vertical direction, R the radial distance from the center of the cylinder, and azimuthal

coordinate ϕ, the following identities define the toroidal coordinates [105]:

$$\eta = \operatorname{arctanh}\left(\frac{2aR}{R^2+y^2+a^2}\right) \qquad \in [0,\infty) \tag{4.20}$$

$$\vartheta = \arctan\left(\frac{2ay}{R^2+y^2-a^2}\right) \qquad \in (-\pi,\pi] \tag{4.21}$$

$$\psi = \phi \qquad \in [0,2\pi) \tag{4.22}$$

In 2D, for a given circle at $(R_0, 0)$ with radius r_0, there is a unique combination of $a = a_0$ and $\eta = \eta_0$ that will produce the same circle[6]. This allows for an easy definition of the boundary of the storage volume compared to other curvilinear coordinate systems. Points within the storage volume are then be restricted to $\eta \geq \eta_0$.

4.3.2 Solution to Laplace's equation

The Laplacian in toroidal coordinates is[7] [122]

$$\begin{aligned}\nabla^2\Phi = \frac{(\cosh\eta-\cos\vartheta)^3}{a^2\sinh\eta}\Bigg[&\frac{\partial}{\partial\eta}\left(\frac{\sinh\eta}{\cosh\eta-\cos\vartheta}\frac{\partial\Phi}{\partial\eta}\right)\\ &+\sinh\eta\,\frac{\partial}{\partial\vartheta}\left(\frac{1}{\cosh\eta-\cos\vartheta}\frac{\partial\Phi}{\partial\vartheta}\right)+\frac{1}{\sinh\eta\,(\cosh\eta-\cos\vartheta)}\frac{\partial^2\Phi}{\partial\psi^2}\Bigg]=0.\end{aligned} \tag{4.23}$$

The solution can be expressed via separation of variables as

$$\Phi(\eta,\vartheta,\psi) = h(\eta,\vartheta)\sum_{n,m=0}^{\infty} U_{nm}(\eta)V_m(\vartheta)W_n(\psi), \tag{4.24}$$

with $h(\eta,\vartheta) = \sqrt{\cosh\eta-\cos\vartheta}$. Both angular solutions turn out to be simple sines and cosines [122],

$$V_m(\vartheta) = a_m\sin(m\vartheta)+b_m\cos(m\vartheta), \qquad m\in\mathbb{N}_0, \tag{4.25}$$

$$W_n(\psi) = c_n\sin(n\psi)+d_n\cos(n\psi), \qquad n\in\mathbb{N}_0, \tag{4.26}$$

and the third coordinate component is

$$U_{nm}(\eta) = A_{nm}P^n_{m-1/2}(\cosh\eta)+B_{nm}Q^n_{m-1/2}(\cosh\eta), \tag{4.27}$$

[6] How to calculate a_0 and η_0 corresponding to a given torus is shown in appendix B.2.1.
[7] See derivation in appendix B.2.2.

with the associated Legendre functions of half integer degree (toroidal harmonics) $P^n_{m-1/2}$ (first kind) and $Q^n_{m-1/2}$ (second kind). The $P^n_{m-1/2}(\cosh\eta)$ have singularities at the focal points, representing sources inside the region of interest [42], and therefore the solution is restricted to terms including only the toroidal harmonics of the second kind, $Q^n_{m-1/2}$[8]. Thus the toroidal potential becomes

$$\Phi(\eta,\vartheta,\psi) = h(\eta,\vartheta)\sum_{n=0}^{\infty}\sum_{m=0}^{\infty} Q^n_{m-1/2}(\cosh\eta)\Big(a_{nm}\sin(m\vartheta)\sin(n\psi) + b_{nm}\sin(m\vartheta)\cos(n\psi) + c_{nm}\cos(m\vartheta)\sin(n\psi) + d_{nm}\cos(m\vartheta)\cos(n\psi)\Big), \tag{4.28}$$

where the individual coefficients of $U(\eta)$, $V(\vartheta)$, and $W(\psi)$ have been combined into four independent coefficients. In practice, of course, only a finite number of coefficients can be determined from a fit, which allows the use of independently indexed coefficients M_k according to

$$M_k = M_{(i-1)j_{\max}+j} = \begin{cases} a_{\lfloor\frac{i}{2}\rfloor\lfloor\frac{j}{2}\rfloor} = a_{nm} = B_{nm}\cdot a_m\cdot c_n, & \text{for } i,j \text{ even},\\ b_{\lfloor\frac{i}{2}\rfloor\lfloor\frac{j}{2}\rfloor} = b_{nm} = B_{nm}\cdot a_m\cdot d_n, & \text{for } i \text{ odd}, j \text{ even},\\ c_{\lfloor\frac{i}{2}\rfloor\lfloor\frac{j}{2}\rfloor} = c_{nm} = B_{nm}\cdot b_m\cdot c_n, & \text{for } i \text{ even}, j \text{ odd},\\ d_{\lfloor\frac{i}{2}\rfloor\lfloor\frac{j}{2}\rfloor} = d_{nm} = B_{nm}\cdot b_m\cdot d_n, & \text{for } i,j \text{ odd}. \end{cases} \tag{4.29}$$

Appropriate base functions can be defined similarly to eq. 4.7[9] as

$$f^V_k(\vartheta) = f_j(\vartheta) = \left\{\begin{array}{ll}\sin\left(\frac{j}{2}\vartheta\right) & \text{for } j \text{ even},\\ \cos\left(\frac{j-1}{2}\vartheta\right) & \text{for } j \text{ odd},\end{array}\right\}, \quad j = \left\lfloor\frac{k}{i_{\max}}\right\rfloor + 1, \tag{4.30}$$

$$f^W_k(\psi) = f_i(\psi) = \left\{\begin{array}{ll}\sin\left(\frac{i}{2}\psi\right) & \text{for } i \text{ even},\\ \cos\left(\frac{i-1}{2}\psi\right) & \text{for } i \text{ odd},\end{array}\right\}, \quad i = k \bmod i_{\max}, \tag{4.31}$$

$$g_k(\eta) = Q^{\lfloor i/2\rfloor}_{\lfloor j/2\rfloor-1/2}(\cosh\eta), \tag{4.32}$$

[8] Appendix B.2.4 shows an in-depth analysis of the behavior of these terms in the limit $\eta\to\infty$ and justifies the exclusion of the toroidal harmonics of the first kind.

[9] Since for $n=0$ or $m=0$ only the cosine terms in eqs. 4.25 and 4.26 are non-zero, $i,j=0$ can be excluded from potential solutions, and indexing is restricted to $i,j\geq 1$.

such that the potential can be written as a simple sum

$$\Phi(\eta,\vartheta,\psi) = h(\eta,\vartheta)\sum_{k=1}^{k_{\max}} M_k\ g_k(\eta)\ f_k^V(\vartheta)\ f_k^W(\psi). \tag{4.33}$$

The vertical (y), radial ($x = R - R_0$), and longitudinal ($z = \phi = \psi$) field components can be expressed in terms of derivatives of the toroidal potential as[10]

$$\begin{aligned} B_y &= \frac{\sinh\eta\sin\vartheta}{a}\frac{\partial\Phi}{\partial\eta} + \frac{1-\cosh\eta\cos\vartheta}{a}\frac{\partial\Phi}{\partial\vartheta} \qquad (4.34)\\ &= \frac{h}{a}\sum_k M_k f_k^W \left(\sinh\eta\sin\vartheta\ \frac{\partial g_k}{\partial\eta}\ f_k^V + \frac{1}{2}\cosh\eta\sin\vartheta\ g_k\ f_k^V \right.\\ &\qquad \left. +(1-\cosh\eta\cos\vartheta)\ g_k\ \frac{\partial f_k^V}{\partial\vartheta}\right),\\ B_x &= \frac{\sinh\eta\sin\vartheta}{a}\frac{\partial\Phi}{\partial\vartheta} - \frac{1-\cosh\eta\cos\vartheta}{a}\frac{\partial\Phi}{\partial\eta} \qquad (4.35)\\ &= \frac{h}{a}\sum_k M_k f_k^W \left(\sinh\eta\sin\vartheta\ g_k\ \frac{\partial f_k^V}{\partial\vartheta} + \frac{1}{2}\sinh\eta\cos\vartheta\ g_k\ f_k^V \right.\\ &\qquad \left. -(1-\cosh\eta\cos\vartheta)\ \frac{\partial g_k}{\partial\eta}\ f_k^V\right),\\ B_z &= -\frac{\cosh\eta-\cos\vartheta}{a\sinh\eta}\ \frac{\partial\Phi}{\partial\psi} = -\frac{h^3}{a\sinh\eta}\sum_k M_k\ g_k\ f_k^V\ \frac{\partial f_k^W}{\partial\psi}. \qquad (4.36) \end{aligned}$$

4.3.3 Direct comparison to 2D polar solution

Eqs. 4.20-4.22 show that even in toroidal coordinates, the azimuthal coordinate ψ is independent of the transverse coordinates η and ϑ, and the azimuthal position only affects the amplitude of certain transverse components of B_y (coupled through n), but not their shape (cf. eq. 4.34). This allows for a direct comparison between the 2D multipole base functions as shown in fig. 4.1, and the transverse behavior of B_y in toroidal coordinates for $f_k^W = 1$. Fig. 4.3 shows the first 12 components ($1 \le k \le 12$), where $n = 0$ and thus $f_k^W = 1$ at any ψ. An azimuthally constant B_y would only contain $n = 0$ terms and one might therefore expect the toroidal terms with $n = 0$ to be equal to the 2D multipoles. There are indeed strong similarities, for example the $k = 6$ term resembles the normal sextupole, the $k = 7$ term the skew sextupole, the $k = 8$ term the normal octupole, and so on. However, a direct comparison (through subtracting or adding a normalized term from/to the corresponding 2D multipole) shows subtle differences, as depicted for select terms in fig. 4.4. In particular,

[10] A detailed derivation of the magnetic field components can be found in appendix B.2.3.

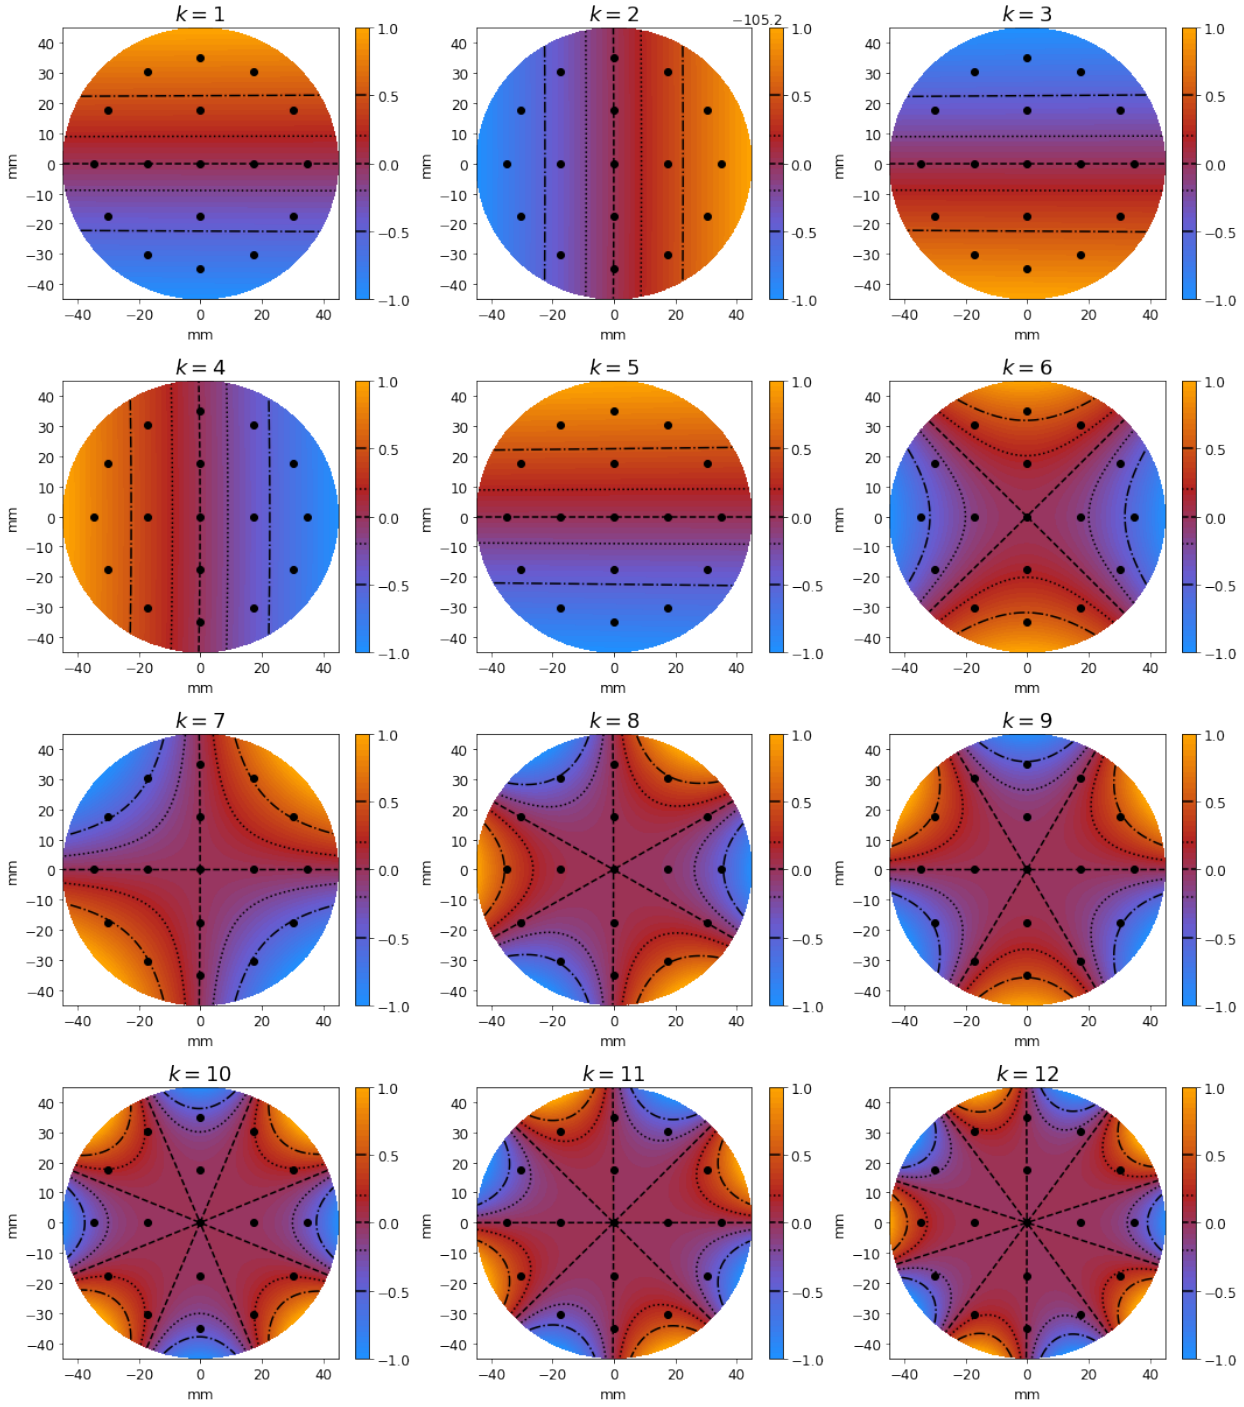
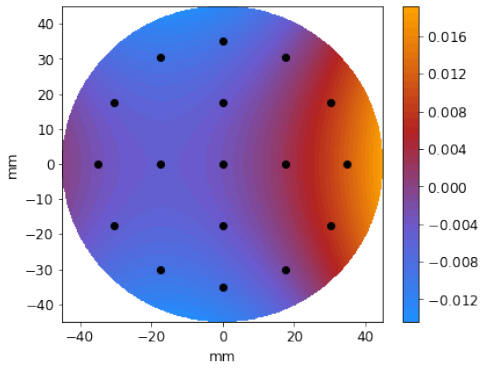
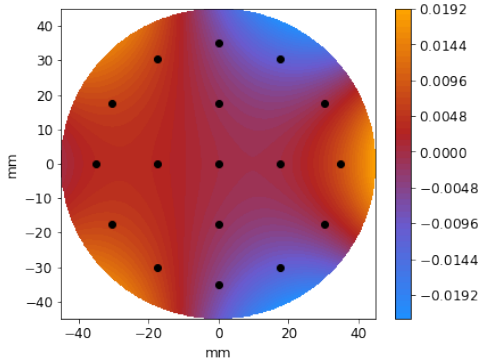


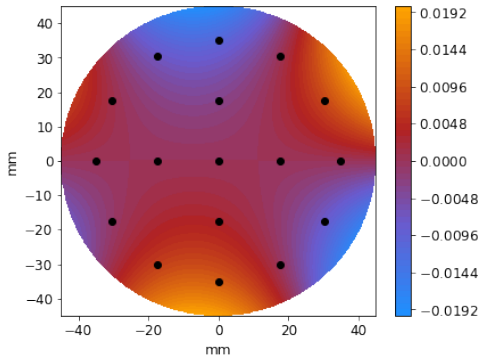
Figure 4.3: Toroidal B_y shape for $n = 0$. Toroidal B_y components for $1 \leq k \leq 12$. The gradient shows the strength of B_y normalized to $|B_y| \leq 1$, with the dots indicating the position of the trolley probes.



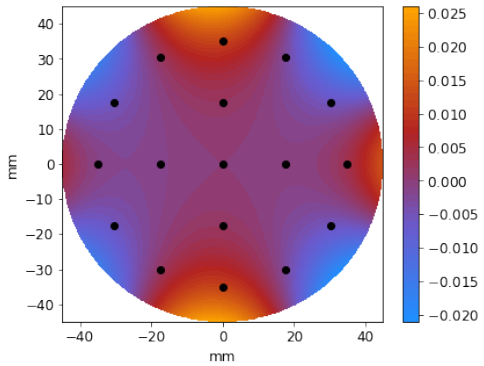
(a) 2D base function f_2 (NQ) and normalized toroidal B_y for $k = 4$, added



(b) 2D base function f_5 (N6) and normalized toroidal B_y for $k = 6$, added

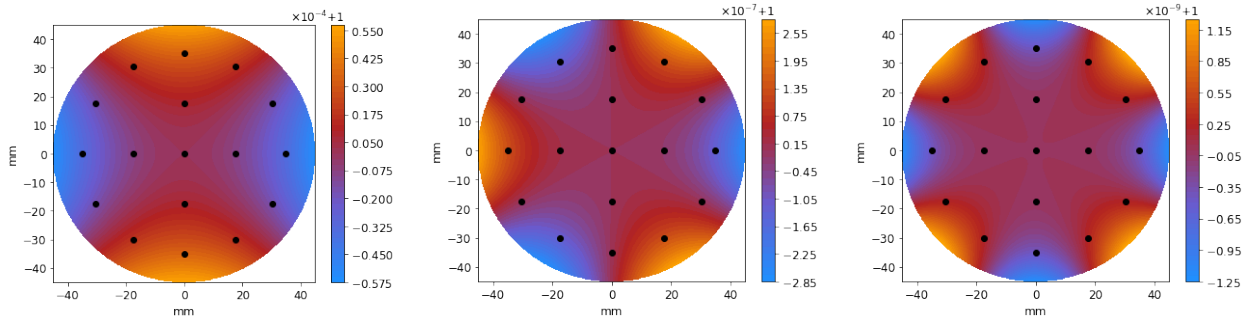


(c) 2D base function f_4 (S6) and normalized toroidal B_y for $k = 7$, subtracted



(d) 2D base function f_7 (N8) and normalized toroidal B_y for $k = 8$, added

Figure 4.4: Direct comparison of 2D and 3D transverse B_y components.



(a) Fit to constant B_y with two terms ($k = 2, 4$)

(b) Fit to constant B_y with three terms ($k = 2, 4, 6$)

(c) Fit to constant B_y with four terms ($k = 2, 4, 6, 8$)

Figure 4.5: Toroidal fit attempts for const. B_y. Attempts to fit a constant B_y with few toroidal terms. The gradient indicates field strength and the dots show the trolley probe positions. The $k = 2$ term provides a constant offset, additional terms are added to eliminate transverse variation. Each added term reduces variation such that the remainder resembles a higher order, but a true constant field requires an infinite number of terms.

while a visual inspection of fig. 4.3 seems to indicate symmetry around the central orbit as in the 2D case, the differences between the 2D and 3D fields in fig. 4.4 show vertical symmetries, but not in the radial direction. This is due to fact that the focal point distance a_0 of the toroidal coordinate system is very slightly to the left of the central orbit.

It is also noteworthy that in the toroidal parameterization, there is no single term equivalent to f_1 producing just a constant, uniform B_y. The $k = 2$ term has transverse variations around a constant offset, but finding a combination of terms that eliminate these variations turns out to be non-trivial. Fig. 4.5 shows attempts to fit[11] $B_y(R, y, \varphi) = 1$ (unit is irrelevant here) with two to four terms only. Using only $k = 2$ and $k = 4$ results in transverse variations of the field that strongly resemble the shape of the $k = 6$ term. Adding this term to the fit then creates variations that look like the $k = 8$ term, and again adding this term to the fit as well shows a shape similar to the $k = 10$ term. This suggests that for a truly uniform field in the entire torus, an infinite number of terms would be required. For this reason, the toroidal fit will subtract the mean of each dataset (the equivalent of m_1 in the 2D multipole parameterization) before performing the linear regression.

A note on magnetic field tracking between trolley runs

As discussed in section 3.2.3, the fixed probes track the magnetic field between trolley runs. The fixed probe stations are located outside of the storage volume, and contain fewer probes

[11] The fit is performed using the same inverse matrix method as described in section 4.2 for the 2D multipole fits.

than the trolley (see fig. 3.10). While they are not in arranged in a symmetry specifically optimized for polar coordinates, their relative geometry can nonetheless be used to determine multiple gradients of the field in x and y, which in turn can provide a simple estimate of some of the 2D multipole strengths at the center of the storage volume [140]. The analytical equations and estimates are as follows[12]:

- Dipole strength $m_1 = A_0 = B_y(r = 0) \approx \langle B \rangle$, where the average is taken across all probes for 6-probe stations, and across the two middle probes for 4-probe stations.
- NQ strength $m_2 = A_1 = r_0 \left.\frac{\partial B_y}{\partial x}\right|_{r=0} \approx r_0 \langle \frac{\Delta_x B}{\Delta x} \rangle_y$. $\Delta_x B$ denotes the difference between the magnetic field measurements of two probes at the same y but different x. For 6-probe stations, only the inner and outer probes are used (see fig. 3.10).
- SQ strength $m_3 = C_1 = r_0 \left.\frac{\partial B_y}{\partial y}\right|_{r=0} \approx r_0 \langle \frac{\Delta_y B}{\Delta y} \rangle_x$. All probes are used in both 4- and 6-probe stations.
- S6 strength $m_4 = C_2 = \frac{1}{2} r_0^2 \left.\frac{\partial^2 B_y}{\partial x \partial y}\right|_{r=0} \approx r_0^2 \frac{\Delta_y (\Delta_x B)}{2 \Delta x \Delta y}$. For 6-probe stations, $\Delta_x B$ is again taken between the inner and outer probes only.
- N6 strength $m_5 = A_2 = \frac{1}{2} r_0^2 \left.\frac{\partial^2 B_y}{\partial x^2}\right|_{r=0} \approx r_0^2 \langle \frac{\Delta_x (\Delta_x B)}{2 (\Delta x)^2} \rangle_y$. Only available for 6-probe stations, since it involves taking the difference between inner/middle and middle/outer probes, and then the difference between those differences. Top and bottom results are averaged. For 4-probe stations, m_5 can be estimated as the average of m_5 from the two nearest neighbors (which are always 6-probe stations) [10].
- S8 strength $m_6 = C_3 = \frac{1}{8} r_0^3 \left.\frac{\partial^3 B_y}{\partial x^2 \partial y}\right|_{r=0} \approx r_0^3 \frac{\Delta_y (\Delta_x (\Delta_x B))}{8 (\Delta x)^2 \Delta y}$. Similarly calculated to m_5, but instead of averaging top and bottom, the difference is taken, and hence also only available for 6-probe stations.

Unfortunately, such direct correlations between coefficients and fixed probe measurements can't be established for the toroidal parameterization. A 3D fit to all fixed probes could theoretically be performed, but due to the limited number of transverse data points it is unlikely to be able to provide sufficiently precise results. Thus the 3D toroidal parameterization is examined only to check and validate the use of the 2D multipole parameterization, and not considered as a possible replacement.

[12] A few fixed probes in the inflector region and the trolley garage region of the ring have special geometries, and the multipoles have to be calculated slightly differently for those probes. Details can be found in [60, 50, 140].

4.4 Magnetic field fitting in toroidal coordinates

Since eq. 4.34 shows a linear dependence on the coefficients M_k, a matrix representing the system of linear equations under the assumption that $B \approx B_y$ can be constructed, similarly to the 2D fits,

$$\mathbf{b} = \mathbf{T} \cdot \mathbf{k}, \tag{4.37}$$

with the vector of magnetic field measurements $\mathbf{b} = \big(B(\eta_1, \vartheta_1, \psi_1), ..., B(\eta_N, \vartheta_N, \psi_N)\big)$ which contains all N measurements in the entire storage ring, the vector of coefficients $\mathbf{k} = (M_1, ..., M_{k_{\max}})^T$, and the matrix $\mathbf{T}$ constructed from eq. 4.34 via

$$\begin{aligned} T_{ij} = \frac{\sqrt{\cosh\eta_i - \cos\vartheta_i}}{a} f_j^W(\psi_i) \Bigg(& \sinh\eta_i \sin\vartheta_i \ \frac{\partial g_j(\eta_i)}{\partial \eta} \ f_j^V(\vartheta_i) \\ & + \frac{1}{2} \cosh\eta_i \sin\vartheta_i \ g_j(\eta_i) \ f_j^V(\vartheta_i) + (1 - \cosh\eta_i \cos\vartheta_i) \ g_j(\eta_i) \ \frac{\partial f_j^V(\vartheta_i)}{\partial \vartheta} \Bigg). \end{aligned} \tag{4.38}$$

This allows for a linear regression fit using the strategy introduced in section 4.2, which will be explored in detail in this section. Additionally, the toroidal parameterization provides the unique opportunity to create a simulated field that is guaranteed to satisfy Laplace's equation, such that fit results can be validated against truth values. To that end, several simulated toroidal field (STF) datasets will be simulated, described in section 4.4.3 and summarized in table 4.2. Each STF starts from a toroidal fit to one of five randomly chosen trolley maps during Run-4/5/6 and, unless otherwise noted, contains only the true magnetic field magnitude values without noise, in order to assess the purely fit-related effects.

4.4.1 Fit performance considerations

4.4.1.1 Data set size

In the 2D multipole fit, each transverse slice has only 17 data points and 12 fit parameters, and the pseudo-inverse $\mathbf{F}^+$ stays the same in each slice. In a 3D fit however, the number of data points and fit parameters is significantly larger, since the entire azimuthal range is included in the fit. During a trolley measurement, the trolley moves roughly 0.083 degrees per second, and the NMR probes are read out in sequence at a rate of about 34Hz, or about 2Hz per probe. This means a single trolley measurement has around $N_d \approx 17 \cdot 8,700 \approx 148,000$ data points, which corresponds to an array of size $\sim$ 1.2MB at double precision. Thus the matrix $\mathbf{T}$ describing the system of linear equations for the toroidal fit will have size $N_k \cdot 1.2$MB, where N_k is the number of parameters in the fit. Even with a modest number of 10 transverse and 100 longitudinal parameters, this is more than 1GB, and as will be

shown in the section 4.4.2, the actual number of fit parameters is $N_k \approx 3600$, which would bring the size of $\mathbf{T}$ to over 4GB. When calculating the pseudo-inverse via the SVD, the peak memory will need to be able to hold $\mathbf{T}$ ($N_d \times N_k$), its SVD $\mathbf{U}$ ($N_d \times N_k$), $\mathbf{D}$ (N_k diagonal), and $\mathbf{V}$ ($N_k \times N_k$), as well as $\mathbf{T}^+$ ($N_k \times N_d$). The total memory requirement for constructing the pseudo-inverse explicitly would thus be just over 13GB, and run the risk of exhausting the computational resources on a standard machine. To mitigate this issue, two changes are implemented:

- The 2D multipole fit can only be performed in cross-sectional "slices" of the storage ring, and therefore the sequential data of the original trolley probe measurements has to be interpolated to a common set of azimuthal positions for all probes. Since in a trolley run, each azimuthal position also has an associated time stamp, the standard procedure is to linearly interpolate each probe's data to a 1s time grid[13], and use numerical integration to find the average measured NMR frequency (and thus, field) in that 1s bin [140]. This results in about 4300 "slices" in a typical trolley run, reducing the number of data points for each probe to about half of the original 9000 measured positions. The toroidal fit will be performed on that same grid of interpolated data points, to decrease the memory requirement for $\mathbf{T}$ and any derived matrices by about a factor of 2.

- In the case of the 2D multipole fit, the pseudo-inverse $\mathbf{F}^+$ is, like $\mathbf{F}$, independent of azimuthal position. Calculating it once enables a simultaneous fit of all transverse "slices" at once, and it even remains unchanged for different trolley runs or STFs. In the case of the toroidal fit however, the entire trolley run or STF is fitted at once, and due to slight changes in the azimuthal locations between different datasets, each toroidal fit is entirely unique. Therefore, it is actually unnecessary to compute the pseudo-inverse explicitly, and $\mathbf{k}$ can be computed directly from the SVD of $\mathbf{T}$ via [86]

$$\mathbf{k} = \sum_{i=1}^{r} \frac{\mathbf{u}_i^T \mathbf{b}}{d_i} \mathbf{v}_i, \tag{4.39}$$

  where r is the rank[14] of $\mathbf{T}$, and $\mathbf{u}_i$ and $\mathbf{v}_i$ are the columns of $\mathbf{U}$ and $\mathbf{V}$ in the SVD

[13] The advantage of using a uniform grid in time rather than azimuth is that this allows for direct correlation between the trolley measurements and the fixed probe measurements at that time, which is a required step during each trolley run.

[14] As a note, if $\mathbf{T}$ does not have full rank, then the solution to $\min(\|\mathbf{T}\cdot\mathbf{k}-\mathbf{b}\|)$ is not unique, and $\mathbf{k} = \mathbf{T}^+\cdot\mathbf{b}$ will return the solution with the smallest norm $\|\mathbf{k}\|$ [86]. This means of the N_k coefficients, $N_k - r$ are in the null-space of $\mathbf{T}$ and changing them does not change the fit result for B_y, but may influence B_x or B_z. In this case, the objective is to approximate $|\vec{B}|$ as best as possible with B_y, and the value of the other two field components or whether the fit can reconstruct them accurately is irrelevant.

respectively. Efficient implementations for this strategy are readily available (e.g. `numpy.linalg.lstsq`), significantly reducing the peak memory needed to compute $\mathbf{k}$, as well as reducing the computational cost by $\mathcal{O}(N_d \cdot N_k \cdot r)$ operations.

4.4.1.2 Numerical stability and dynamic range

Algorithms to compute associate Legendre functions of half-integer degree are not generally included in standard computation packages, but they can be expressed via hypergeometric functions. The toroidal harmonics of the second kind with arguments larger than 1 (which $\cosh\eta$ exclusively is in the muon g-2 setup) can be written as [42, 3]

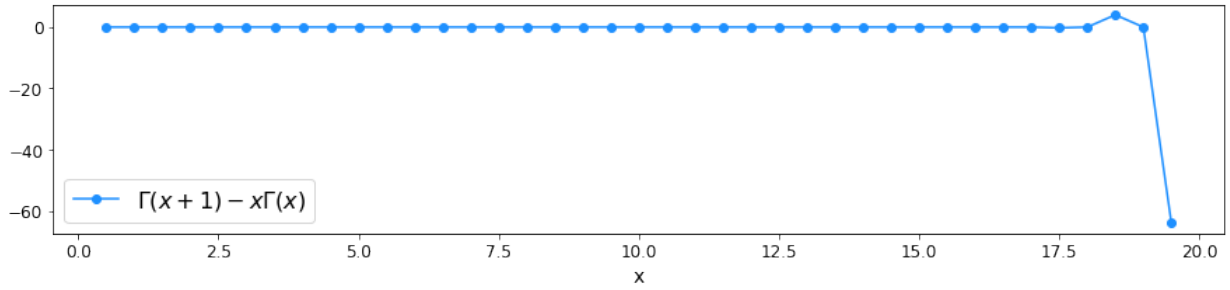
$$\begin{aligned} Q^n_{m-1/2}(\cosh\eta) = (-1)^n \sqrt{\frac{\pi}{2}}\, \frac{\Gamma(m+n+\frac{1}{2})}{2^m\,\Gamma(m+1)}\, \frac{(\tanh\eta)^n}{(\cosh\eta)^{m+\frac{1}{2}}} \\ \times\ {}_2F_1\left(\frac{1}{2}(n+m+\frac{1}{2}), \frac{1}{2}(n+m+\frac{3}{2}); m+1; \frac{1}{(\cosh\eta)^2}\right). \end{aligned} \tag{4.40}$$

It is beneficial to consider the behavior of each term in this expression across the range of possible input values to ensure numerical stability and optimize performance of the toroidal fit. In particular, SciPy's built-in gamma function (`scipy.special.gamma`) does not fulfill the recurrence relation $\Gamma(x+1) = x\Gamma(x)$ for $n \gtrsim 15$, as illustrated in fig 4.6a. An easy solution to this issue is to let the respective M_k absorb all coordinate-independent terms of eq. 4.40, such that

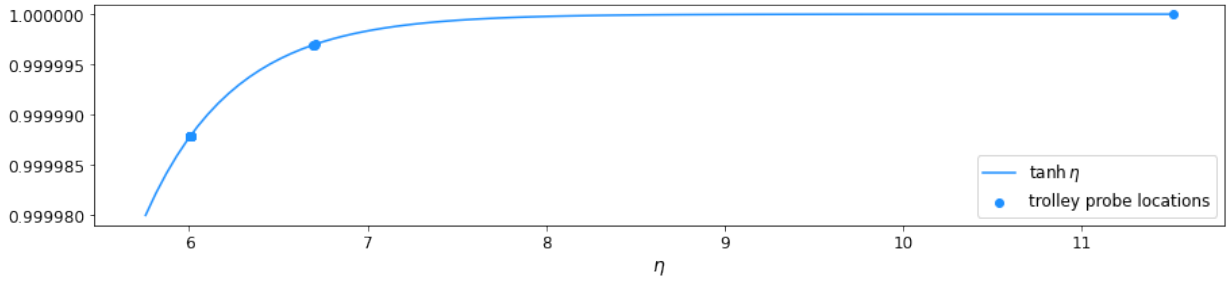
$$\tilde{M}_k = M_k\ (-1)^n \sqrt{\frac{\pi}{2}}\, \frac{\Gamma(m+n+\frac{1}{2})}{2^m\,\Gamma(m+1)}. \tag{4.41}$$

The boundary of the source-free region of the g-2 storage ring is uniquely defined by the focal point position $a_0 = \sqrt{R_0^2 - r_0^2} \approx 7111.86\text{m}$ and $\eta_{\min} = \operatorname{arccosh}\left(\frac{R_0}{r_0}\right) \approx 5.76$ (cf. appendix B.2.1), where $R_0 = 7.112\text{m}$ and $r_0 = 45\text{mm}$ as usual. The largest η value is determined by the data points closest to the focal point, which is the central trolley probe at $R = R_0$ or $x = 0$, and thus $\eta_{\max} \approx 11.51$. Across this range of η values, $\tanh\eta$ and the hypergeometric function are close to 1 at all times (see fig. 4.6b and 4.6d respectively). However, $\cosh\eta$ takes values between $\cosh\eta_{\min} \approx 158$ and $\cosh\eta_{\max} \approx 49,955$, shown in fig. 4.6c, and increasing even more quickly with larger m. While the range of m values in the toroidal fit is unlikely to be large enough to cause subnormal numbers, it is nonetheless prudent to consider normalization. Eq. 4.34 shows a factor of $h = \sqrt{\cosh\eta - \cos\vartheta} \approx \sqrt{\cosh\eta}$ (for $\eta \gg 1$) multiplying each term in the sum, effectively reducing the power of $\cosh\eta$ to $(\cosh\eta)^{-m}$ in terms containing ${Q^n_{m-1/2}}$[15]. Thus a normalization factor of $(\cosh\eta_0)^m$ can be extracted from M_k, with an appropriately chosen η_0, such that the quantity calculated by

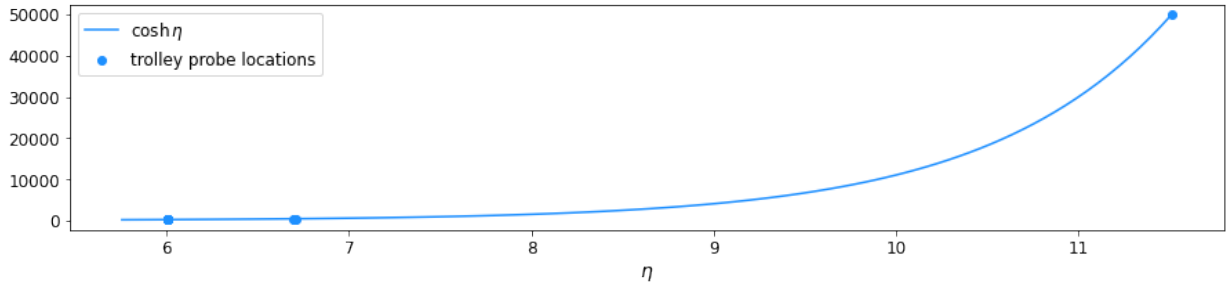
[15] Or its derivative, as will be shown shortly.



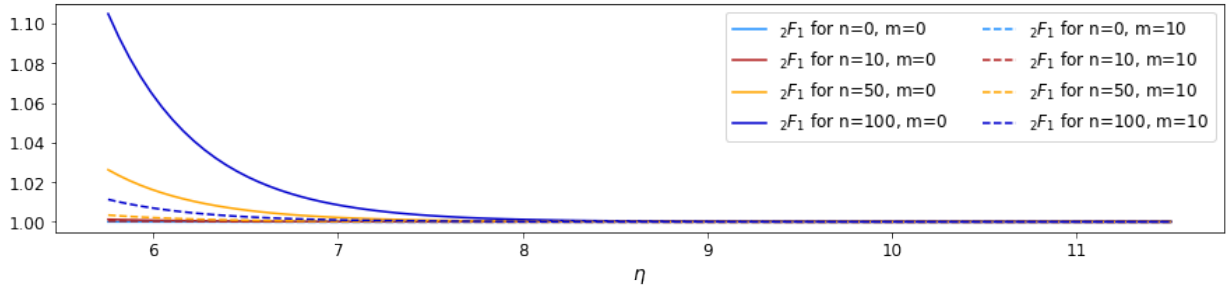
(a) Test of Γ function recurrence relation for integer and half-integer values



(b) Values of $\tanh \eta$ across the η-range of the storage ring



(c) Values of $\cosh \eta$ across the η-range of the storage ring



(d) Values of the hypergeometric function ${}_2F_1\left(\frac{1}{2}(n+m+\frac{1}{2}), \frac{1}{2}(n+m+\frac{3}{2}); m+1; \frac{1}{(\cosh\eta)^2}\right)$ across the η-range of the storage ring for a range of n and m values

Figure 4.6: Tests of functions in $Q^n_{m-1/2}$ expression, across relevant input ranges.

the fit algorithm instead of $Q^n_{m-1/2}$ becomes

$$\tilde{Q}^n_m(\eta,\eta_0) = \frac{(\tanh\eta)^n(\cosh\eta_0)^m}{(\cosh\eta)^{m+\frac{1}{2}}}\ {}_2F_1\left(\frac{1}{2}(n+m+\frac{1}{2}),\frac{1}{2}(n+m+\frac{3}{2});m+1;\frac{1}{(\cosh\eta)^2}\right)$$
$$= (-1)^n\sqrt{\frac{2}{\pi}}\frac{2^m\ \Gamma(m+1)}{\Gamma(m+n+0.5)}(\cosh\eta_0)^m\ Q^n_{m-1/2}(\cosh\eta), \tag{4.42}$$

and the parameter to be fit is

$$\tilde{M}_k = M_k\ (-1)^n\sqrt{\frac{\pi}{2}}\ \frac{\Gamma(m+n+\frac{1}{2})}{2^m\ \Gamma(m+1)}\frac{1}{(\cosh\eta_0)^m}. \tag{4.43}$$

Although this normalization factor cannot reduce the dynamic range of the matrix that

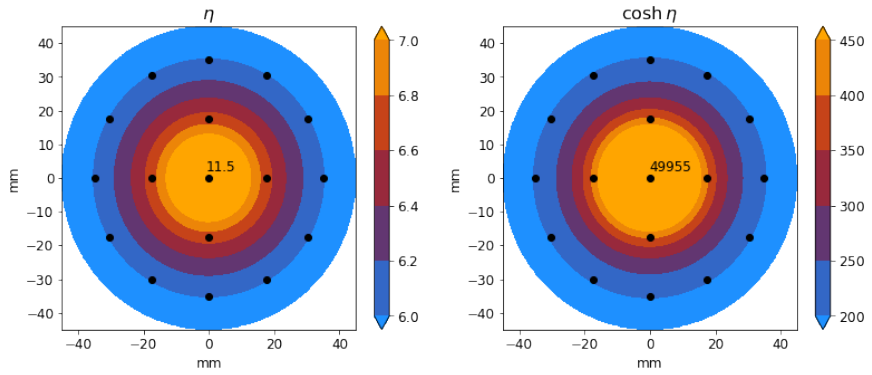


Figure 4.7: Range of η and $\cosh\eta$ values inside the storage ring. η (left) and $\cosh\eta$ (right) values inside the Muon g-2 storage ring, with dots indicating the trolley probe positions. Values for the center probe are written in text. All probes except the center probe are grouped fairly closely in η, and have $\cosh\eta$ values on the same order of magnitude, whereas the center probe's $\cosh\eta$ value is two orders of magnitude larger.

represents the system of linear equations (i.e. the orders of magnitude between the largest and smallest entry of the matrix), it can be chosen such that a majority of terms are close to $\mathcal{O}(1)$. Since the transverse position of the data points is limited to the 17 probe positions, the true spread of η and $\cosh\eta$ values that appear in the fit matrix is much smaller for all but the center probe. Fig. 4.7 shows the η and $\cosh\eta$ values in the transverse plane, with the NMR probe positions marked. The outermost probes (6-17) all have values close to $\eta \approx 6.0$, and probes 2-5 are clustered around $\eta \approx 6.8$, leading to $\cosh\eta$ values of about 200 and 400-450 respectively, two orders of magnitude smaller than the value at the center probe. This means that especially for larger m (higher order transverse components), the corresponding term in B_y at the central probe location will be vanishingly small. Similarly to the 2D multipoles, the higher m order terms in the toroidal parameterization show

variation mainly towards the outside of the storage volume (see fig. 4.3), thus it is in fact expected that the probe at the center of the storage ring can't help resolve those coefficients. A value of $\eta_0 = 6.5$ is chosen for the toroidal fit, such that $\frac{\cosh \eta_0}{\cosh \eta_{\max}} \approx 6.7 \cdot 10^{-3}$. At $m_{\max} = 6$ (see section 4.4.2), the value at the center probe is $\mathcal{O}(10^{-15})$ smaller than at the outer probe locations, indistinguishable from rounding errors. While this is not concerning in a qualitative sense, a large dynamic range in $\mathbf{T}$ increases the condition number $\kappa(\mathbf{T})$ significantly, and leads to an ill-conditioned system, which will be discussed below in 4.4.1.3.

The adjustments discussed above ensure numerical stability in the computation of $Q^n_{m-1/2}$, but the calculation of the derivative, $\frac{\partial Q^n_{m-1/2}}{\partial \eta}$, also requires careful examination. Derivatives of the associated Legendre functions can be calculated via recurrence relations in both m and n [3]. The following expressions are the recurrence relations that involve only $Q^n_{m-1/2}$ and one other Q shifted by ± 1 in either n or m:

$$\frac{\partial Q^n_{m-1/2}(\cosh \eta)}{\partial \eta} = \begin{cases} \frac{n}{\tanh \eta} Q^n_{m-1/2} + Q^{n+1}_{m-1/2}, & (4.44) \\ (m+n-0.5)(m-n+0.5) Q^{n-1}_{m-1/2} - \frac{n}{\tanh \eta} Q^n_{m-1/2}, & (4.45) \\ \frac{m-n+0.5}{\sinh \eta} Q^n_{m+1/2} - \frac{m+0.5}{\tanh \eta} Q^n_{m-1/2}, & (4.46) \\ \frac{m-0.5}{\tanh \eta} Q^n_{m-1/2} - \frac{n+m-0.5}{\sinh \eta} Q^n_{m-3/2}. & (4.47) \end{cases}$$

For a given n and m, the normalized $\tilde{Q}^n_m$ in eq. 4.42 is just $Q^n_{m-1/2}$ multiplied by a constant factor and as such the derivative is scaled in the same way. The equivalent recurrence relations for $\frac{\partial \tilde{Q}^n_m}{\partial \eta}$ are then

$$\frac{\partial \tilde{Q}^n_m(\eta, \eta_0)}{\partial \eta} = \begin{cases} \frac{n}{\tanh \eta} \tilde{Q}^n_m - (m+n+0.5) \tilde{Q}^{n+1}_m, & (4.48) \\ (n-m-0.5) \tilde{Q}^{n-1}_m - \frac{n}{\tanh \eta} \tilde{Q}^n_m, & (4.49) \\ \frac{(m-n+0.5)(m+n+0.5)}{2(m+1) \sinh \eta \cosh \eta_0} \tilde{Q}^n_{m+1} - \frac{m+0.5}{\tanh \eta} \tilde{Q}^n_m, & (4.50) \\ \frac{m-0.5}{\tanh \eta} \tilde{Q}^n_m - \frac{2m \cosh \eta_0}{\sinh \eta} \tilde{Q}^n_{m-1}. & (4.51) \end{cases}$$

In each of the equations above, there is a potential for catastrophic cancellation[16] for certain ranges of m, n, and/or η. To evaluate which recursion relations are the most numerically stable, the analytical results for each of the four cases are compared to a second order finite differences estimate of the gradient (using `numpy.gradient` on $10,000$ points between $\eta_{\min}$ and $\eta_{\max}$), the relative residuals of which are shown in fig. 4.8. It is immediately clear that the $(m+1)$-based derivative (eq. 4.50) delivers significantly worse results that the other three options, not only because the relative difference to the gradient is several orders of magnitude larger, but also because the residual is changes drastically over the η-range, being the worst at small η. Since all but the center probe have relatively small η values, this would affect the toroidal fit greatly. Additionally, the other three residuals decrease if the number of points between $\eta_{\min}$ and $\eta_{\max}$ is increased, but for the $(m+1)$-based derivative, the relative residual doesn't change. Since the finite differences gradient is expected to be closer to the true gradient with smaller step size, this suggests that the $(n+1)$-, $(n-1)$-, and $(m-1)$-based derivatives are more accurate expressions. Comparing the four derivatives to each other (fig. 4.9) also confirms that the $(n \pm 1)$- and $(m-1)$-based derivatives agree well with each other, whereas the relative differences between the $(m+1)$-based derivative and each of the others is large in particular for small η.

Since three of the four derivative expressions seem to perform equally well, the choice is made based on practical considerations: If the $(m-1)$-based equation were used, the derivative of $\tilde{Q}_0^n$ would require some special treatment, as the solutions to the toroidal Laplace equation are only defined for $m, n \geq 0$. There are expressions that relate Legendre functions of negative order or degree back to some positive counterpart, but this would still impact the m- and n-dependent factors in eq. 4.51. Since the $(n+1)$-based expression can be used on any n and m, this will be the choice in the toroidal fit. To avoid the calculation of $m_{\max} \cdot N$ extra terms (for $\tilde{Q}_m^{n_{\max}+1}$), the derivative of $\tilde{Q}_m^{n_{\max}}$ will be determined through the $(n-1)$-based recurrence relation eq. 4.49.

4.4.1.3 Truncated SVD

As discussed above, the matrix $\mathbf{T}$ is both very large in size ($\mathcal{O}(10^7$ entries with 3600 columns), which increases the chances for numerical rank deficiency[17], and has a large dynamic range

[16] The relative error of a number $x = a - b$ is [95]

$$\frac{\Delta x}{x} \leq \max(|\Delta a|, |\Delta b|) \frac{|a| + |b|}{|a - b|}.$$

Catastrophic cancellation occurs when $|a - b| \ll |a| + |b|$, as there is a potential for large amplification of the input errors.

[17] Due to rounding errors and floating point precision, a true singular value of zero is often not present, especially in very large matrices. Any singular values that are smaller than ϵ_m are numerically indistinguishable

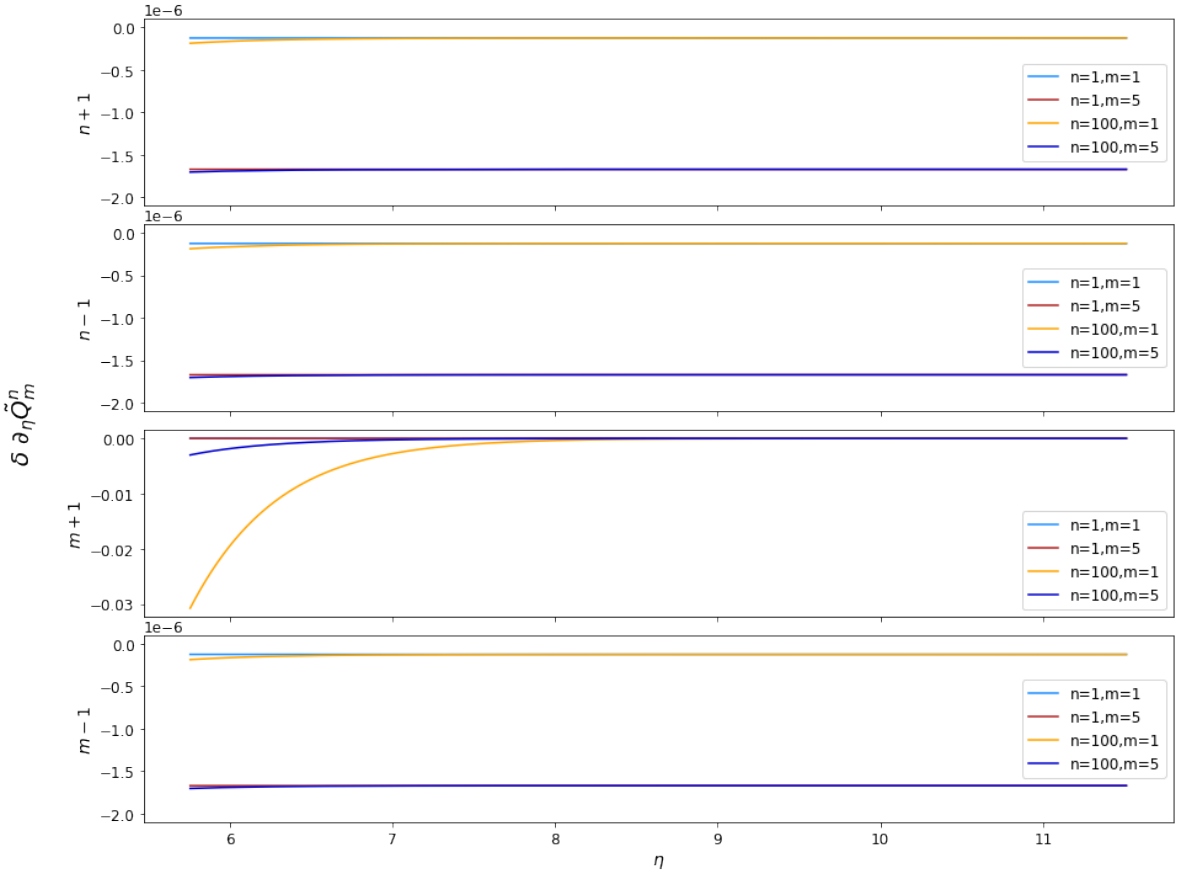


Figure 4.8: Relative differences between analytical derivatives and numerical gradient. Relative differences between the four analytical expressions for $\frac{\partial \tilde{Q}_m^n}{\partial \eta}$ (eqs. 4.48-4.51) and the numerically calculated gradient for $10,000$ η values. $(n \pm 1)$- and $(m-1)$-based expressions show consistent values of $\mathcal{O}(10^{-6})$, whereas the $(m+1)$-based derivative's residual is significantly larger for small η.

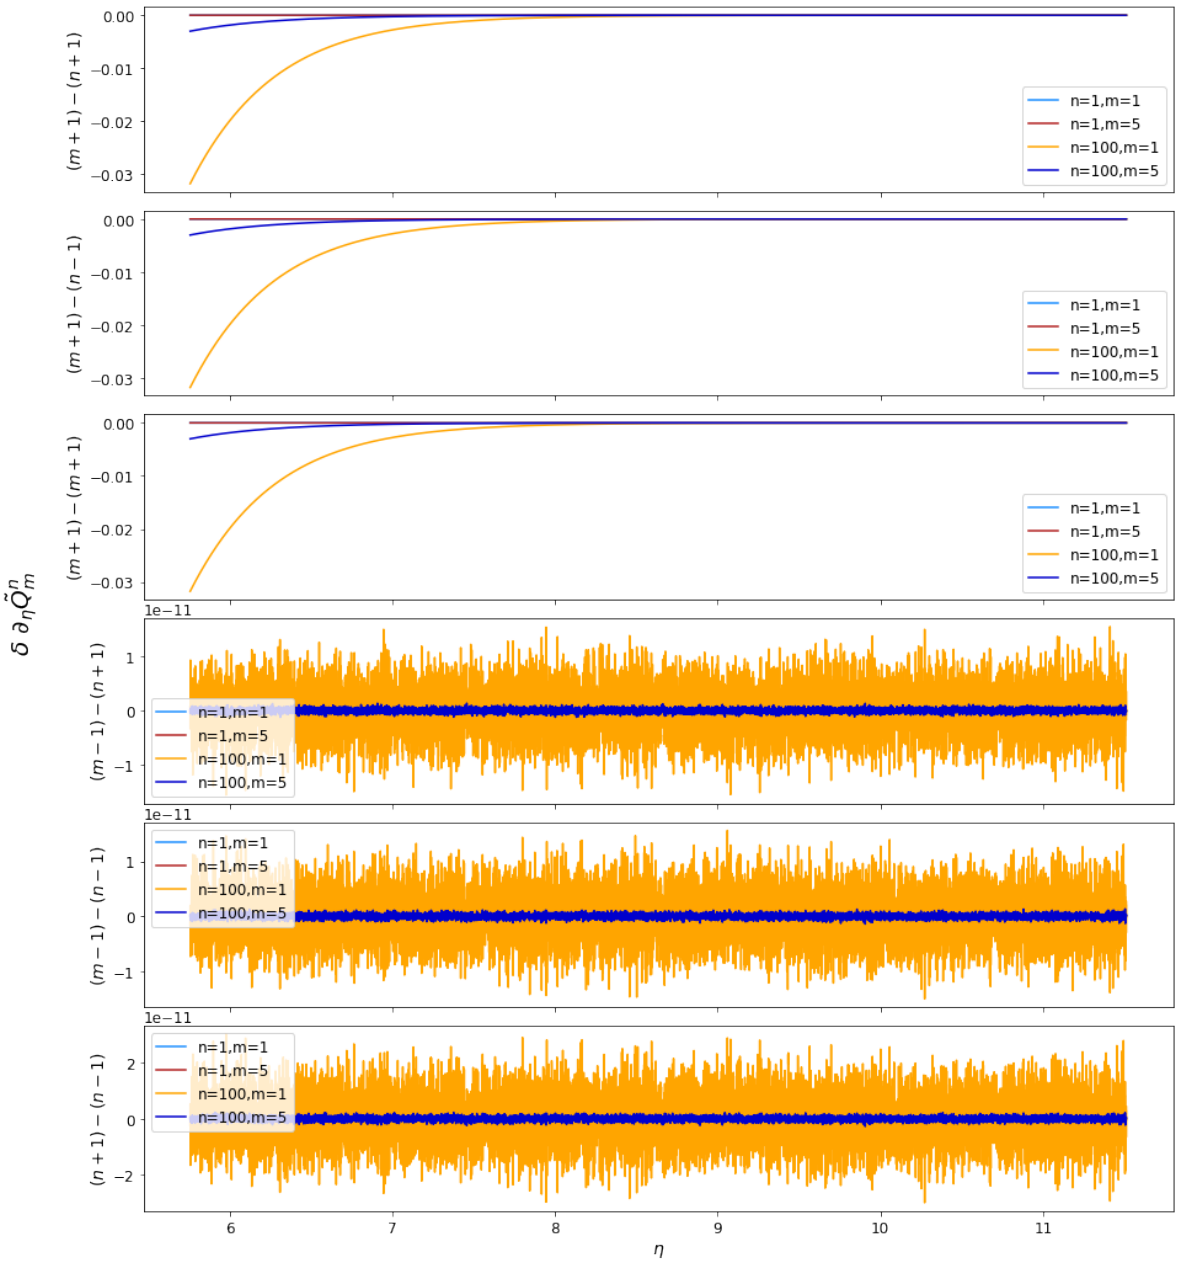


Figure 4.9: Relative differences between each of the analytical derivatives, show that the $(m+1)$-based derivative does not agree with the other three expressions.

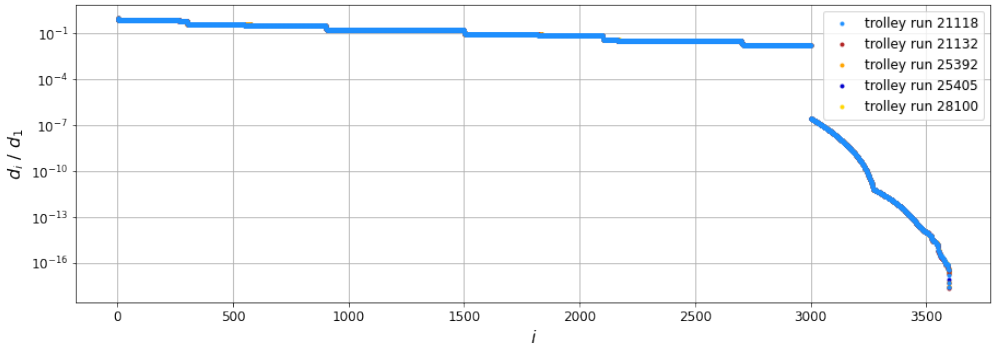


Figure 4.10: Singular value spectrum of $\mathbf{T}$. Normalized singular value spectra for $\mathbf{T}$ constructed for five different trolley runs, which agree extremely well with each other. A significant gap is present for all exactly between $i = 3000$ and $i = 3001$.

(entries differ by about 15 orders of magnitude), which also increases the difference between largest and smallest non-zero singular value and thus the condition number $\kappa(\mathbf{T})$. Eq. 4.16 gave the error of a *full-rank* pseudo-inverse, but a very similar relation holds for rank-deficient matrices as well [135],

$$\frac{\|(\mathbf{T}+\mathbf{\Delta T})^{+}-\mathbf{T}^{+}\|}{\|\mathbf{T}^{+}\|} \leq \frac{1+\sqrt{5}}{2}\,\kappa(\mathbf{T})\frac{\|\mathbf{\Delta T}\|}{\|\mathbf{T}\|} \lesssim \mathcal{O}(\epsilon_m)\;\kappa(\mathbf{T}), \tag{4.52}$$

as long as the perturbation $\mathbf{\Delta T}$ does not *change* the rank of $\mathbf{T}$. Thus it is clear that a large condition number allows for potentially large errors in the pseudo-inverse due to very small changes in the original matrix, which could lead to large errors in the coefficients M_k and the reconstructed magnetic field values. In order to limit the error introduced by the fit algorithm, a regularization method has to be introduced. It is helpful to take a look at the singular value distribution of $\mathbf{T}$, shown for the five trolley runs in fig. 4.10 (normalized by the largest singular value d_1), which reveals a consistent and significant gap between d_{3000} and d_{3001}. An effective regularization choice in this case is the truncated SVD [92], which constrains the condition number by setting all singular values $d_{i>\hat{r}} = 0$, thus defining the rank $\hat{r}$ matrix [92, 86]

$$\mathbf{T}_{\hat{r}} = \mathbf{U}\cdot\mathrm{diag}(d_1,...,d_{\hat{r}},0,...,0)\cdot\mathbf{V}^T = \sum_{i=1}^{\hat{r}} d_i\mathbf{u}_i\cdot\mathbf{v}_i^T, \tag{4.53}$$

where $\mathbf{U}$ and $\mathbf{V}$ are part of the SVD of the full $\mathbf{T}$. Alternatively to choosing a specific effective rank $\hat{r}$, the truncation can be performed based on a relative cut-off value $\delta_{\hat{r}}$, such

from zero, and numerical rank deficiency means such small singular values are present.

that all $d_i < \delta_{\hat{r}}\ d_1$ are set to zero. This functionality is already implemented in most least squares solvers. In this case, based on the singular value spectra of the five trolley runs that show $0.015 \leq d_{3000} \leq 0.017$ and $2.7{\cdot}10^{-7} \leq d_{3001} \leq 2.8{\cdot}10^{-7}$, $\delta_{\hat{r}} = 10^{-6}$ is an appropriate choice which also limits the condition number $\kappa(\mathbf{T}_{\hat{r}}) = \frac{d_1}{d_{\hat{r}}} \leq \delta_{\hat{r}}^{-1} = 10^6$. As long as the gap between $d_{\hat{r}}$ and $d_{\hat{r}+1}$ is larger than the perturbations $\mathbf{\Delta T}$, which is easily true here since $\|\mathbf{\Delta T}\| \leq \mathcal{O}(\epsilon_m)\ d_1 \ll d_{\hat{r}} - d_{\hat{r}+1} = \mathcal{O}(10^{-2})\ d_1$, the perturbation of the truncated pseudo-inverse is then bounded by [92]

$$\frac{\|(\mathbf{T}_{\hat{r}} + \mathbf{\Delta T}_{\hat{r}})^+ - \mathbf{T}_{\hat{r}}^+\|}{\|\mathbf{T}_{\hat{r}}^+\|} \leq 3\frac{\eta_{\hat{r}}}{(1-\eta_{\hat{r}})(1-\eta_{\hat{r}}-\Omega_{\hat{r}})} \lesssim \kappa(\mathbf{T}_{\hat{r}})\ \mathcal{O}(\epsilon_m) \approx \mathcal{O}(10^{-10}), \tag{4.54}$$

where $\Omega_{\hat{r}} = \frac{d_{\hat{r}+1}}{d_{\hat{r}}} \approx \mathcal{O}(10^{-5})$ and

$$\eta_{\hat{r}} = \kappa(\mathbf{T}_{\hat{r}})\frac{\|\mathbf{\Delta T}\|}{\|\mathbf{T}\|} \leq \kappa(\mathbf{T}_{\hat{r}})\mathcal{O}(\epsilon_m). \tag{4.55}$$

For the truncated linear system $\mathbf{T}_{\hat{r}}{\cdot}\mathbf{k}_{\hat{r}} = \mathbf{b}$, coefficients can easily and reliably be computed via the pseudo-inverse (or directly via the SVD), and since $\mathbf{T}_{\hat{r}}$ has full rank, their perturbation bound according to eq. 4.17 is

$$\frac{\|\mathbf{\Delta k}_{\hat{r}}\|}{\|\mathbf{k}_{\hat{r}}\|} \leq \kappa(\mathbf{T}_{\hat{r}}) \left(\frac{\|\mathbf{\Delta T}_{\hat{r}}\|}{\|\mathbf{T}_{\hat{r}}\|} + \frac{\|\mathbf{\Delta b}\|}{\|\mathbf{b}\|}\right) + \mathcal{O}(\epsilon_m^2) \lesssim \kappa(\mathbf{T}_{\hat{r}})\ \mathcal{O}(\epsilon_m). \tag{4.56}$$

Any changes to the reconstructed $\mathbf{b}_{\hat{r}}^{\text{fit}} = \mathbf{T}{\cdot}\mathbf{k}_{\hat{r}}$ due to small perturbations in $\mathbf{T}$, $\mathbf{T}_{\hat{r}}$, $\mathbf{k}_{\hat{r}}$, or $\mathbf{b}$ are well constrained as well,

$$\frac{\|\mathbf{\Delta b}_{\hat{r}}^{\text{fit}}\|}{\|\mathbf{b}\|} \leq \frac{\|\mathbf{\Delta b}\|}{\|\mathbf{b}\|} + \kappa(\mathbf{T}_{\hat{r}}) \left(\frac{\|\mathbf{\Delta T}\|}{\|\mathbf{T}\|} + \frac{\|\mathbf{\Delta T}_{\hat{r}}\|}{\|\mathbf{T}_{\hat{r}}\|} + \frac{\|\mathbf{\Delta b}\|}{\|\mathbf{b}\|}\right) \lesssim \mathcal{O}(\epsilon_m)\ \Big(1 + \kappa(\mathbf{T}_{\hat{r}})\Big). \tag{4.57}$$

However, the truncation of the SVD itself introduces a bias in the reconstruction of $\mathbf{b}$ [92, 86], namely

$$\mathbf{b}^{\text{fit}} - \mathbf{b}_{\hat{r}}^{\text{fit}} = \mathbf{T}\cdot\mathbf{k} - \mathbf{T}\cdot\mathbf{k}_{\hat{r}} = \mathbf{T}\cdot(\mathbf{T}^+\cdot\mathbf{b} - \mathbf{T}_{\hat{r}}^+\cdot\mathbf{b}) = \sum_{i=\hat{r}+1}^{\text{rank}(\mathbf{T})} \mathbf{u}_i(\mathbf{u}_i^T\cdot\mathbf{b}). \tag{4.58}$$

Generalized bounds on the norm of this are potentially very large (e.g. $\frac{\|\mathbf{b}^{\text{fit}}-\mathbf{b}^{\text{fit}}_{\hat{r}}\|}{\|\mathbf{b}\|} \leq \kappa(\mathbf{T})$), so it is helpful to instead calculate directly from data

$$\frac{\|\mathbf{b}^{\text{fit}}-\mathbf{b}^{\text{fit}}_{\hat{r}}\|}{\|\mathbf{b}\|} = \frac{1}{\|\mathbf{b}\|}\sqrt{\sum_{i=\hat{r}+1}^{\text{rank}(\mathbf{T})}(\mathbf{u}_i^T\cdot\mathbf{b})^2}, \tag{4.59}$$

which for all five trolley runs and $\delta_{\hat{r}} = 10^{-6}$ is between $1.2\cdot10^{-7}$ and $1.5\cdot10^{-7}$. While this is an upper bound to how much each individual data point could be biased, the more important conclusion is that only $\mathcal{O}(10^{-7})$ of the norm of $\mathbf{b}$ lies outside of the subspace defined by $\mathbf{T}_{\hat{r}}$. This means the truncated SVD captures the behavior of $\mathbf{b}$ almost entirely, and the remainder lies in directions with small singular values, which cause large amplification factors and extreme sensitivity to small perturbations, making those components practically indistinguishable from noise. Thus this choice of regularization is well justified, and the reconstruction performance will be tested on the STFs to confirm that the effect of truncation bias is negligible.

4.4.2 Free parameters

In any fit function that is based on an infinite series (like Taylor expansions or, in this case, multipole expansions), the number of parameters in the fit is in principle variable. Unless there is some prior knowledge about the underlying signal that is being examined and/or the noise in the data, the choice is somewhat arbitrary. In the toroidal fit, the number of parameters is determined by $i_{\max}$ and $j_{\max}$, where $n = \lfloor\frac{i}{2}\rfloor$ and $m = \lfloor\frac{j}{2}\rfloor$ as introduced in eqs. 4.29-4.31. As mentioned in section 4.3.3, n mainly modulates the variation along azimuth, and m exclusively determines the transverse shape of the field. This allows for independent determinations of the ideal $i_{\max}$ and $j_{\max}$. The same five trolley runs that serve to construct the STFs will be used here.

For the transverse direction, the ability to fit the shape is still restricted by the limited number of data points. With 17 trolley probes, $j_{\max} \leq 16$ is required to retain an "overconstrained" system. To isolate the dependence on j, each transverse slice of a trolley run can be taken as a separate dataset and fitted using only $n = 0$ terms, which reduces the model to a much smaller number of parameters. With this procedure, the reduced χ^2 statistic as well as an F-test can be used to guide the choice of $j_{\max}$. The reduced χ^2 statistic [34]

$$\chi^2_\nu = \frac{\chi^2}{\nu} = \frac{1}{N_d - N_p}\frac{\|\mathbf{b}-\mathbf{b}_{\text{fit}}\|^2}{\sigma^2}, \tag{4.60}$$

where $\nu = N_d - N_p$ is the number of degrees of freedom calculated as the difference between number of data points N_d and the number of model parameters N_p, σ is the uncertainty of the data points (assumed to be the same for all points), and $\|\mathbf{b} - \mathbf{b}_{\text{fit}}\|^2$ is the squared sum of residuals, is a standard measure of goodness of fit. The model that minimizes the residual of the least squares problem will also minimize the (regular) χ^2 statistic. When normalized by the degrees of freedom, a value $\chi^2_\nu \approx 1$ indicates a model that is well-suited to describe the data, whereas values larger than one indicate underfitting, and values smaller than one can indicate overfitting [34]. To evaluate whether adding additional model parameters will improve the model at a statistically significant level, an F-test can be performed. This test compares the observed value of the F-statistic [34],

$$F_\chi = \frac{\chi^2(N_p) - \chi^2(N_p + 1)}{\frac{\chi^2(N_p+1)}{N_d - N_p - 1}}, \tag{4.61}$$

to the F-distribution f, which describes the range of values of F under the hypothesis that any improvements in the fit are due to statistical fluctuations rather than the added parameter. Then the probability of the observed F-statistic being due to to random noise is given by the integral probability (also known as p-value) [34],

$$p = \int_{F_\chi}^{\infty} p_f(f; 1, N_d - N_p - 1)\mathrm{d}f, \tag{4.62}$$

where p_f is the probability density function of f, 1 indicates one added model parameter, and $N_d - N_p - 1$ the degrees of freedom of the larger model. This means a large p-value suggests the added model parameter does not improve the fit in a meaningful way, and a small p-value is a sign that the fit does indeed benefit from the added complexity.

In the toroidal fit, not all model parameters (i.e. coefficients) are truly restricting the degrees of freedom of the system, which is indicated by $\mathbf{T}$ not having full rank (either because the corresponding basis functions are not linearly independent - which can be the case numerically even if not mathematically - or because they lie in the nullspace of $\mathbf{T}$). Using the truncated SVD will remove even more coefficients from the parameter space by artificially reducing the rank of $\mathbf{T}$. Thus what is required to determine the reduced χ^2 statistic is the *effective* number of parameters, $N_p = \text{rank}(\mathbf{T})$ [93].

Fig. 4.11a shows both the reduced χ^2 values for $2 \leq j_{\max} \leq 16$ for each trolley run, as well as the p-values for adding one parameter at a time starting at $j_{\max} = 3$ (compared to $j_{\max} = 2$. The F-statistic is undefined if adding a parameter to the model does not increase the rank if $\mathbf{T}$, as is the case for going from $j_{\max} = 2$ to $j_{\max} = 3$, and from $j_{\max} = 6$ to $j_{\max} = 7$, hence there is also no p-value. For $12 \leq j_{\max} \leq 14$, the χ^2_ν values are minimal and

$\mathcal{O}(1)$, and the p-values are starting to trend upwards at $j_{\max} \geq 12$. Therefore $j_{\max} = 12$ is chosen as the number of transverse parameters, having both a sufficiently small p-value, and the smallest χ^2_ν for 4 out of the 5 trolley runs. For comparison, fig. 4.11b shows the same quantities for the 2D multipole parameterization (calculated in the same way on the same data), which show the reduced χ^2 minimum at $i_{\max} = 12$[18], consistent with what was determined as the optimal number of parameters in the Muon g-2 experiment [51]. While the p-value is larger than 0.5 at $i_{\max} = 11$, it drops again at $i_{\max} = 12$, which indicates that even if the previous parameter did not significantly improve the model, this one does.

For the longitudinal direction, the problem can effectively be reduced to fitting a Fourier series to the azimuthal spread of each of the trolley probes. Due to the large number of data points however, neither the χ^2_ν nor the F-test provide reliable guidance to find the optimal $i_{\max}$. Instead, for each of the five trolley runs, a visual inspection of FFTs of all trolley probes (shown in appendix C, fig. C.2) is used to determine $n_{\max}$[19]. The FFTs show a consistent leveling off around $n \lesssim 150$, which indicates the start of the noise floor and means fitting higher frequencies is likely overfitting. Therefore $j_{\max} = 2n_{\max} = 300$ is chosen for the toroidal fit.

4.4.3 Simulated toroidal fields

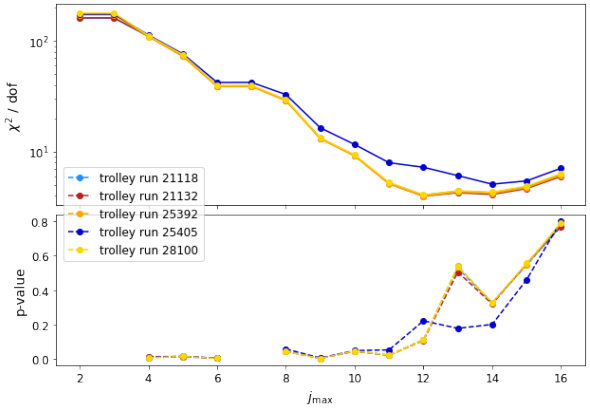
To evaluate the fit results of both the toroidal fit and the 2D multipole fit, several magnetic fields were simulated using the toroidal parameterization, which are summarized in table 4.2. Five measured trolley runs served as inputs, as well as the radial and longitudinal field maps taken in 2016 after passive shimming was completed [126], shown in fig. 4.12. The STFs were constructed as follows:

- Each STF is a toroidal fit to *all three* magnetic field directions simultaneously. First, a given radial and longitudinal field map is linearly interpolated to the trolley run azimuthal positions. Then $\mathbf{T}$ is constructed from eqs. 4.34-4.36 (three matrix rows per position, one for each field direction) and a toroidal fit is performed. The coefficients from this fit are the defining component of the STFs, as they can be used to reconstruct the field not only at the original measurement positions, but anywhere in the ring.

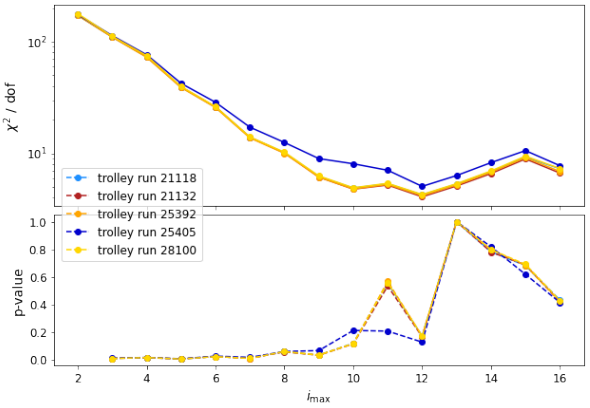
- From each STF's coefficients, the three field components are calculated at positions that are representative of a typical trolley run, and combined to form the magnetic field

[18] The $i_{\max}$ of the 2D parameterization corresponds to the number of included m_i from table 4.1 and should not be confused with the $i_{\max}$ of the toroidal expansion.

[19] When performing an FFT of data distributed on an angular grid, the x-axis will have units of cycles per radians, which is exactly n in eq. 4.31.



(a) Toroidal fit with $n = 0$, $2 \leq j_{\max} \leq 16$, where $m = \lfloor j/2 \rfloor$. χ^2_ν is minimal (and $\mathcal{O}(1)$) between $12 \leq j_{\max} \leq 14$, whereas p-values start to increase at $j_{\max} = 12$.



(b) 2D multipole fit, $2 \leq i_{\max} \leq 16$, corresponding to m_i in table 4.1. χ^2_ν is minimal (and $\mathcal{O}(1)$) at $i_{\max} = 12$, and p-values increase at $i_{\max} \geq 10$ but have another local minimum at $i_{\max} = 12$.

Figure 4.11: F-test and χ^2_ν values for determining the number of transverse parameters. Reduced χ^2 and F-test p-values for five trolley runs, used to determine the ideal number of transverse parameters in the toroidal fit (left) and the 2D multipole fit (right) for comparison.

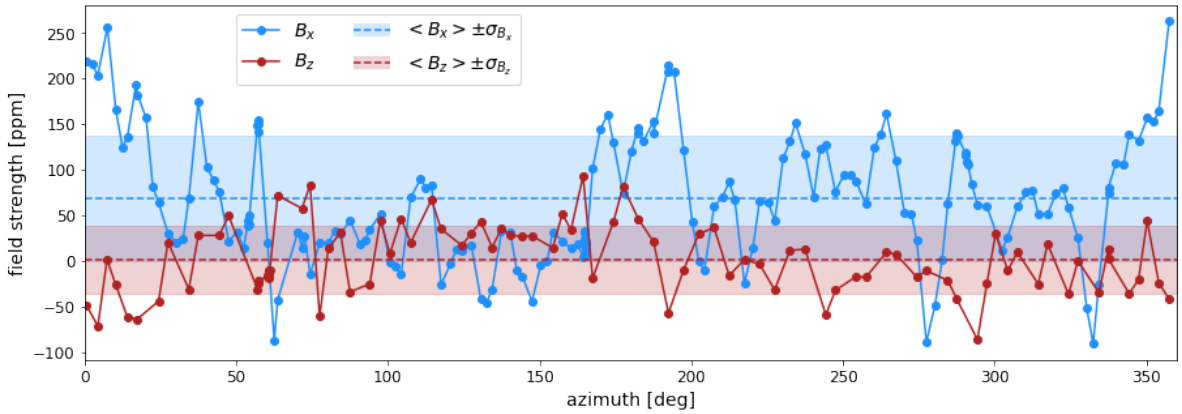


Figure 4.12: 2016 radial and longitudinal field measurements, reproduced from [126]

magnitude that an NMR probe would measure. These simulated trolley runs (as well as the underlying field components) are then used to perform either a 2D multipole fit or another toroidal fit, and the fit performance is assessed using the STF "truth" values for comparison

- **STF1** uses the original and unchanged 2016 radial and longitudinal field measurements as input, with 15 transverse and 350 longitudinal fit parameters for a total of 5250 parameters.

- **STF2** uses the 2016 measurements, but shifted/scaled to different means and variances. The radial field mean in this case would be consistent with the radial field during Run periods 1-3 [10] when active shimming was also present, while the longitudinal field mean is small but larger than measured in 2016. Both variances are significantly larger than the 2016 measurements. This is to simulate the hypothetical case in which a reduction of the radial field mean increases the longitudinal field, and the active shimming has a positive effect on the average field but a negative effect on the variances. The number of parameters is the same as for STF1.

- **STF3** again uses the unchanged 2016 measurement values, but inverts them and shifts them in opposite directions about halfway around the storage ring, to test a different correlation between the radial and longitudinal field components. The parameters are unchanged from STF1 and STF2.

- **STF4** shifts the radial field mean to less than 1ppm, which is consistent with run conditions during Run-4/5/6, and scales the variance down slightly. To test the effect of a large longitudinal field, the longitudinal mean is set to about half of the 2016 radial field mean, but shifted in the opposite direction, with a mildly smaller variance. Again, the number of parameters is the same as for STF1-STF3.

- **STF5** uses small means (2016 measurements shifted to $<$ 1ppm) and variances on the same order of magnitude as the original measurements, but enforces higher order transverse components. To do this, B_y is calculated for the higher orders ($i = 20, 25$ for $j = 1, 12, 53$) with a coefficient of 1. Then the coefficients are chosen such that when the total field of those components is subtracted from the measured field, it has a noticeable effect without completely distorting it. The radial and longitudinal field values of these higher order components are calculated to confirm that they are negligible, and then the remaining B_y as well as B_x and B_z are fitted with 8 transverse parameters and the usual 350 longitudinal parameters. The total vertical field is the result of the fit plus the higher order components.

- **STF6** uses the same strategy as STF5, with a higher order longitudinal components ($j = 522$ with $i = 5, 12$). The fit of the remainder is performed with 15 transverse and 100 longitudinal components. Both STF4 and STF5 are introduced to investigate the efficacy of the toroidal fit if higher orders than the included parameters are present.

- **STF7** is the only simulated field that does not rely on the 2016 radial and longitudinal field measurements as input. Instead, a random number generator (`numpy.random.normal`) is used to create data samples for a similar number of longitudinal points as in the 2016 measurements, which are then again linearly interpolated to match the trolley measurement positions. This procedure was chosen to avoid too much "high frequency noise" in the data, which would appear if the random data were generated directly for the trolley positions. The means are consistent with Run-4/5/6 conditions, and the variances slightly larger than the original 2016 measurements, which is likely an overestimation. The number of parameters is the same as for STF1-STF4.

Graphs of the radial and longitudinal field input data used for each STF, as well as the resulting field components and magnitude, can be found in appendix C.

Name	STF1	STF2	STF3	STF4	STF5	STF6	STF7
Trolley run	21118	21118	21132	21132	25392	25405	28100
$\langle B_R \rangle$ [ppm]	68.50	20.00	68.50	0.95	0.50	−0.20	0.95
σ_{B_R} [ppm]	68.89	150.00	68.89	50.00	50.00	50.00	75.30
$\langle B_\phi \rangle$ [ppm]	0.14	5.00	0.14	−30.00	−0.20	0.50	0.20
σ_{B_ϕ} [ppm]	36.93	100.00	36.93	25.00	50.00	50.00	41.14
N_{azimuth}	4333	4333	4330	4330	4330	4330	4329
$i_{\max}$	350	350	350	350	350	100	350
$j_{\max}$	15	15	15	15	8	15	15
higher order i					20, 25	5, 12	
higher order j					1, 12, 53	522	

Table 4.2: Summary of the simulated toroidal fields, created via toroidal fits of trolley runs and radial and longitudinal field data. STF5 and STF6 contain artificially fixed higher order components, whose parameters are listed in the last two lines.

4.4.4 Toroidal fit results

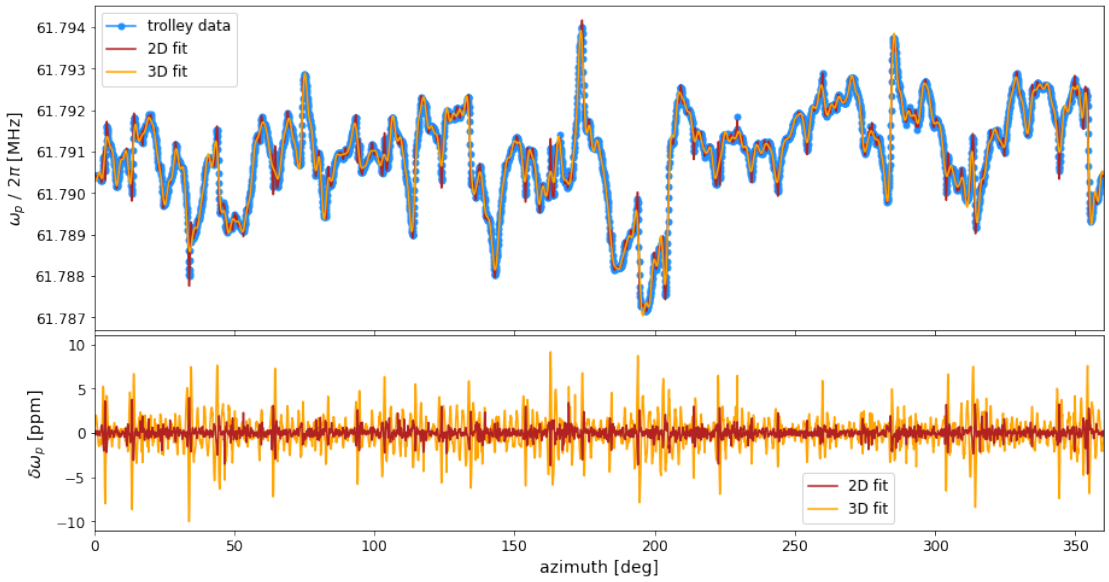
The quantitative results of the toroidal fits of all five trolley runs, as well as the seven STFs, will be presented in direct comparison with the 3D multipole fits in the next section (4.5). Here however, the qualitative performance of the fit algorithm will be briefly summarized. As mentioned already in section 4.4.1.3, the residuals to the measured data that is being fit turns out to be unaffected by the cut-off ratio of the singular values at the sub-ppb level, which confirms the stability of the fit algorithm and reliability of the results.

Toroidal fits to the five trolley runs all produce almost identical results. As an example, the fit results for the central probe and one of the outer probes for trolley run 25392 is shown in fig. 4.13 (alongside the 2D fit for comparison). The main feature of the toroidal fit is a smoothing effect in azimuth, due to the limit on the number of longitudinal parameters. This leads to larger residuals than is the case for the 2D multipole fit, which is especially noticeable for the outer probes (see fig. 4.13). As will be shown in the quantitative analysis below, this does lead to slightly worse fit performance compared to the 2D fit, although not to a level that impacts the general reliability of the toroidal fit.

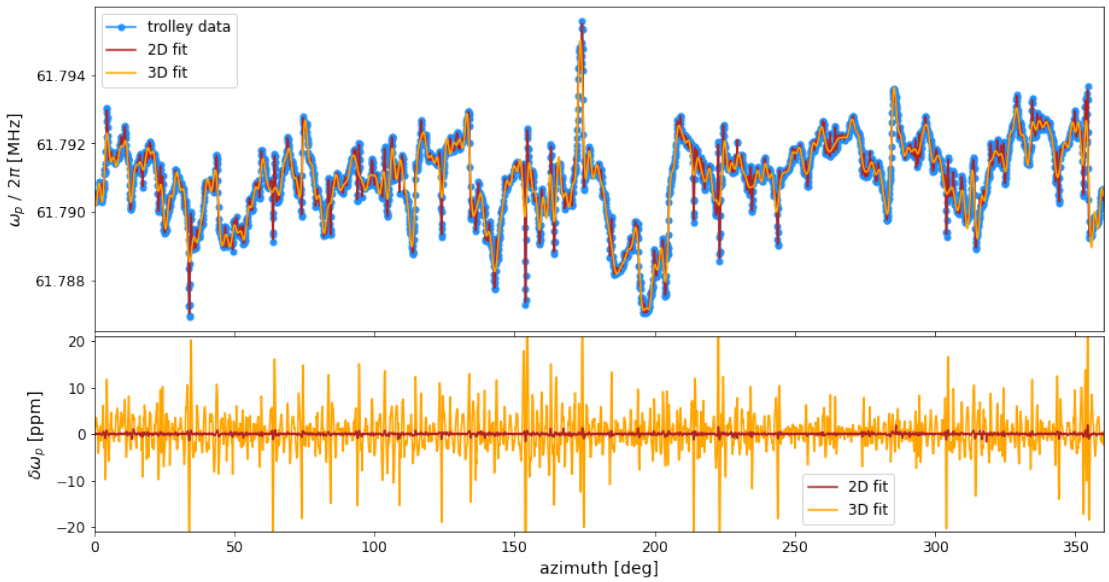
Similarly to the trolley runs, the toroidal fits of the STFs show very little variation for STF1-STF4 and STF7, which suggests that the fit result is not or barely dependent on the size, shape, or variance of the radial and longitudinal field components. Some results for STF1 are shown as an example in fig. 4.14. Since the simulated data contains only the true field magnitude values and no noise, the residuals are smaller than for the trolley runs (by about a factor of 2). The fits to STF5 and STF6 both show some interesting behaviors due to the artificially injected higher order components. In the case of the transverse higher order components (STF5), the mid-range probes 2-5 (cf. fig. 3.10) appear uniquely qualified to reveal the presence of those components. Typically, the mid-range probes follow the center probe fairly closely, and the outermost probes are the most sensitive to deviations from a constant dipole field. However, as shown in fig. 4.15, in this case the residuals of probe 4 (as a representative example) clearly exhibit a modulation of the same shape as the transverse higher order components that are present in STF5 (cf. fig C.7), whereas the toroidal fit residuals of the outer probes are dominated by higher frequency artifacts that obscure the shape.

In the case of the longitudinal higher order components (STF6), it is the outermost probes who can detect their presence, with the toroidal fit residuals showing the high frequency signal of the injected higher order terms (cf. fig. C.8). Additionally, since in this case the remaining longitudinal terms included in STF6 are fewer than the fit paramters, the central probe fit is able to reconstruct the field at that location extremely well. Fig. 4.16 shows both the center probe fit as well as one of the outer probes as an example.

The results of these two simulated cases suggest that the toroidal fit can potentially indicate when higher order components are present that are currently not included in the fit (as long as they are larger than the noise level on the data). The quantitative analysis below will further show that despite these higher order terms, the toroidal fit can still reconstruct the field to an acceptable level.

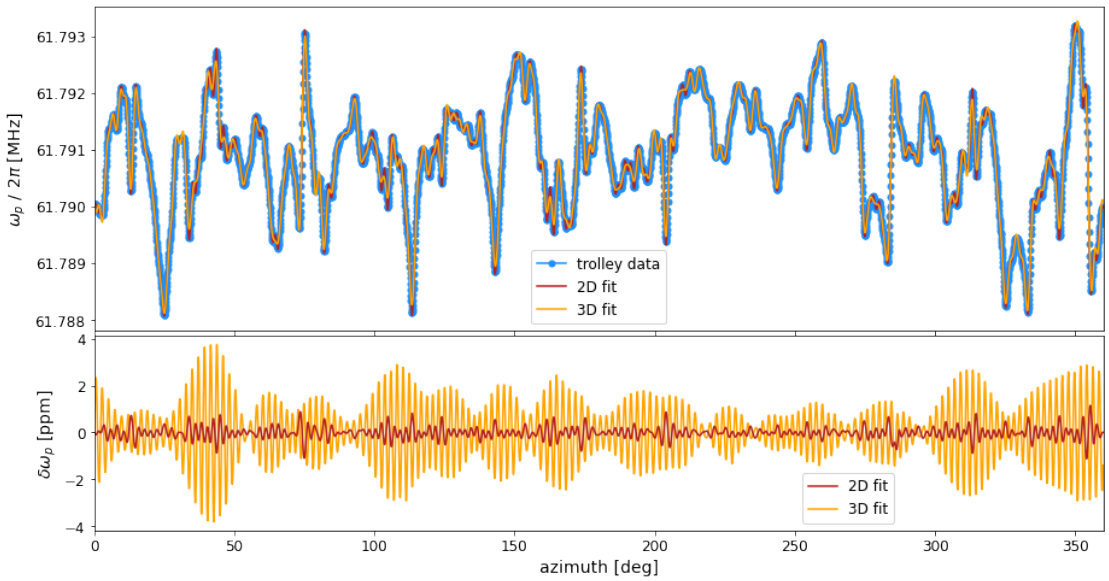


(a) Central trolley probe (probe 1)

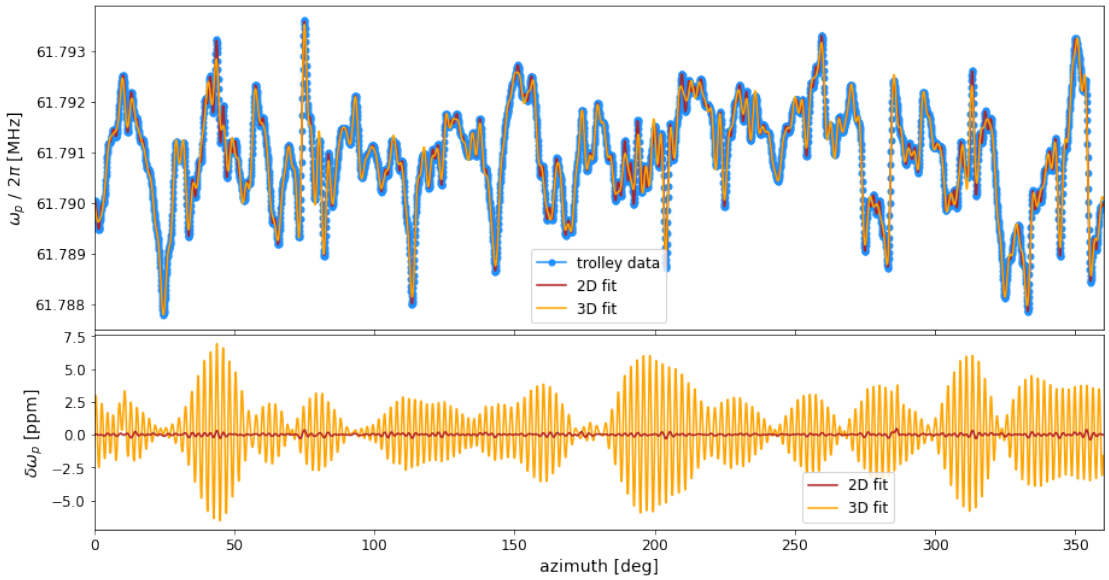


(b) Bottom trolley probe (probe 6, one of the outermost probes)

Figure 4.13: Example fit results for two trolley probes. Exemplary results of toroidal and 2D multipole fits of trolley run 25392, probes 1 (center) and 6 (bottom). For each probe, the top shows the absolute values of the measured ω_p, overlayed with the reconstructed fields from the fits, whereas the bottom shows the residuals between each of the fits and the measured data, relative to B_0 as usual. The toroidal fit produces a smoothed out result of the input data, even compared to the 2D fit, which is also visible in the increased residuals, in particular for the outer probes.

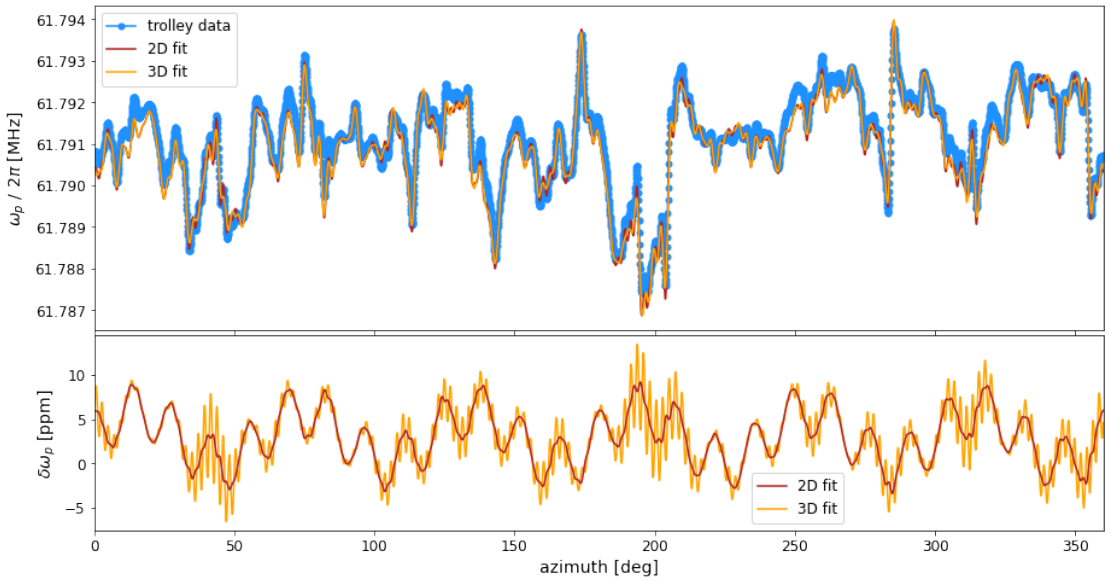


(a) Central trolley probe (probe 0)

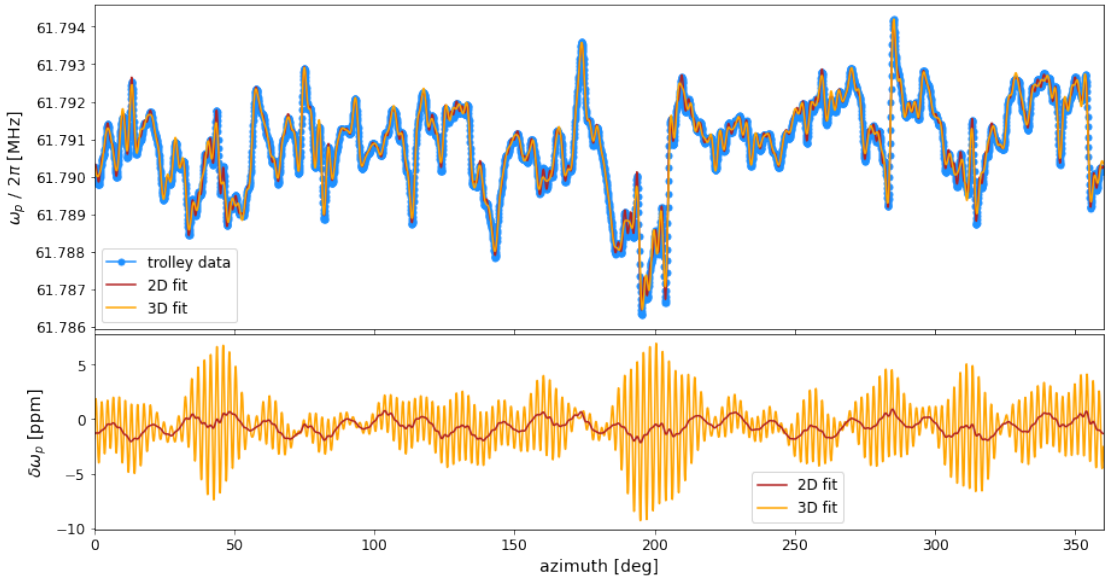


(b) Top trolley probe (probe 12, one of the outermost probes)

Figure 4.14: Example fit results for STF1. Exemplary results of toroidal and 2D multipole fits of STF1, at locations of probe 1 (center) and 12 (top). For each probe location, the top shows the absolute values of the true simulated ω_p, overlayed with the reconstructed fields from the fits, whereas the bottom shows the residuals between each of the fits and the simulated data, relative to B_0 as usual. The results are similar to the trolley run fits, with the residuals being about a factor of 2 smaller due to the lack of noise in the simulated data.

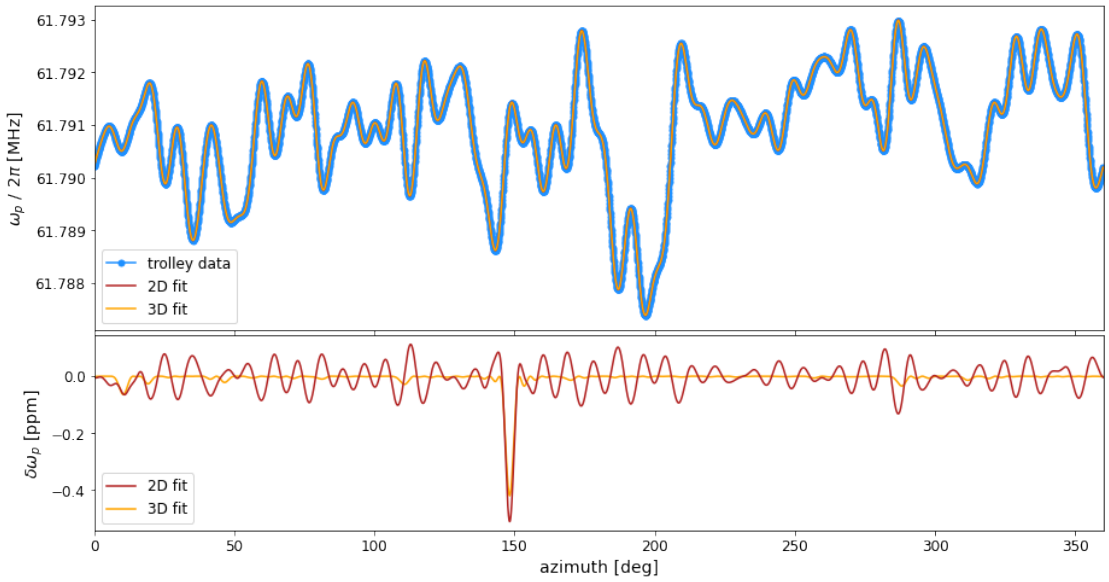


(a) Trolley probe between top and central probes (probe 4, one of the mid-range probes)

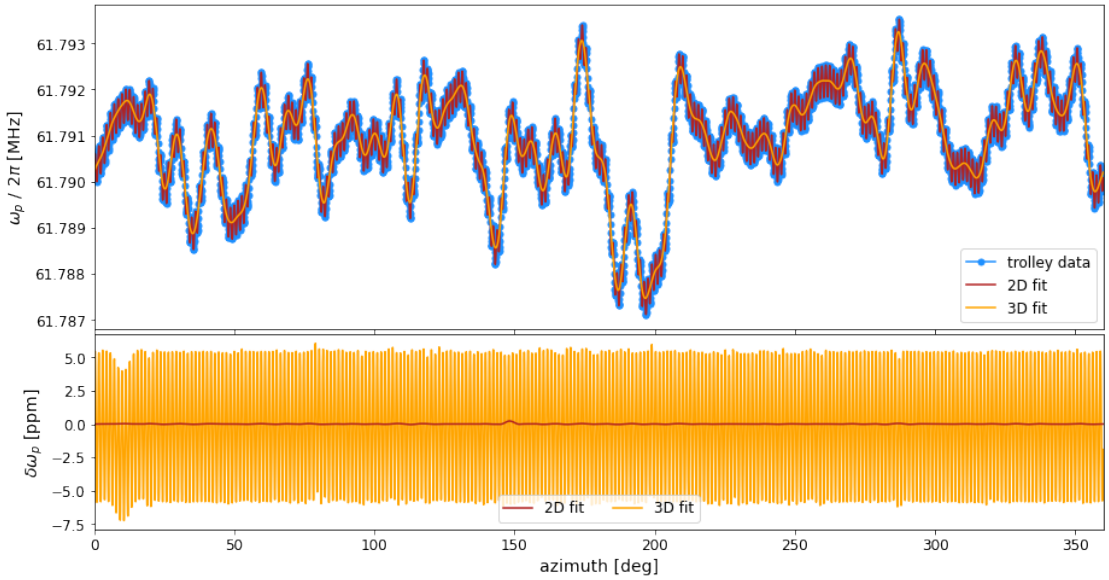


(b) Top trolley probe (probe 12, one of the outermost probes)

Figure 4.15: Example fit results for STF5. Exemplary results of toroidal and 2D multipole fits of STF5, at locations of probe 4 (between center and top) and 12 (top). For each probe location, the top shows the absolute values of the true simulated ω_p, overlayed with the reconstructed fields from the fits, whereas the bottom shows the residuals between each of the fits and the simulated data, relative to B_0 as usual. The 2D fit residuals are significantly larger than for STF1, and in particular for the mid-range probe location, both the 2D and toroidal residuals show a similar shape as the transverse higher order components that are present in STF5 (cf. fig C.7).



(a) Central trolley probe (probe 0)

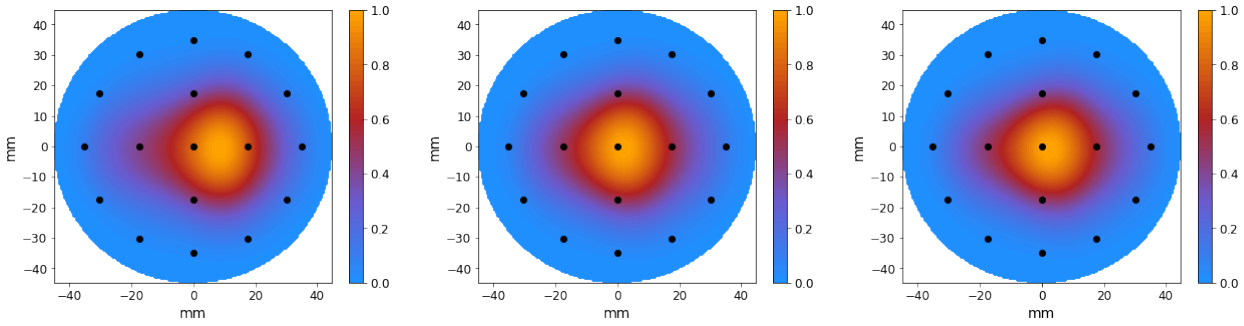


(b) Left trolley probe (probe 15, one of the outermost probes)

Figure 4.16: Example fit results for STF6. Exemplary results of toroidal and 2D multipole fits of STF6, at locations of probe 1 (center) and 15 (left). For each probe location, the top shows the absolute values of the true simulated ω_p, overlayed with the reconstructed fields from the fits, whereas the bottom shows the residuals between each of the fits and the simulated data, relative to B_0 as usual. The central residuals are significantly smaller for both fits than any of the other STFs, and the toroidal fit in particular follows the simulated data extremely closely, whereas the 2D fit residuals seem to exhibit the same shape as the radial field at the outer probe locations (cf. fig. C.8). On the other hand, the toroidal residuals at the outer probe locations clearly show the longitudinal higher order component.

4.5 Validation of 2D multipole fits

This section presents a direct comparison between 2D multipole and toroidal fit results, as well as two checks of systematic effects essential to the 2D multipole fit procedure using the STFs. The ultimate quantity of interest is the total spatial field average - with the caveat that this should be the *muon-weighted* average (as touched on in the previous chapter, section 3.2.1). The procedure to determine the true muon-weighted field average is non-trivial, and will be explored in depth in chapter 5. For the purpose of the evaluation of the fit results, a simpler but approximated approach will be used. As in other systematic studies of the field, weights for each probe are calculated from RF-dataset[20] averages of the beam distribution within a 2mm radius area around each probe location (in the transverse plane) [51]. The three relevant beam distributions are shown in fig. 4.17: the noRF average for trolley runs 21118 and 21132, and the derived STFs STF1-STF4; the xRF average for trolley runs 25392 and 25405, as well as STF5 and STF6; and the xyRF6 average for trolley run 28100/STF7. Thus in the following, any results will be reported as "mean", which indicates a standard, unweighted mean along transverse and azimuthal directions, or as "avg", which designates the estimated muon-weighted average (calculated as the weighted average in the transverse plane using the weights derived from the beam distributions, followed by a standard mean along azimuth). If a variance or standard deviation is listed, it is the usual, unweighted version.



(a) noRF dataset (b) xRF dataset (c) xyRF6 dataset

Figure 4.17: Average RF-dataset beam distributions. Average beam distributions for noRF, xRF, and xyRF6 datasets, normalized. Black dots indicate trolley probe positions.

[20] A table defining the RF-datasets can be found in appendix A.

4.5.1 Effect of non-vertical field components

Before considering the validity and performance of the fitting algorithms, it is prudent to verify the critical underlying assumption that $|\vec{B}| \approx B_y$, which is possible using the STFs. In particular, while eq. 4.1 of course holds, it needs to be shown that localized larger non-vertical field values don't have unexpected impacts. To assess the effects of the radial and longitudinal field contributions, the standard arithmetic mean or all data points, and the approximated muon-weighted average[21] of $|\vec{B}| - B_y$ of each of the STFs are shown in table 4.3, for two sets of data points: Once for the simulated trolley measurement positions only, and once for a dense grid of points (0.5 mm between points) in the transverse plane for each azimuthal position in the simulated trolley measurement, which is a better representation of the whole muon storage volume. The mean values for both sets stay $\mathcal{O}(10^{-8})B_0$, although the dense grid shows larger means for most STFs, since this takes values at the edge of the storage ring into account, and the radial field in particular tends to get larger with increased distance from the central orbit. The muon-weighted average shows that, as expected, those larger differences between $|\vec{B}|$ and B_y at the edges are negligible due to the small number of muons that reach those areas. Thus overall the assumption is valid at the level stated in eq. 4.1.

$\lvert B\rvert - B_y$	**trolley measurement positions**		**fine transverse grid**	
Dataset	mean [ppb]	weighted avg [ppb]	mean [ppb]	weighted avg [ppb]
STF1	17	9	29	7
STF2	28	20	40	18
STF3	17	9	30	7
STF4	13	5	26	3
STF5	3	3	3	3
STF6	10	4	17	3
STF7	15	7	28	5

Table 4.3: $|\vec{B}| - B_y$ arithmetic mean and muon-weighted average for each STF calculated for the set of simulated trolley measurement positions, and a denser grid of data points capturing the entire muon storage volume. Values confirm eq. 4.1.

[21] The muon-weighted average is used here to gauge the error if the fit reconstructs B_y instead of $|\vec{B}|$.

4.5.2 Difference between 2D and 3D fits

	2D fit residuals		3D fit residuals	
Dataset	mean [ppb]	weighted avg [ppb]	mean [ppb]	weighted avg [ppb]
STF1	0	−8	0	−4
STF2	0	−8	0	−4
STF3	0	−8	0	−4
STF4	0	−8	0	−4
STF5	0	9	0	9
STF6	0	−5	0	−5
STF7	0	−9	0	−7

Table 4.4: Results of 2D and 3D fits to simulated toroidal fields

	2D fit residuals		3D fit residuals	
trolley run	mean [ppb]	weighted avg [ppb]	mean [ppb]	weighted avg [ppb]
21118	0	5	0	9
21132	0	5	0	9
25392	0	9	0	10
25405	0	1	0	2
28100	0	6	0	9

Table 4.5: Results of 2D and 3D fits to trolley runs

The main purpose of the toroidal field parameterization is to validate the 2D multipole fit results against the STFs, and to compare the fit performance of the toroidal and 2D fit. A summary of this comparison is given in table 4.4, where the residuals of the 2D and 3D fit results to the true $|\vec{B}|$ values for the simulated trolley measurement positions are reported (again as standard mean across all points, and muon-weighted average). All standard means of both fits are several orders of magnitude smaller than 1ppb and therefore considered 0. The four STFs that have larger non-vertical field components show a consistent discrepancy of 4ppb between the weighted averages of the 2D and 3D fit results, whereas this difference reduces to half that for STF7, and for the two STFs with the smallest non-vertical field

contributions the fits agree.
In table 4.5, a summary of the residuals of both fits are shown for each of the five original trolley runs, i.e. for each of the 17 trolley probes at the interpolated azimuthal positions. The notable difference compared to the STF fits is the fact that the toroidal fit now has consistently larger residuals than the 2D fit, when that was the opposite for the STFs. Overall the order of magnitude of the residuals is the same between fits to the real trolley data and fits to the STFs, which signals that both fits are able to recover the shape of the magnetic field despite noise in the data, although further studies, especially simulated trolley measurements that include position uncertainties as well as other noise sources, would be warranted to fully understand the behavior of the toroidal fit here. Considering the differences between the 2D and 3D fits for both the simulated datasets and the measured trolley run data, it seems safe to assume they generally agree to within $\pm 5\,\text{ppb}$.

4.5.3 Azimuthal position interpolation

Another test that is possible with the STFs is a cross-check of the interpolation to common azimuthal positions for all probes. To do this, the sequential readout of the trolley probes is simulated with an average trolley speed of 0.083 degrees per second, and a 34 Hz readout frequency (2 Hz per probe, giving a data point roughly every $2 \cdot 10^{-3}$ degrees). The field values are then interpolated to the azimuthal positions of the trolley run that served as input for the corresponding STF, using the same principle of linear interpolation and numerical integration outlined in section 4.4.1.1 [140]. Due to small variations in real-life trolley speed during trolley runs, the 1 s grid does not equal a uniform grid in azimuth. While the simulated sequential sampling *is* constant in azimuth, the same effect is achieved by interpolating to the non-uniformly spaced sample points provided by the original trolley run. Then a 2D or 3D fit can be performed using the interpolated data, and be compared to the *true* magnetic field values at those azimuthal positions. The results are given in table 4.6 in the same format as above. The standard means are larger than for a fit directly to the true data (although still sub-ppb and considered 0), but the muon-weighted averages are exactly the same. This is an excellent confirmation of the validity of this approach.

	2D fit residuals		**3D fit residuals**	
Dataset	mean [ppb]	avg [ppb]	mean [ppb]	avg [ppb]
STF1	0	−8	0	−4
STF2	0	−8	0	−4
STF3	0	−8	0	−4
STF4	0	−8	0	−4
STF5	0	9	0	9
STF6	0	−5	0	−5
STF7	0	−9	0	−7

Table 4.6: Test of 2D and 3D fits to interpolated data (simulation)

CHAPTER 5

Muon Weighted Field Average

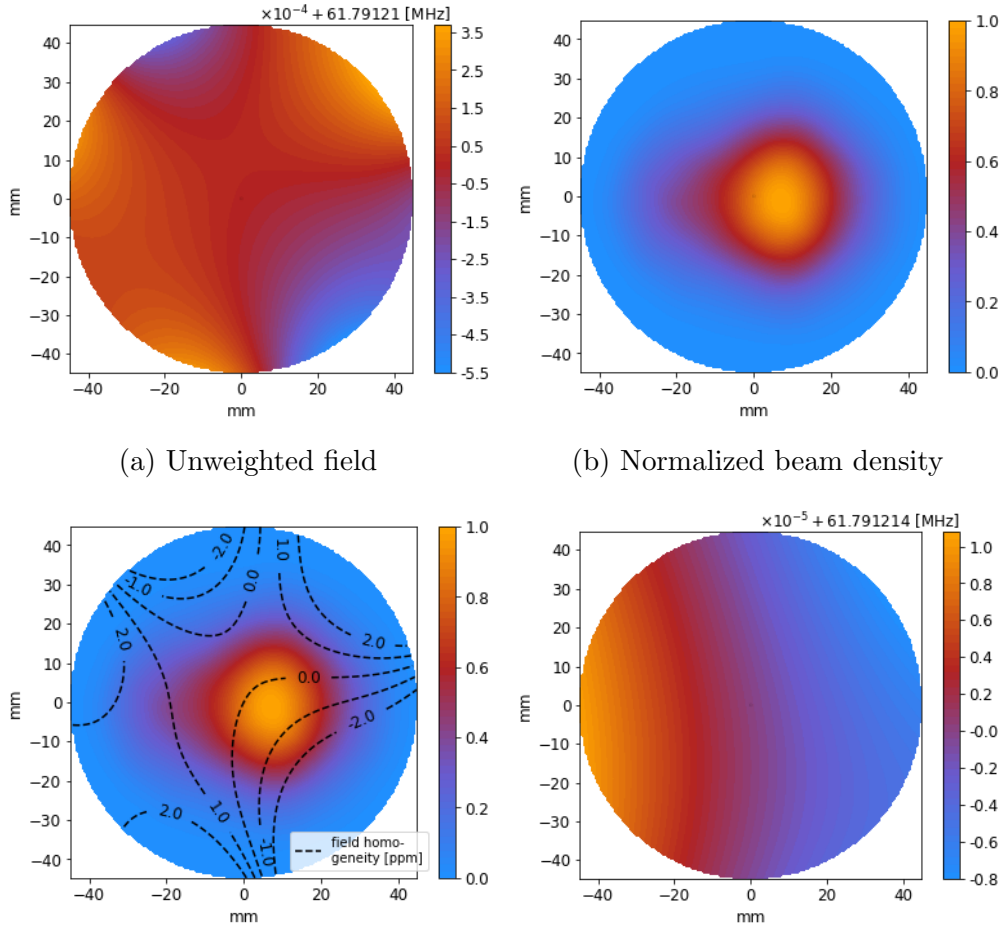
As was touched on in chapter 3.2.1, to determine the muon g-factor via eq. 3.7, the magnetic field *that the muons sample* is needed. While the ensemble of muons samples approximately the entire storage volume, the muon density varies significantly in the transverse plane, and a simple spatial average of the field would not provide the appropriate result. The procedure of averaging the magnetic field in a way that accounts for the distribution of the muon beam is called "muon weighting", because it is effectively a weighted average of the field. Specifically, the field has to be weighted by those muons that are detected and enter the omega-a analysis. This means taking into account spatial and temporal muon number variations, as well as applying data quality cuts (DQC) coherently to all types of data. DQC are used to ensure consistent muon storage conditions across combined datasets [6] and typically include selection criteria from all detectors as well as several monitor systems. Fig. 5.1 shows an example (at an arbitrary time and location) of the field and beam before the muon weighting, and the resulting weighted field, to illustrate the principle.

Remember that, to first order (i.e. neglecting velocity-dependent terms in eq. 3.5, and assuming $B = |\vec{B}| \simeq B_y$ as in chapter 4), the anomalous precession frequency is (cf. eq. 3.3)

$$\omega_a = -a_\mu \frac{q}{m} \tilde{B}, \tag{5.1}$$

where $\tilde{B}$ denotes the muon-weighted magnetic field.
A single muon on trajectory $\vec{X}(t)$ in the storage volume with instantaneous anomalous precession frequency ω_a will acquire a phase $\Delta\phi(t)$ between its spin and momentum from time t_0 to t [118] given by

$$\Delta\phi(t) = \phi(t) - \phi(t_0) = \int_{t_0}^{t} \omega_a(t')\mathrm{d}t' = -a_\mu \frac{q}{m} \int_{t_0}^{t} B(\vec{X}(t'))\mathrm{d}t'. \tag{5.2}$$



(a) Unweighted field

(b) Normalized beam density

(c) Normalized beam density overlayed with unweighted field homogeneity (ppm relative to mean)

(d) Weighted field

Figure 5.1: Muon weighting illustration. Example cross-sectional view, illustrating the muon weighting: The unweighted field (a) and beam distribution (b) are combined to form the muon-weighted field (d). Fig. (c) overlays the unweighted field and beam to show that most of the muons experience a magnetic field within ±1ppm of the mean, and almost all within ±2ppm. The average of the entire cross-sectional slice weighted by the muon density is equal to the unweighted (area) average of the weighted field in (d).

The magnetic field along the trajectory can equivalently be written as an integral over the entire storage volume V, using a delta function for the single muon:

$$B(\vec{X}(t')) = \int_V B(\vec{x})\ \delta(\vec{x} - \vec{X}(t'))\mathrm{d}^3x \tag{5.3}$$

For an ensemble of muons which sample the storage ring field many times during their lifetimes, the delta function can be replaced with the probability density Π of the muon ensemble, and the ensemble averaged phase becomes [118]

$$\langle\phi(t)\rangle_\mu = -a_\mu \frac{q}{m}\int_{t_0}^{t}\int_V B(\vec{x})\ \Pi(\vec{x},t')\mathrm{d}^3x\mathrm{d}t' + \langle\phi(t_0)\rangle. \tag{5.4}$$

During the Run-1 analysis, it has been shown that the muon probability density's dependence on time in fill is a negligible effect [118], but Π could still be changing slowly over time due to beam dynamics effects. Additionally, once integration time extends over many fills, B acquires a time dependence. While it is reasonable to assume B to be constant in time over one muon's lifetime, other effects like temperature fluctuations in the magnet hardware, imperfect performance of the feedback system, and general long-term drifts cause changes in the magnetic field over longer timescales. Thus ω_a becomes [118]

$$\omega_a(t) = \frac{\mathrm{d}\langle\phi(t)\rangle_\mu}{\mathrm{d}t} = -a_\mu\frac{q}{m}\int_V B(\vec{x},t)\ \Pi(\vec{x},t)\mathrm{d}^3x. \tag{5.5}$$

Comparing this to eq. 5.1, the muon-weighted magnetic field $\tilde{B}$ is

$$\tilde{B}(t) = \int_V B(\vec{x},t)\ \Pi(\vec{x},t)\mathrm{d}^3x. \tag{5.6}$$

For the eventual determination of a_μ through eq. 3.7, both ω_a and $\tilde{\omega}_p$ will be averaged over datasets spanning several months, i.e. the quantity of interest is the muon ensemble average of the magnetic field,

$$\langle\tilde{B}\rangle_{\mathrm{dataset}} = \int_{t=0}^{T_{\mathrm{dataset}}}\tilde{B}(t)\mathrm{d}t = \int_{t=0}^{T_{\mathrm{dataset}}}\int_V B(\vec{x},t)\ \Pi(\vec{x},t)\mathrm{d}^3x\ \mathrm{d}t. \tag{5.7}$$

The muon probability density can be approximated by using the muon decay vertex distribution $M(\vec{x},t)$, as

$$\Pi(\vec{x},t) = \frac{M(\vec{x},t)}{\int_T\int_V M(\vec{x},t)\mathrm{d}^3x\mathrm{d}t}, \tag{5.8}$$

where

$$M(\vec{x},t) = \sum_{i=1}^{N_\mu} \delta(\vec{x}-\vec{x}_i, t-t_i), \tag{5.9}$$

with the decay vertex position and time of the i-th muon $(\vec{x}_i, t_i)$ and total number of muons N_μ. The normalizing integral in eq. 5.8 must be taken over the same time and volume as the averaging integral in eq. 5.7, in which the density $\Pi(\vec{x},t)$ is used, i.e.

$$\begin{aligned}\langle\tilde{B}\rangle_{\text{dataset}} &= \int_{t=0}^{T_{\text{dataset}}} \int_V B(\vec{x},t)\ \frac{M(\vec{x},t)}{\int_{t=0}^{T_{\text{dataset}}}\int_V M(\vec{x},t)\mathrm{d}^3x\mathrm{d}t}\mathrm{d}^3x\ \mathrm{d}t \\ &= \int_{t=0}^{T_{\text{dataset}}} \int_V \frac{B(\vec{x},t)\ M(\vec{x},t)}{N_{\text{dataset}}}\mathrm{d}^3x\ \mathrm{d}t,\end{aligned} \tag{5.10}$$

where N_{dataset} is the number of muons that decay inside the storage volume (V) between $t=0$ and $t=T_{\text{dataset}}$.

To avoid having to find an analytic expression for the time dependence of both the magnetic field and the muon distribution, the time integral is replaced by a sum of time bins in which time-independence of those quantities can be reasonably assumed. These time bins are called "run groups" (indicated by subscript rg), and section 5.2.1 will describe in more detail how they are constructed. Then the muon-weighted magnetic field becomes

$$\begin{aligned}\langle\tilde{B}\rangle_{\text{dataset}} &= \frac{1}{N_{\text{dataset}}}\sum_{\text{run groups}}\int_{T_{rg}}\int_V B(\vec{x},t)\ M(\vec{x},t)\mathrm{d}^3x\ \mathrm{d}t \\ &= \sum_{\text{run groups}}\frac{N_{rg}}{N_{\text{dataset}}}\int_V B_{rg}(\vec{x})\ \int_{T_{rg}}\frac{M_{rg}(\vec{x})}{N_{rg}}\ \mathrm{d}t\ \mathrm{d}^3x \\ &= \sum_{\text{run groups}}\frac{N_{rg}}{N_{\text{dataset}}}\int_V B_{rg}(\vec{x})\ \Pi_{rg}(\vec{x})\ \mathrm{d}^3x = \sum_{\text{run groups}}\frac{N_{rg}}{N_{\text{dataset}}}\tilde{B}_{rg}\end{aligned} \tag{5.11}$$

This is effectively a weighted average of the muon-weighted field per run group $\tilde{B}_{rg}$, using the number of muon decays in each run group as weights. The most accurate count of muon decays comes from the calorimeters (CTAG count (CTAGs)), and a **C**alorimeter **TAG**ged (CTAG)-weighted average also ensures the magnetic field is weighted according to which muons enter the omega-a analysis.

Similarly to the time bins, the azimuthal part of the volume integral will be split into a sum of azimuthal bins in order to utilize the 2D multipole parameterization introduced in section 4.1, eq. 4.6 [116]. The muon decay rate is uniform in azimuth, which allows for the run

group magnetic field to be written as the arithmetic mean of the azimuthal bins:

$$\begin{aligned}
\tilde{B}_{rg} &= \int_V B(\vec{x})\ \Pi(\vec{x})\mathrm{d}^3x = \int_A \int_0^{2\pi} \frac{B(x,y,\phi)\ M(x,y,\phi)}{N_{rg}}\ \mathrm{d}\phi\ \mathrm{d}x\ \mathrm{d}y \\
&= \int_A \sum_{i=1}^{N_\phi} \int_{\phi_i}^{\phi_{i+1}} \frac{B(x,y,\phi)\ M(x,y,\phi)}{N_\phi N_{rg,\Delta\phi}}\ \mathrm{d}\phi\ \mathrm{d}x\ \mathrm{d}y \\
&\approx \frac{1}{N_\phi} \sum_{i=1}^{N_\phi} \int_A B_{\Delta\phi_i}(x,y) \int_{\phi_i}^{\phi_{i+1}} \frac{M(x,y,\phi)}{N_{rg,\Delta\phi}}\ \mathrm{d}\phi\ \mathrm{d}x\ \mathrm{d}y \\
&= \frac{1}{N_\phi} \sum_{i=1}^{N_\phi} \int_A B_{\Delta\phi_i}(x,y)\ \Pi_{\Delta\phi_i}(x,y)\ \mathrm{d}x\ \mathrm{d}y. \qquad (5.12)
\end{aligned}$$

Note that $B_{\Delta\phi_i}(x,y) = \langle B(x,y)\rangle_{\Delta\phi_i}$ is the average across the azimuthal bin of the 2D magnetic field (see eq. 4.6), and $\Pi_{\Delta\phi_i}(x,y)$ is the normalized muon distribution summed over the azimuthal bin. A convenient choice for the size of $\Delta\phi$ is 5°, which coincides with the average azimuthal coverage of the fixed probes. Using the 2D multipole parameterization, one can now define "beam multipole projections", i.e. describe the muon probability density in terms of factors multiplying the multipole strengths [116]:

$$\begin{aligned}
\tilde{B}_{rg} &= \frac{1}{N_\phi} \sum_{i=1}^{N_\phi} \int_A B_{\Delta\phi_i}(x,y)\ \Pi_{\Delta\phi_i}(x,y)\ \mathrm{d}x\ \mathrm{d}y \\
&= \frac{1}{N_\phi} \sum_{i=1}^{N_\phi} \int_A \sum_{j=1}^{13} m_j(\phi_i) f_j(x,y)\ \Pi_{\Delta\phi_i}(x,y)\ \mathrm{d}x\ \mathrm{d}y \\
&= \frac{1}{N_\phi} \sum_{i=1}^{N_\phi} \sum_{j=1}^{13} m_j(\phi_i) \int_A f_j(x,y)\ \Pi_{\Delta\phi_i}(x,y)\ \mathrm{d}x\ \mathrm{d}y \\
&\equiv \frac{1}{N_\phi} \sum_{i=1}^{N_\phi} \sum_{j=1}^{13} m_j(\phi_i)\ k_j(\phi_i), \qquad (5.13)
\end{aligned}$$

with

$$k_j(\Delta\phi_i) = \int_A f_j(x,y)\ \Pi_{\Delta\phi_i}(x,y)\ \mathrm{d}x\ \mathrm{d}y \tag{5.14}$$

$$= \int_A f_j(x,y)\ \frac{M_{\Delta\phi_i}(x,y)}{\int_A M_{\Delta\phi_i}(x,y)\ \mathrm{d}x\ \mathrm{d}y}\ \mathrm{d}x\ \mathrm{d}y$$

$$= \int_A f_j(x,y)\ \frac{1}{N_{\Delta\phi_i}} \sum_{n=1}^{N_{\Delta\phi_i}} \delta(x - x_n, y - y_n)\ \mathrm{d}x\ \mathrm{d}y$$

$$= \frac{1}{N_{\Delta\phi_i}} \sum_{n=1}^{N_{\Delta\phi_i}} \int_A f_j(x,y)\ \delta(x - x_n, y - y_n)\ \mathrm{d}x\ \mathrm{d}y$$

$$= \frac{1}{N_{\Delta\phi_i}} \sum_{n=1}^{N_{\Delta\phi_i}} f_j(x_n, y_n) = \langle f_j(x,y) \rangle_{\Delta\phi_i}, \tag{5.15}$$

where ϕ_i can be interpreted as the center of the azimuthal bin $\Delta\phi_i$, but $m_j(\phi_i)$ is in fact the azimuthally averaged magentic field moment in the bin, and $k_j(\phi_i)$ is determined from all muon decays within the bin. An alternative naming convention for the beam multipole projections is to use I_n for "normal" beam projections and J_n for "skew" projections, analogous to the A_n and B_n parameters of the magnetic field. Table 5.1 shows a side-by-side comparison of the magnetic multipole moments and the corresponding beam multipole projections.

Ideally, the muon distribution $M(x,y)$ would be determined from data in each of the azimuthal bins $\Delta\phi_i$. In practice however, the calorimeters aren't able to provide granular enough position data to reliably reconstruct the beam shape. The tracking detectors can provide the muon decay distribution, but only at two positions along the ring. Therefore, the beam distributions at the tracker locations have to be extrapolated to other points in azimuth using beam dynamics information. The tracking detectors are able to reconstruct only a few percent tracks relative to the number of CTAGs, thus $\Pi_{\Delta\phi_i}(x,y)$ in eq. 5.14 must be normalized with the number of tracks rather than CTAGs.

Fig. 5.2 shows an overview of the steps of the muon weighting process, including necessary inputs from other parts of the experiment. Briefly summarized, the procedure is as follows:

1. Define run groups based on tracker data, taking into account both the number of reconstructed tracker hits to minimize statistical uncertainty, as well as the duration of the run groups to be able to approximate the magnetic field and muon distribution

n	moment	field	$k_i, I_n/J_n$	beam projection
0	Dip	m_1	$k_1 = I_0$	1.0
1	NQ	$m_2\bar{x}$	$k_2 = I_1$	$\langle\bar{x}\rangle$
1	SQ	$m_3\bar{y}$	$k_3 = J_1$	$\langle\bar{y}\rangle$
2	N6	$m_5(\bar{x}^2 + \bar{y}^2)$	$k_5 = I_2$	$\langle\bar{x}^2\rangle + \langle\bar{y}^2\rangle$
2	S6	$2m_4\bar{x}\bar{y}$	$k_4 = J_2$	$2\langle\bar{x}\bar{y}\rangle$
3	N8	$m_7(\bar{x}^3 - 3\bar{x}\bar{y}^2)$	$k_7 = I_3$	$\langle\bar{x}^3\rangle - 3\langle\bar{x}\bar{y}^2\rangle$
3	S8	$m_6(3\bar{x}^2\bar{y} - \bar{y}^3)$	$k_6 = J_3$	$3\langle\bar{x}^2\bar{y}\rangle - \langle\bar{y}^3\rangle$
4	N10	$m_8(\bar{x}^4 - 6\bar{x}^2\bar{y}^2 + \bar{y}^4)$	$k_8 = I_4$	$\langle\bar{x}^4\rangle - 6\langle\bar{x}^2\bar{y}^2\rangle + \langle\bar{y}^4\rangle$
4	S10	$4m_9(\bar{x}^3\bar{y} - \bar{x}\bar{y}^3)$	$k_9 = J_4$	$4\langle\bar{x}^3\bar{y}\rangle - 4\langle\bar{x}\bar{y}^3\rangle$
5	N12	$m_{10}(\bar{x}^5 - 10\bar{x}^3\bar{y}^2 + 5\bar{x}\bar{y}^4)$	$k_{10} = I_5$	$\langle\bar{x}^5\rangle - 10\langle\bar{x}^3\bar{y}^2\rangle + 5\langle\bar{x}\bar{y}^4\rangle$
5	S12	$m_{11}(5\bar{x}^4\bar{y} - 10\bar{x}^2\bar{y}^3 + \bar{y}^5)$	$k_{11} = J_5$	$5\langle\bar{x}^4\bar{y}\rangle - 10\langle\bar{x}^2\bar{y}^3\rangle + \langle\bar{y}^5\rangle$
6	N14	$m_{12}(\bar{x}^6 - 15\bar{x}^4\bar{y}^2 + 15\bar{x}^2\bar{y}^4 - \bar{y}^6)$	$k_{12} = I_6$	$\langle\bar{x}^6\rangle - 15\langle\bar{x}^4\bar{y}^2\rangle$ $+15\langle\bar{x}^2\bar{y}^4\rangle - \langle\bar{y}^6\rangle$
6	S14	$m_{13}(6\bar{x}^5\bar{y} - 20\bar{x}^3\bar{y}^3 + 6\bar{x}\bar{y}^5)$	$k_{13} = J_6$	$6\langle\bar{x}^5\bar{y}\rangle - 20\langle\bar{x}^3\bar{y}^3\rangle + 6\langle\bar{x}\bar{y}^5\rangle$

Table 5.1: Comparison of field multipole moments and beam multipole projections, in cartesian coordinates. $\bar{x} = \frac{x}{r_0}$ and $\bar{y} = \frac{y}{r_0}$ are the normalized radial and vertical distance from the center of the storage volume, with $r_0 = 45$ mm. The dipole projection $k_1 = I_0 = 1.0$ by definition.

as time-independent (see section 5.2.1).

2. Extract muon decay vertex distributions at the azimuthal positions of the tracker stations for each run group (see section 5.2.3).

3. Produce virtual trolley runs for the run group time intervals T_{rg} (see section 5.1).

4. Extrapolate the beam distributions at the tracker locations to azimuthal bins along the entire storage ring, using beam dynamics information from other parts of the experiment and simulation (see section 5.3).

5. Calculate beam multipole projections per azimuthal bin and run group, as discussed above.

6. Calculate the weighted field per run group by multiplying the field multipole strengths

with their corresponding beam projections in each azimuthal bin, and taking the azimuthal mean.

7. Calculate the total weighted field for a desired dataset as the CTAG-weighted average of all run groups in the dataset.

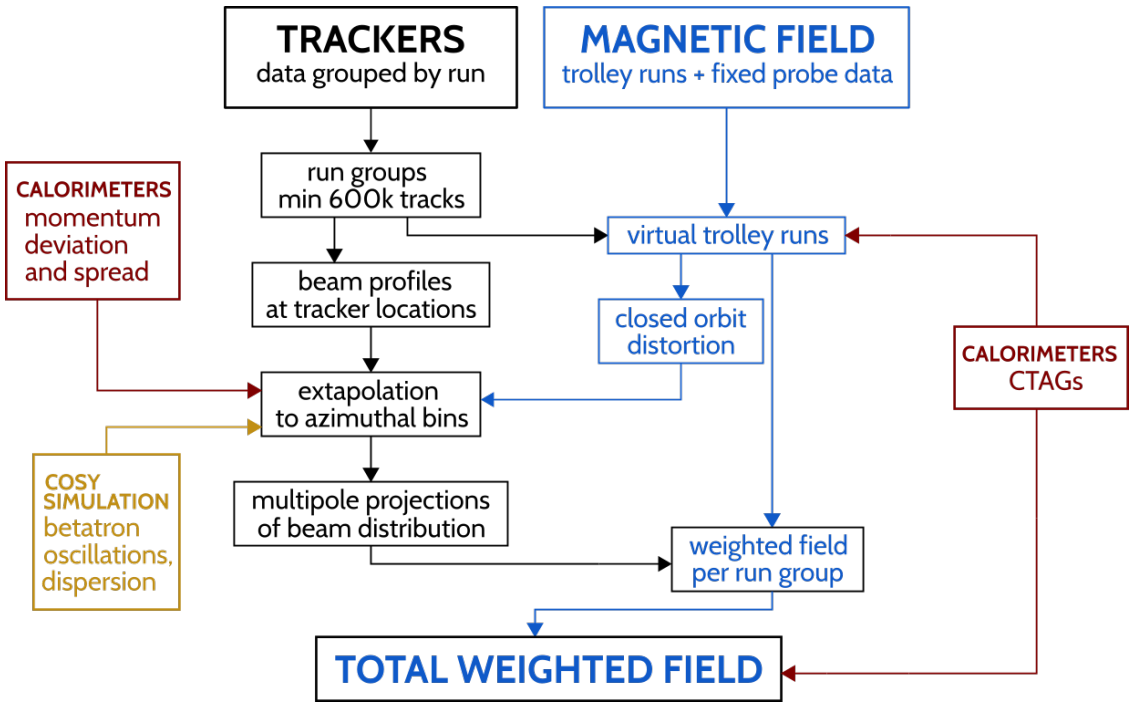


Figure 5.2: Flow chart of muon weighting steps and inputs, reproduced from [69]. Details can be found in sections 5.1 (magnetic field related steps), and 5.2-5.3 (tracker/beam related steps).

5.1 Magnetic field binning

In section 3.2.3 the concept of virtual trolley runs (VTRs) was introduced. VTRs are spatial field maps constructed for times in between physical trolley runs, using fixed probe data to track the first few multipole moments across time (cf. also section 4.1). During each trolley run, the trolley and fixed probe moments are synchronized by determining the calibration offset $\vec{c}_s$ of each fixed probe station s according to [140, 10]

$$\vec{c}_s = \vec{m}_{tr}(\phi_s, t_s) - \mathbf{J}_s \vec{m}_{fp,s}(t_s). \tag{5.16}$$

The vector $\vec{m}_{tr} = (m_1, ..., m_{13})^T$ is constructed from the multipole strengths determined from the trolley measurements, specifically at the location ϕ_s of a fixed probe station s.

Similarly, the vector $\vec{m}_{fp,s}$ contains the multipole strengths from the fixed probe station s (up to 5 elements for 6-probe stations), measured at time t_s when the trolley passes ϕ_s. $\mathbf{J}_s$ is a transformation matrix that correlates fixed probe and trolley moment bases, and depends on station geometry. $J_{s,ij} = \frac{\partial(m_{tr})_i}{\partial(m_{fp,s})_j} = 0$ for $i \geq \max(j)$, i.e. moments that aren't tracked by a fixed probe aren't calibrated. Naively one would expect the remainder of $\mathbf{J}_s$ to be the identity matrix, however due to the geometric arrangement of the fixed probe stations, there is some aliasing of higher order moments into lower orders, in particular the normal sextupole (N6) will create a false dipole signal. Thus $\mathbf{J}_s$ contains some constant, station-dependent off-diagonal elements.

Due to drifts in the untracked higher order moments (which alias into the tracked moments) [140], the synchronization constant $\vec{c}_s$ varies with each trolley run (which leads to the sync offset mentioned in section 3.2.3). The majority of trolley pairs can in principle be interpolated from either end, and thus a linear interpolation in time is used, such that a "virtual" (i.e. interpolated) trolley measurement at fixed probe station s at an arbitrary time t (between trolley run calibration times t_1 and t_2 at that station) can be written as

$$\vec{m}_{vtr}(\phi_s, t) = \mathbf{J}_s \vec{m}_{fp}(t) + \frac{|t - t_2|}{t_2 - t_1}\vec{c}_{s,1} + \frac{|t - t_1|}{t_2 - t_1}\vec{c}_{s,2}. \tag{5.17}$$

There are situations when data between two trolley runs can only be interpolated from one side, for example if there are problems with data taking during one of the trolley runs, or the surface coil configuration is changed between two trolley runs [140]. In those cases, only the synchronization constant from one trolley run is used, and a larger systematic uncertainty is assigned.

For a full virtual trolley run, more spatial data points than the 72 fixed probe locations are needed. For most analysis purposes, a simple linear interpolation in space is used [50, 60]. A virtual trolley measurement at time t and azimuthal location ϕ between fixed probe stations $s1$ at ϕ_1 and $s2$ at ϕ_2 is then

$$\vec{m}_{vtr}(\phi, t) = \frac{|\phi - \phi_2|}{\phi_2 - \phi_1}\vec{m}_{vtr}(\phi_1, t) + \frac{|\phi - \phi_1|}{\phi_2 - \phi_1}\vec{m}_{vtr}(\phi_2, t). \tag{5.18}$$

Since the time interpolation step as well as the transformation matrix $\mathbf{J}_s$ are linear, virtual trolley runs for time intervals can be obtained by simply averaging the respective fixed probe data beforehand.

As part of the muon weighting procedure, a VTR is produced for every run group, which typically span one to a few hours. DQC-passing fixed probe data is combined into

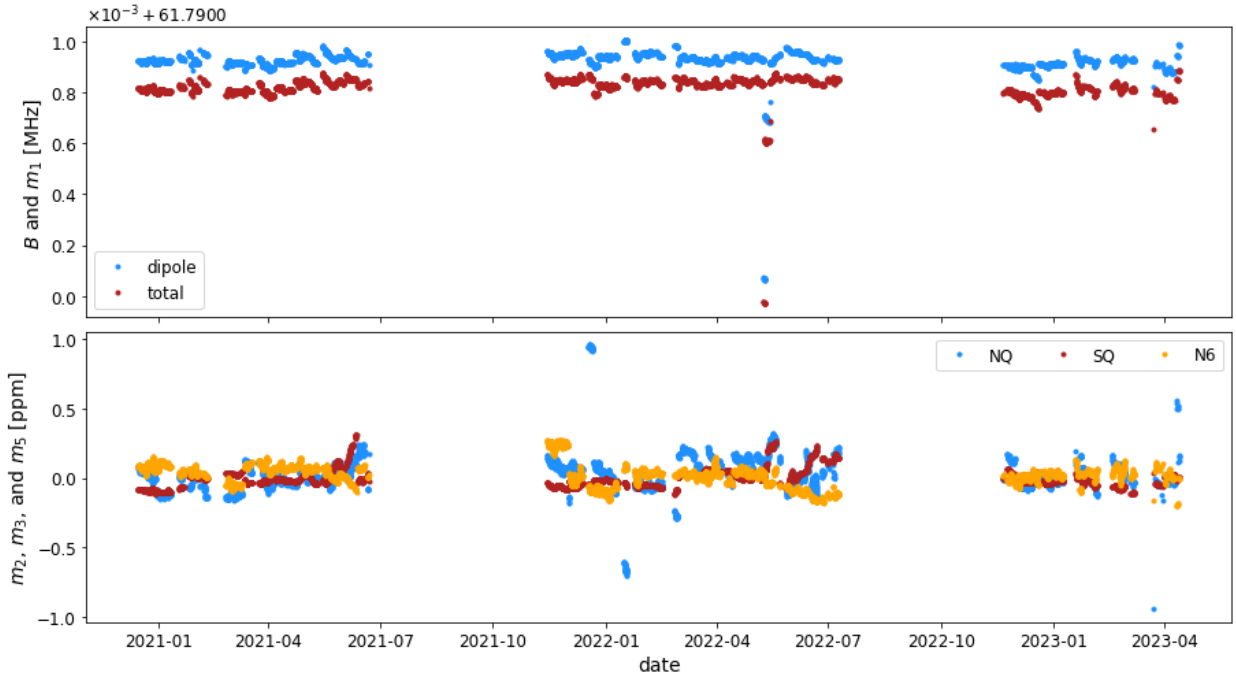


Figure 5.3: Unweighted field vs. time. Overview of unweighted field (azimuthally averaged) over time: Total field and dipole component on the top, and the next three most important multipole contributions. As will be shown in section 5.4, the beam projections onto the normal quadrupole and sextupole as well as the skew quadrupole are about an order of magnitude larger than the remaining multipole components. The three clusters of data correspond to the three run periods (4, 5, and 6).

a CTAG-weighted average across the duration of a run group [117], then interpolated to a VTR with 72×10^4 equally spaced azimuthal points [69] according to eq.s 5.16-5.18. At this point, the closed orbit distortion is calculated, which will be explained in section 5.3. Additionally, a quadratic and a cubic spatial interpolation is also calculated. Each of the three different interpolated VTRs is then averaged into 5° azimuthal bins, and the arithmetic mean of the three interpolations is taken as the final (binned) VTR, which will be combined with the beam projection weights according to eq. 5.13.

Fig. 5.3 shows the azimuthal average of each VTR (i.e. the mean of the 72 azimuthal bins) in Run-4/5/6 over time, for the total field strength and select multipoles. A summary of time-averaged field values for each RF-dataset will be shown in section 5.4 along with the results of the muon weighting.

5.2 Beam data binning

The muon beam distribution is typically more stable than the magnetic field, especially during Runs 4-6 when run conditions were well established. Thus the criteria of approximately "time-independent" run groups is in principle limited by the variation of the magnetic field. However, a sufficient number of reconstructed muon decay vertices from the tracking detectors ("tracks") have to be accumulated in order to limit the statistical uncertainty on the muon weighted field.

5.2.1 Run group formation

As stated previously, it is crucial to weigh the magnetic field by the specific muons that enter the ω_a analysis. Aside from DQC, an in-fill time filter also needs to be applied, as only positrons with in-fill decay times between $t_{\text{start}} \approx 30$ μs and $t_{\text{end}} \approx 650$ μs enter in the ω_a analysis [98]. The basic strategy to construct run groups is then to define time intervals in which the number of muon decay vertices that fulfill these requirements is large enough to have negligible statistical uncertainty, while also keeping the time intervals as short as possible.

In Run-1, it has been shown that a minimum of 6×10^5 DQC-passing tracks per run group (from both tracker stations combined) results in less than 1 ppb statistical error on the weighted field [118]. Due to the large data volume in Runs-4/5/6 and the computational demands of iterating over the tracker data, run groups are also required (except for a handful of cases detailed below) to contain entire DAQ runs[1]. Additionally, run groups must be contained within trolley pairs (to match a VTR to the run group) and omega-a datasets [117]. Finally, large time gaps between DQC-passing tracks should ideally be outside run group bounds in order to minimize field variation over time.
The Run-2/3 muon weighting analysis employed a straightforward algorithm to define run groups, which can be summarized as follows:

1. Define fixed boundaries including omega-a dataset bounds, trolley pair bounds, and time gaps > 4h.

2. Within each fixed boundary, list the DAQ runs containing DQC-passing tracks from earliest to latest.

[1]Each run period contains about 400 times as many DAQ subruns as DAQ runs. In Run-4, tracker data exists for $2,366$ runs, which contain a total of $951,480$ subruns; in Run-5, for $3,169$ runs with $1,337,034$ subruns; and in Run-6, for $1,519$ runs with $600,858$ subruns.

3. Going run by run, starting with the earliest run, accumulate runs into a run group until required minimum number of tracks (6×10^5) is achieved, then start a new run group.

4. If the last run group before a fixed boundary doesn't contain enough tracks, combine it with the previous one. (This step does not check if the previous run group lies across a fixed boundary.)

While this strategy did achieve all the listed criteria for Run-2/3, it still required manual inspection of and occasional changes to the run group definitions afterwards, since there would still be large time gaps within run groups that could be avoided. Additionally, for Run-6, there are six trolley pairs that contain less than the required number of tracks across the entire trolley pair. In the above algorithm, these trolley pairs will be grouped with the preceding trolley pair due to the last step.

A new algorithm was developed for the Run-4/5/6 muon weighting analysis, which aims to eliminate the need for manual editing of run groups and thus improve consistency and reproducibility, as well as optimize run group distribution and duration. The new strategy is as follows [69]:

1. Trolley pair and omega-a dataset bounds are enforced by creating run groups independently for each trolley pair. Additional boundaries can be defined beforehand as needed and will also be adhered to (see table 5.2 in section 5.2.2 for a list of additional hard-coded boundaries).

2. Within each trolley pair, runs are separated into "above-target" runs (i.e. runs already containing the required minimum tracks, that can form a run group by themselves) and "sub-target" runs (runs that need to be combined with other runs into run groups).

3. Any above-target runs that aren't adjacent to sub-target runs are defined as run groups.

4. The remaining runs are split into sets of consecutive runs. Each set begins and ends with either an above-target run (that is adjacent to a sub-target run on one side), or a border. Above-target runs that are flanked on both sides by sub-target runs are contained within these sets.

5. A recursive function is called on each set, that evaluates all possible permutations of forming valid run groups (where each run group contains the required minimum number of tracks). This function performs a modified minimax algorithm, and minimizes the

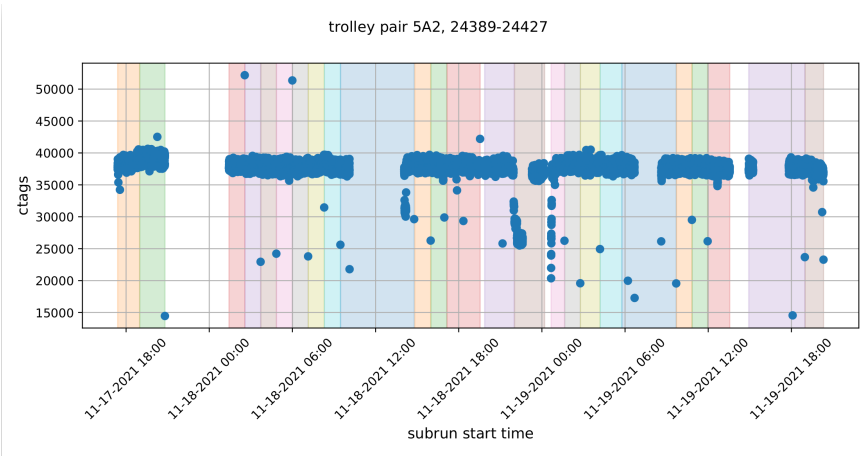
largest duration(s) of the run groups formed (i.e. it minimizes the total time of the longest run group, if there are multiple configurations with the same longest duration it minimizes the time of the second longest run group, and so on).

The recursive function's objective is to minimize the overall duration of all run groups while prioritizing the longest ones. This means it avoids large time gaps automatically wherever possible, since including a time gap in a run group significantly increases that run group's length. Additionally, this leads to more evenly distributed run groups (i.e. reduces variance in the run group duration), because decreasing the length of one run group must increase the length of another.

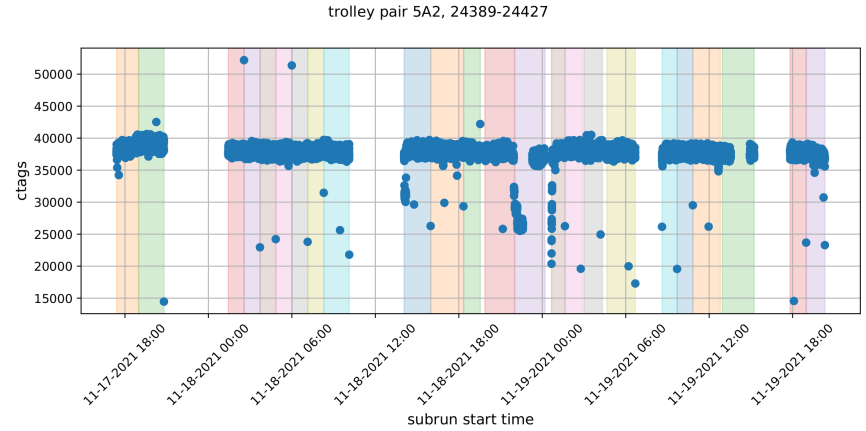
Two specific examples of how run group formation is improved with the new algorithm are provided in figs. 5.4 and 5.5. Fig. 5.4 shows that with the new algorithm, multiple time gaps that were previously included are now excluded. Fig. 5.5 illustrates that where the Run-2/3 algorithm created a very long run group from a number of runs with few tracks, the new Run-4/5/6 strategy is able to create two shorter ones by grouping some of the sub-target runs with the previous above-target run.
Histograms of the run group durations for the two algorithms per run period (4/5/6) are shown in fig. 5.6. A clear decrease of run groups longer than $10,000$s can be observed for Runs 4 and 5, and the standard deviation of the run group length $\sigma_{T_{rg}}$ reduces from 2375s to 2065s for Run-4, and from 2268s to 1691s for Run-5. Due to some trolley pairs in Run-6 which, as a whole, don't contain the required minimum number of tracks, the Run-2/3 algorithm does not produce a full list of valid run groups for Run-6 (creating run groups across trolley pair boundaries instead). The Run-4/5/6 algorithm is designed to prioritize adhering to the trolley pair boundaries, as this is necessary for VTRs. Thus there is an increased number of run groups longer than $30,000$s in Run-6, leading to only a slight decrease of $\sigma_{T_{rg}}$ from 3858s to 3785s.

5.2.2 Non-standard run groups

As mentioned above, a few issues were identified that required either forced run group splits (i.e. hard-coded run group boundaries) or subrun splits of a detector DAQ run. These cases are listed in table 5.2. Large field jumps should be excluded from run groups as they clearly violate the assumption of time-independence of the magnetic field. Similarly, a surface coil set point change requires a run group split to ensure reliability of the corresponding VTR. Lastly, in one instance during Run-5, the border between two consecutive trolley runs



(a) Run-2/3 unedited run group definitions for trolley pair 5A2



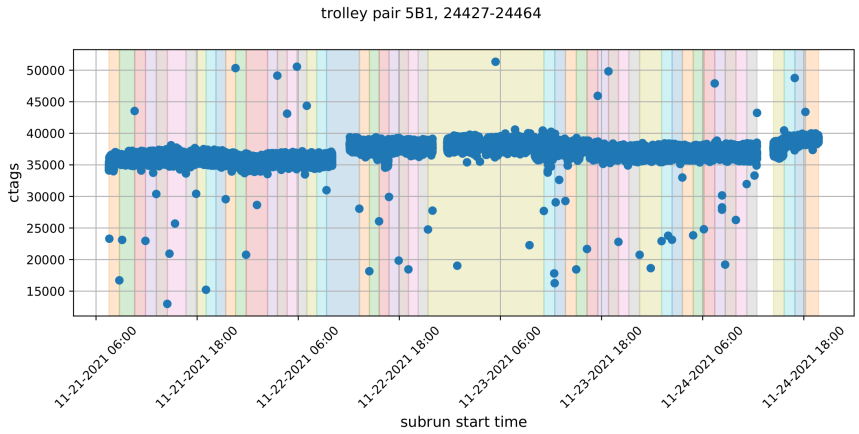
(b) Run-4/5/6 (unedited) run group definitions for trolley pair 5A2

Figure 5.4: Run group algorithm comparison for trolley pair 5A2. Comparison of run groups defined by Run-2/3 algorithm (top) and Run-4/5/6 algorithm (bottom) for trolley pair 5A2, reproduced from [69]. The new algorithm excludes three time gaps (one of which is about 3h long) automatically that were previously included, and includes one shorter time gap to allow for better grouping of a sub-target run.

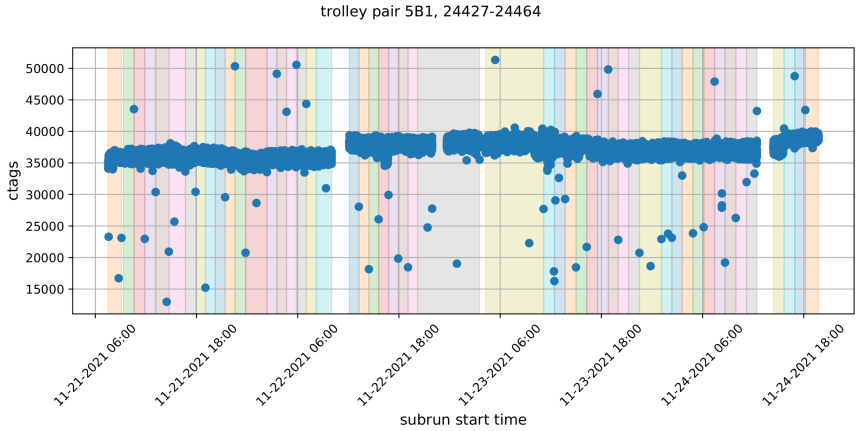
intersects with a detector DAQ run[2], which has to be split by subrun to observe the trolley pair bounds.

In addition to these non-standard run group boundaries, the six Run-6 trolley pairs that contain less than the minimum number of tracks also require special attention. Since the number of required tracks is chosen such that the statistical uncertainty of the total weighted

[2]The two trolley runs are not defined between two physical trolley measurements as is otherwise the case. The border between them corresponds to a change in fixed probe FID parameters, which breaks the synchronization between fixed probes and trolley.



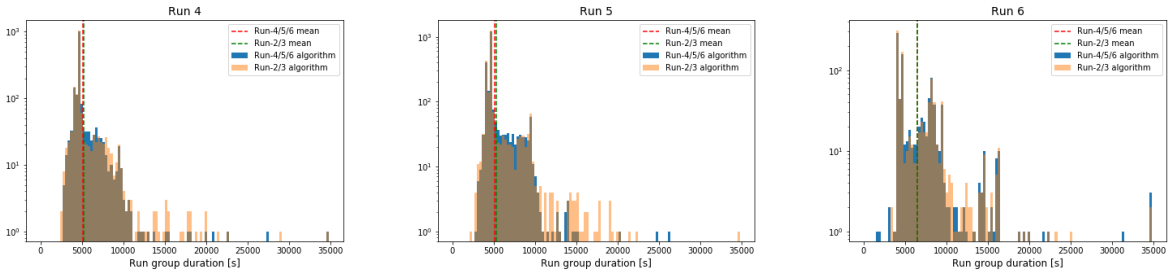
(a) Run-2/3 unedited run group definitions for trolley pair 5B1



(b) Run-4/5/6 (unedited) run group definitions for trolley pair 5B1

Figure 5.5: Run group algorithm comparison for trolley pair 5B1. Comparison of run groups defined by Run-2/3 algorithm (top) and Run-4/5/6 algorithm (bottom) for trolley pair 5B1, reproduced from [69]. Previously, a number of subsequent sub-target runs had been added up to one long run group. The new Run-4/5/6 algorithm creates two medium length run groups instead, with about half of the sub-target runs joined to the preceding above-target run.

field will be negligible (< 1 ppb), evaluating the statistical uncertainty of those six sub-target run groups will determine if further action is needed. An upper limit estimate (reproduced in detail in appendix D.1) shows that while each individual run group has a statistical uncertainty on the order of a few ppb, their limited size also means they each contribute less than 10^{-3} ppb to the total weighted field, and will not increase the overall statistical uncertainty to a significant level. Thus the six trolley pairs, listed with the results in table 5.3, are kept as-is, each forming one run group with $< 6 \times 10^5$ tracks.



(a) Run-4 (b) Run-5 (c) Run-6

Figure 5.6: Comparison of run group durations, for the new Run-4/5/6 algorithm vs the Run-2/3 algorithm, reproduced from [69].

trolley pair	**intervention**	**reason**
4G1	hard-coded run group boundary	large field jump during field run
4J1	hard-coded run group boundary	large field jump during field run
5Q3	subrun split of omega-a run	trolley pair end during omega-a run,
5Q4	subrun split of omega-a run	trolley pair start during omega-a run,
5Q5	subrun split of omega-a run and hard-coded run group boundary	set-point change
6F4	hard-coded run group boundary	large field jump during field run

Table 5.2: List of required hard-coded run group splits during Run-4/5/6

5.2.3 Beam distribution reconstruction

The tracker stations are identified by the number of the calorimeter they are placed in front of, 12 (at 180° from the inflector exit) and 18 (at 270°) respectively. The field coordinate system references the 12 o'clock position of the storage ring as 0° azimuth, which is 18.35° upstream of the inflector exit[3]. Their mean azimuthal sensitivity is a few degrees upstream of their location, thus the azimuthal locations of the reconstructed beam distributions in the field coordinate system are $\theta_{st} \approx 192.3°, 282.4°$ for station 12 and 18 respectively [118]. Both stations have a narrow azimuthal sensitivity root mean square (RMS) width (4.9° and 4.8°) and can reasonably be treated as localized at ϕ_{12} and ϕ_{18} [6]. For each run group, the transverse muon beam distributions $M_{st}(x, y)$ at station $st = 12, 18$ (also called

[3]A summary of the different coordinate systems used in this analysis and how they relate to each other can be found in appendix A.

Trolley pair	total tracks	σ_{stat} [ppb]	CTAG fraction	uncertainty contribution [ppt]
6F6	193,403	2.48	2.57×10^{-4}	≤ 0.638
6G3	414,785	1.70	5.57×10^{-4}	≤ 0.949
6J3	275,790	2.08	3.74×10^{-4}	≤ 0.778
6J6	436,037	1.71	5.78×10^{-4}	≤ 0.988
6J7	147,507	2.92	1.95×10^{-4}	≤ 0.570
6J9	278,176	2.12	3.74×10^{-4}	≤ 0.789

Table 5.3: Sub-target trolley pairs: Trolley pairs in Run-6 that do not contain the desired minimum number of tracks to form a run group, with total number of DQC-passing tracks, statistical uncertainty estimate, CTAG fraction of dataset xyRF6, and upper bound to their contribution to the total statistical uncertainty of dataset xyRF6 (calculated as the product of CTAG fraction and statistical uncertainty $\sigma_{\tilde{B}_{rg}}$).

tracker profiles) are accumulated as histograms with 0.5 mm wide bins in x and y from each station's DQC-passing reconstructed muon decay vertex position[4].

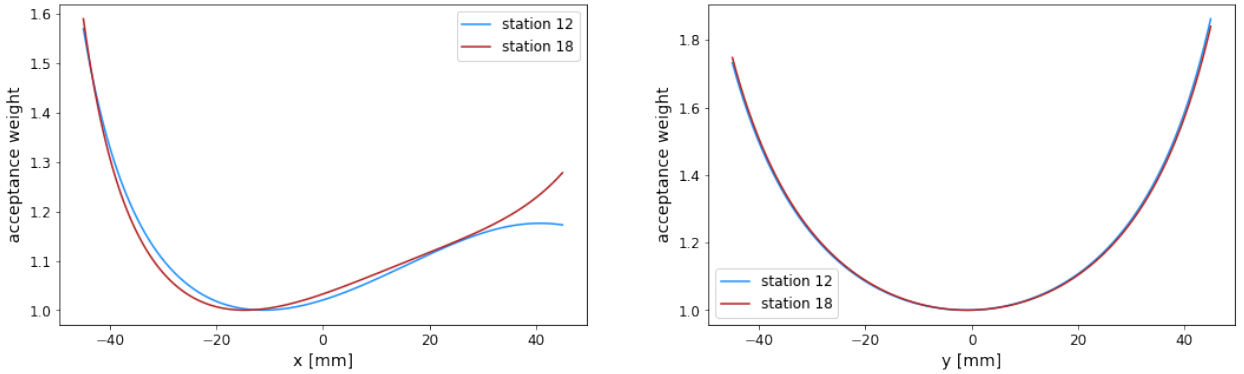
Corrections for tracker resolution and acceptance have to be applied to obtain beam distributions that are representative of the true muon distribution inside the storage volume. The resolution correction, which slightly narrows the tracker profiles [118], is performed using an iterative unfolding method [62] in x and y independently to account for the vertex positions that can't be reconstructed the tracking detectors[5]. The acceptance correction is based on spatial acceptance functions determined from simulation (shown in fig. 5.7) and broadens the tracker profiles to account for the lower acceptance towards the edges of the detector. The radial acceptance function is asymmetric and therefore also changes the radial mean of the distributions [118]. Fig 5.8 shows an example of a beam histogram before and after applying resolution and acceptance corrections.

5.3 Beam profile extrapolation

The tracker profiles are a representation of the muon beam distribution at ϕ_{12} and ϕ_{18}. To extrapolate to other azimuthal positions (then called beam profiles), knowledge about beam

[4]Details on the reconstruction algorithm can be found in chapter 6.

[5]Due to the geometric arrangement of the straw tubes in the tracking stations and the limitations of the track reconstruction algorithm, not all vertex positions can be recovered accurately. Therefore the bin contents in the raw reconstructed beam profiles have to be "smeared out" to fill the intermediate spaces.



(a) Radial acceptance function

(b) Vertical acceptance function

Figure 5.7: Radial and vertical acceptance functions for beam profile reconstruction. The acceptance of each tracker station is simulated independently and parameterized as a 4th order polynomial for the radial acceptance, and a second order polynomial for the vertical acceptance. Hits further from the center generally are weighted more strongly to account for the reduces likelihood of detection.

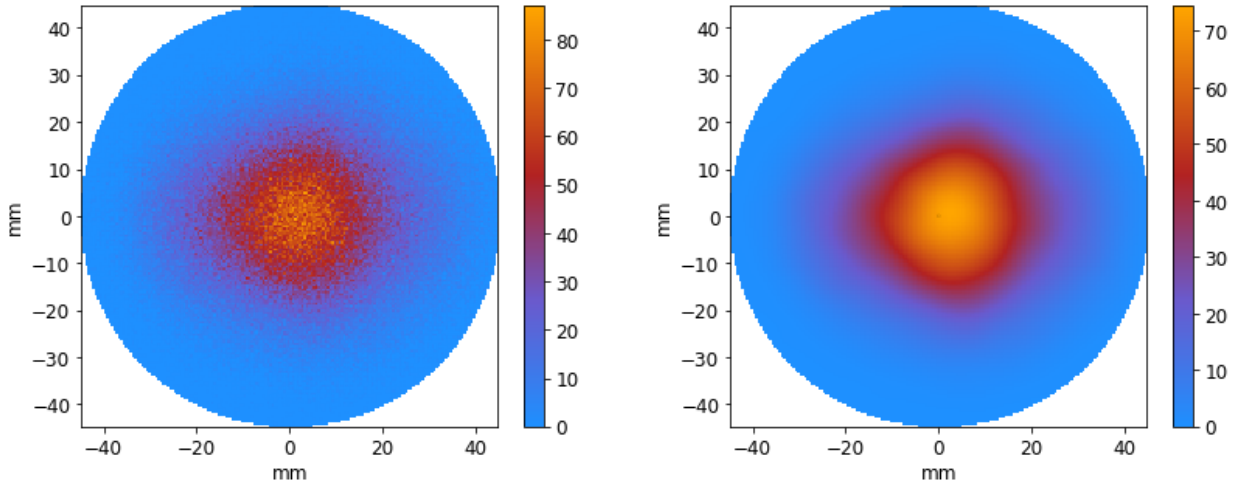
dynamics effects that move the beam coherently is necessary. The dominant beam dynamics effects include [116]:

- **Dispersion** - The magnetic field is designed to keep muons with momentum p_0 on the central orbit of the storage ring. Since the degree of deflection a charged particle experiences in a magnetic field is momentum-dependent, muons with relative momentum deviation $\delta = \frac{p-p_0}{p_0}$ will have stable orbits with different radii. The ensemble of muons in the beam with mean momentum offset $\langle\delta\rangle$ and momentum offset width σ_δ "disperse" from the central orbit. Additionally, the non-continuous nature of the ESQs introduces azimuthal variations to the radial dispersion function $D_x(\phi)$.

- **Closed orbit distortion** (COD) - Azimuthal variations in the vertical field lead to small changes of the curvature of a muon's path, a distortion of the otherwise circular closed orbit[6], affecting the beam's mean position as a function of azimuth.

- **Betatron oscillations** - The muons enter the storage ring radially off-center through the inflector exit. The kickers provide a short magnetic "kick" to center the beam on the central orbit, but any muon that has a transverse momentum component after the kick (i.e. it's momentum is not perfectly tangential to the central orbit) will experience betatron motion. Betatron motion is an oscillation around the muon's equilibrium

[6]Azimuthal variations of the radial field similarly create a vertical distortion to the closed orbit. Due to the experiment's inability to measure the radial dipole field directly, the vertical COD is only considered as an uncertainty, see section 5.5.



(a) Raw station 18 histogram

(b) Corrected station 18 histogram

Figure 5.8: Illustration of beam reconstruction corrections. Example of a station 18 raw (uncorrected) tracker histogram, and the same data after applying resolution and acceptance corrections. The color gradient indicates number of hits in a 0.5×0.5mm bin. The change from "grainy" to "smooth" is due to the resolution correction, while the acceptance correction causes a subtle shift in the beam distribution's shape.

radius due to the ESQs vertically focusing force and the inherent weak horizontal focusing of the storage ring. The frequency of betatron oscillations is typically not a multiple of the cyclotron frequency, which is why the betatron motion is a function of total path length s rather than azimuthal position ϕ. The *betatron function* $\beta(\phi)$ however, a measure of the focusing strength at ϕ, is a function of the ring structure and thus periodic.

A single muon's trajectory can be modeled in transverse coordinates x and y, as a function of total path length s and azimuthal angle ϕ, as [118]

$$x(s) = \delta D_x(\phi) + \Delta x_{COD}(\phi) + x_\beta(s), \tag{5.19}$$

$$y(s) = \Delta y_{COD}(\phi) + y_\beta(s), \tag{5.20}$$

where the betatron oscillations in x and y can be describes by [118]

$$x_\beta(s) = \sqrt{\epsilon_x \beta_x(\phi)}\cos(\nu_x s/R - \phi_x), \tag{5.21}$$

$$y_\beta(s) = \sqrt{\epsilon_y \beta_y(\phi)}\cos(\nu_y s/R - \phi_y). \tag{5.22}$$

ϕ_x and ϕ_y are determined from initial conditions. $\nu_x = \omega_x/\omega_c = \sqrt{1-n}$ and $\nu_y = \omega_y/\omega_c = \sqrt{n}$ are the horizontal and vertical tunes, which represent the number of betatron oscillations per revolution, and can be calculated in a continuous quad approximation from the effective field index $n = \frac{R_0}{vB_0}\frac{\partial E_y}{\partial y}$ [9]. The field index is mainly influenced by the voltage of the ESQs, and in Runs 4-6, $n \approx 0.108$. The emittances ϵ_x and ϵ_y represent the phase space area of the beam.

For an ensemble of muons who sample the ring many times during their lifetime, the mean transverse positions are then

$$\langle x \rangle(\phi) = \langle \delta \rangle D_x(\phi) + \Delta x_{COD}(\phi), \tag{5.23}$$

$$\langle y \rangle(\phi) = \Delta y_{COD}(\phi), \tag{5.24}$$

and the beam widths are

$$\sigma_x(\phi) = \sqrt{\epsilon_x \beta_x(\phi) + D_x^2(\phi)\sigma_\delta^2}, \tag{5.25}$$

$$\sigma_y(\phi) = \sqrt{\epsilon_y \beta_y(\phi)}. \tag{5.26}$$

To determine the beam mean and width at an arbitrary azimuthal position, inputs from other parts of the experiment are required: the COD can be determined from the virtual trolley runs; the average momentum deviation and spread are calculated as part of the CERNext E-field correction [45]; and the beam functions (radial and vertical betatron, radial dispersion) come from a simulation using a COSY-based model of the Muon g-2 ring fields [138]. Each of these inputs will be described in more detail in the following sections.

Practically, beam profiles at ϕ_i are constructed by shifting and scaling each of the tracker profiles using eqs. 5.23-5.26, such that [69]

$$\langle x \rangle(\phi) = \Delta x_{COD}(\phi) + \langle \delta \rangle D_x(\phi) + \Big(\langle x \rangle(\phi_{st}) - \Delta x_{COD}(\phi_{st}) - \langle \delta \rangle D_x(\phi_{st}) \Big), \tag{5.27}$$

$$\langle y \rangle(\phi) = \langle y \rangle(\phi_{st}), \tag{5.28}$$

$$x(\phi) - \langle x \rangle(\phi) = \frac{\sigma_x(\phi)}{\sigma_x(\phi_{st})} \Big(x(\phi_{st}) - \langle x \rangle(\phi_{st}) \Big), \tag{5.29}$$

$$y(\phi) - \langle y \rangle(\phi) = \frac{\sigma_y(\phi)}{\sigma_y(\phi_{st})} \Big(y(\phi_{st}) - \langle y \rangle(\phi_{st}) \Big). \tag{5.30}$$

The vertical COD is assumed to be zero in the absence of reliable radial magnetic field information, and a corresponding uncertainty added (cf. section 5.5). The horizontal and

vertical emittances can be calculated directly from the tracker profiles via

$$\epsilon_x = \frac{1}{\beta_x(\phi_{st})} \Big(\sigma_x^2(\phi_{st}) - D_x^2(\phi_{st})\sigma_\delta^2\Big), \tag{5.31}$$

$$\epsilon_y = \frac{1}{\beta_y(\phi_{st})} \sigma_y^2(\phi_{st}), \tag{5.32}$$

such that the tracker profiles are scaled as follows:

$$x(\phi) = \langle x\rangle(\phi) + \sqrt{\frac{\beta_x(\phi)}{\beta_x(\phi_{st})} + \Big(D_x^2(\phi) - \frac{\beta_x(\phi)}{\beta_x(\phi_{st})} D_x^2(\phi_{st})\Big)\frac{\sigma_\delta^2}{\sigma_x^2(\phi_{st})}} \Big(x(\phi_{st}) - \langle x\rangle(\phi_{st})\Big), \tag{5.33}$$

$$y(\phi) = \langle y\rangle(\phi) + \sqrt{\frac{\beta_y(\phi)}{\beta_y(\phi_{st})}} \Big(y(\phi_{st}) - \langle y\rangle(\phi_{st})\Big). \tag{5.34}$$

The beam profiles of the two tracker stations are extrapolated to other azimuthal positions by shifting and scaling the histograms according to the above equations, with inputs described below. Each tracker profile is shifted and scaled independently, and the muon beam distribution at angle ϕ is taken as the sum of the two extrapolated histograms,

$$M_\phi(x, y) = M_{12}\big(x(\phi), y(\phi)\big) + M_{18}\big(x(\phi), y(\phi)\big). \tag{5.35}$$

Fig. 5.9 shows an example of the mean and width of the extrapolated beam for one run group, for each station independently as well as their sum, overlayed with the true mean and width measured at the tracker locations.

5.3.1 Momentum deviation

The mean momentum deviation $\langle\delta\rangle$ and momentum spread σ_δ are provided from the fast rotation analysis performed as part of the CERNext E-field correction [45]. At early times in a fill (significantly before the typical analysis start time of $t_0 \approx 30$ µs), each calorimeter sees a distinct positron pulse as the muon beam passes by. Each calorimeter will observe the pulse repeating at the cyclotron frequency, stemming from muons with nominal momentum p_0, but also at other frequencies depending on the the momentum distribution of the muons in the fill. This is called the "fast rotation" signal, and the mean momentum deviation as well as the momentum spread can be extracted from the Fourier spectrum of the fast rotation signal. Due to debunching of the muon beam, only a limited amount of data can be collected in each fill, and the fast rotation analysis is performed on the scale of omega-a

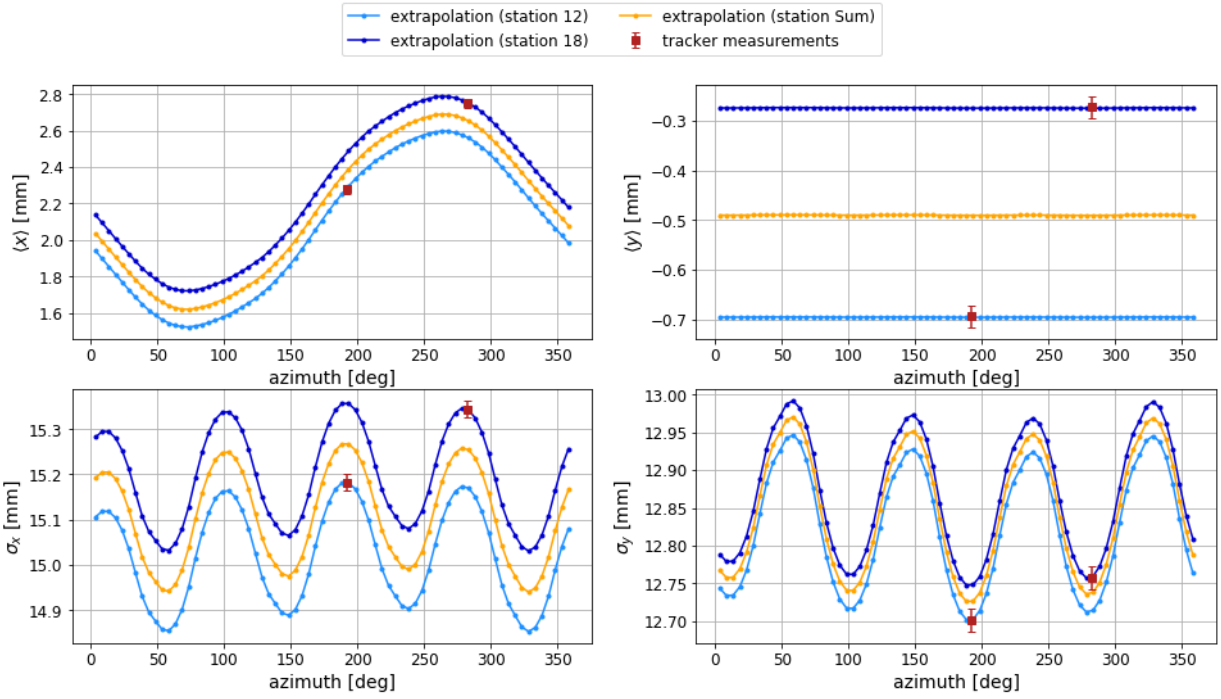


Figure 5.9: Example of beam extrapolation results, for a single run group. Mean position on top (radial left, vertical right), and respective RMS width on bottom. The extrapolation of each single station shows the mean position matched to the station measurements, whereas the station sum is effectively a weighted mean of the two stations. In this case, station 12 and 18 each registered a similar number of tracks, and the station sum is centered between the two.

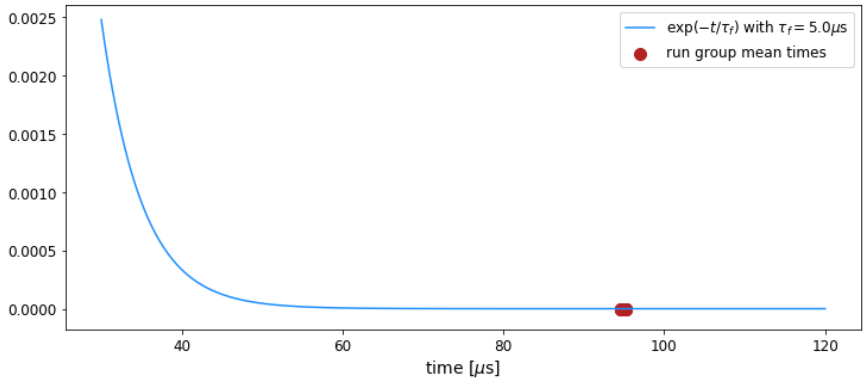


Figure 5.10: Time-dependence of the Run-3b COSY beam dynamics functions, with $\tau = 5.0\mu s$. Mean time in fill for every run group of trolley pair 4A5 is also plotted, clearly showing that the beam dynamics functions will have reached near constant values.

datasets to have sufficient statistics. Table 5.4 shows the CTAG-weighted averages for the RF-datasets of the mean momentum deviation $\langle\delta\rangle$ and the momentum deviation spread σ_δ during Run-4/5/6.

	noRF	xRF	xyRF5	xyRF6
mean momentum deviation $\langle\delta\rangle$ $[10^{-3}]$	0.396	0.173	0.210	0.183
momentum deviation spread σ_δ $[10^{-3}]$	1.173	1.143	1.159	1.163

Table 5.4: Summary of momentum deviation and spread for RF datasets. CTAG-weighted averages of mean momentum deviation and momentum spread for the RF datasets in Run-4/5/6.

5.3.2 Betatron and dispersion functions

The Muon g-2 COSY simulation module uses optical beam transfer maps to calculate a variety of beam dynamics functions, taking realistic fields and the discrete structure of the ESQs into account [139]. It has been extensively tested against measured data as well as other simulation packages that use traditional numerical integration methods. The radial and vertical betatron functions ($\beta_x(\phi)$, $\beta_y(\phi)$), as well as the radial dispersion function ($D_x(\phi)$), are obtained per omega-a dataset on a 1° azimuthal grid using measured field maps as input. The beam functions technically vary with the time in fill, but it was shown in Run-1 that the in-fill time dependence is a negligible effect for the muon weighted field, and instead, each run group's mean time in fill is used to evaluate the beam function values [118]. In previous

muon weighting analyses, each of the beam functions was parameterized as

$$f_b(\phi) = a_f(\phi) + b_f(\phi)\exp(-t/\tau_f), \tag{5.36}$$

where f_b is any of the beam dynamics functions, and $a_f(\phi)$, $b_f(\phi)$, and τ_f are function- and, for a_f and b_f, azimuth-dependent parameters determined from a fit to a series of simulated beam function values over time in fill. For Run-3b (the run period immediately preceding Run-4), $\tau_{\beta_x} = \tau_{\beta_y} = \tau_{D_x} = 5.0\mu s$, such that each of the beam dynamics functions reaches approximately constant values after about 60µs. Fig. 5.10 shows the time dependent term only of f_b, $\exp(-t/\tau_f)$, for the Run-3b value of 5.0µs. As an example, the mean in-fill times of all run groups in trolley pair 4A5 and their corresponding exponential value are also displayed, illustrating that the mean times vary little on a function that is already almost constant at that point. Because of this, for the Run-4/5/6 muon weighting analysis, the time-dependent parameterization of the beam functions is replaced by their value at $t = 94$ µs. The validity of this approach has been confirmed by evaluating the full functional form (eq. 5.36) of the Run-3b beam functions for each run group in a randomly chosen trolley pair (4A5), which lead to sub-ppt differences [106], far below the relevant order of magnitude. Figs. 5.11-5.12 show a direct comparison of the three beam dynamics functions in the Run-3b parameterization and evaluated at $t = 94$ µs, and the Run-4/5/6 parameterization (here using the values of dataset 4A), as well as the residuals of the Run-3b functions at each run group's mean time in fill with respect to $t = 94$ µs.

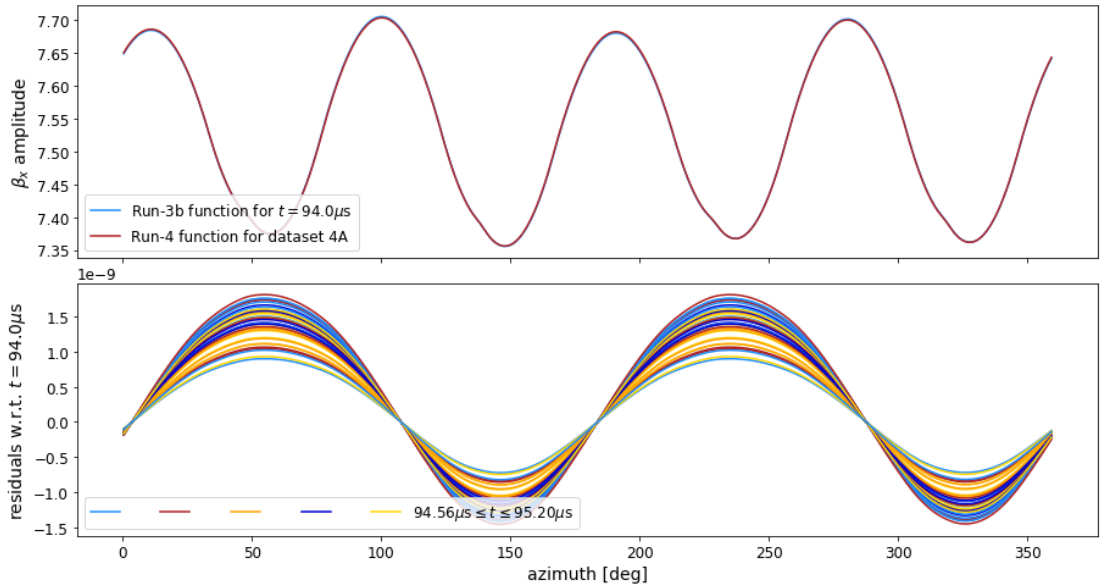
5.3.3 Radial closed orbit distortion

Assuming continuous ESQs and constant muon momentum, the radial closed orbit distortion Δx_{COD} can be approximated by a Fourier series expansion [89],

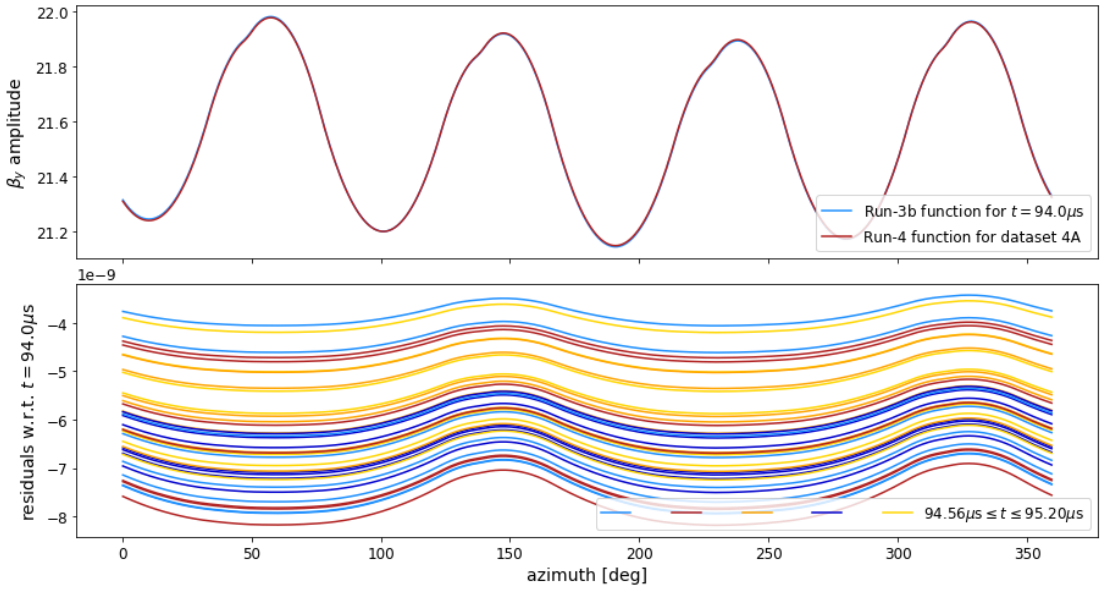
$$\Delta x_{COD} = -R_0 \sum_{N=0}^{\infty} \frac{\cos(N\phi + \phi_N)}{\nu_x^2 - N^2} \frac{b_N}{B_0}, \tag{5.37}$$

with the storage ring radius $R_0 = 7.112$ m, nominal magnetic field strength $B_0 = 1.4513$ T, horizontal tune $\nu_x = \sqrt{1-n} \approx 0.944$ (effective field index $n \approx 0.108$), and the Fourier amplitude b_N and phase ϕ_N of the dipole moment (i.e. m_1). The Fourier amplitude and phase are obtained for each run group from a standard Fourier transform of the linearly (in space) interpolated VTR.

Due to the inverse proportionality to $\nu_x^2 - N^2$ and the horizontal tune's size $\nu_x \lesssim 1$, the most significant term is $N = 1$, while all other terms vanish quickly. In the muon weighting



(a) Radial betatron function



(b) Vertical betatron function

Figure 5.11: Betatron functions vs azimuth. Each plot depicts on top the betatron amplitude as a function of azimuthal angle in the Run-4/5/6 parameterization (here for dataset 4A), and the Run-3b parameterization evaluated at $t = 94.0\mu s$, showing very good agreement. Below are the residuals of the Run-3b betatron function evaluated for all 4A5 run groups' mean times with respect to the $t = 94.0\mu s$ result. Differences are on the order of 10^{-9} and thus negligible.

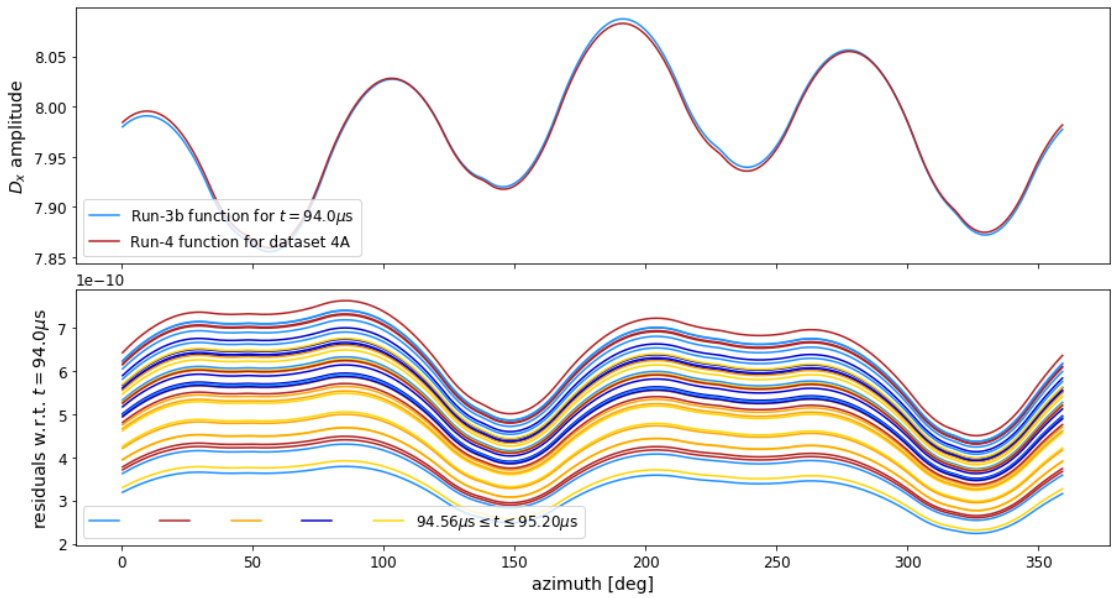


Figure 5.12: Radial dispersion function vs azimuth. On top: radial dispersion amplitude as a function of azimuthal angle in the Run-4/5/6 parameterization (here for dataset 4A), and the Run-3b parameterization evaluated at $t = 94.0$µs, showing very good agreement. Below: residuals of the Run-3b dispersion function evaluated for all 4A5 run groups' mean times with respect to the $t = 94.0$µs result. Differences are on the order of 10^{-10} and thus negligible.

analysis, only the dominant term is considered, such that [69]

$$\Delta x_{COD} \simeq \frac{R_0}{n}\frac{b_1}{B_0}\cos(\phi + \phi_1) \tag{5.38}$$

5.4 Run-4/5/6 muon weighting results

multipole	unweighted field	beam projections	weighted field
Dipole	61.790925 MHz	1.0	61.790925 MHz
NQ	-4.97 ppb	0.083872	7.46 ppb
SQ	-8.89 ppb	-0.021990	0.22 ppb
N6	55.42 ppb	0.047575	2.04 ppb
S6	-21.20 ppb	-0.001889	0.36 ppb
N8	21.42 ppb	-0.006151	0.04 ppb
S8	67.30 ppb	-0.002583	-0.18 ppb
N10	13.79 ppb	0.001841	-0.26 ppb
S10	66.88 ppb	0.001121	0.06 ppb
N12	-716.14 ppb	-0.000010	0.03 ppb
S12	-119.77 ppb	-0.000276	0.04 ppb
N14	-1067.74 ppb	0.000673	-0.72 ppb
S14	0.00 ppb	0.000055	0.00 ppb
Total	**61.790820 MHz**		**61.790926 MHz**

Table 5.5: Summary of muon weighting results for noRF dataset, including multipole contributions and total field (reported as $\tilde{\omega}_p/2\pi$). Dataset averages are CTAG weighted across run groups and azimuthally averaged. Dipole beam projections are 1.0 by definition.

Once the tracker profiles have been extrapolated to each azimuthal bin, the beam multipole projections can be calculated per eq. 5.15. The dominant beam projections coincide with the largest field multipoles: the NQ projection, which couples mostly to the mean radial position $\langle x \rangle$ of the beam, followed by the NS projection (coupling to both radial mean and width) and the SQ projection (coupling to the vertical mean). With the beam projections, and the binned VTRs, the weighted field per run group is calculated according to eq. 5.13. The total weighted field for each RF-dataset is the CTAG-weighted average of the weighted field per run group, eq. 5.11. Alternatively and equivalently, the sum over the 13 multipoles can be performed last, which allows every multipole to be weighted separately before combining their contributions into the total weighted field. Tables 5.5-5.8 summarize the results for each of the RF-datasets, where the RF-dataset average of m_j (unweighted field), k_j (beam projections), and $m_j\ k_j$ (weighted field) is calculated as the standard arithmetic mean in azimuth, and the CTAG-weighted average across run groups. Both the total field

multipole	unweighted field	beam projections	weighted field
Dipole	61.790944 MHz	1.0	61.790944 MHz
NQ	55.37 ppb	0.039456	0.66 ppb
SQ	-39.32 ppb	-0.017674	0.73 ppb
N6	25.66 ppb	0.031178	-0.58 ppb
S6	-8.73 ppb	0.000366	-0.08 ppb
N8	17.34 ppb	-0.002624	-0.13 ppb
S8	70.33 ppb	-0.001375	-0.10 ppb
N10	11.98 ppb	0.002887	-0.07 ppb
S10	86.33 ppb	0.000598	0.07 ppb
N12	-726.38 ppb	0.000640	-0.47 ppb
S12	-126.94 ppb	-0.000262	0.04 ppb
N14	-1067.86 ppb	-0.000270	0.28 ppb
S14	0.00 ppb	-0.000003	0.00 ppb
Total	**61.790839 MHz**		**61.790944 MHz**

Table 5.6: Summary of muon weighting results for xRF dataset, including multipole contributions and total field (reported as $\tilde{\omega}_p/2\pi$). Dataset averages are CTAG weighted across run groups and azimuthally averaged. Dipole beam projections are 1.0 by definition.

multipole	unweighted field	beam projections	weighted field
Dipole	61.790916 MHz	1.0	61.790916 MHz
NQ	88.97 ppb	0.047985	9.76 ppb
SQ	59.57 ppb	-0.009856	-0.48 ppb
N6	-19.52 ppb	0.035582	-1.76 ppb
S6	87.70 ppb	0.001103	0.12 ppb
N8	48.98 ppb	-0.002898	0.02 ppb
S8	71.72 ppb	-0.000729	-0.06 ppb
N10	1.04 ppb	0.003254	-0.13 ppb
S10	84.37 ppb	0.000572	0.05 ppb
N12	-727.49 ppb	0.000723	-0.52 ppb
S12	-127.32 ppb	-0.000196	0.03 ppb
N14	-1059.02 ppb	0.000062	-0.07 ppb
S14	0.00 ppb	0.000012	0.00 ppb
Total	**61.790824 MHz**		**61.790917 MHz**

Table 5.7: Summary of muon weighting results for xyRF5 dataset, including multipole contributions and total field (reported as $\tilde{\omega}_p/2\pi$). Dataset averages are CTAG weighted across run groups and azimuthally averaged. Dipole beam projections are 1.0 by definition.

as well as the individual multipole contributions are listed. Note that due to in particular azimuthal variations in the field and beam projections, the dataset-averaged weighted field is not necessarily close to the dataset-averaged unweighted field multiplied by the dataset-

multipole	unweighted field	beam projections	weighted field
Dipole	61.790908 MHz	1.0	61.790908 MHz
NQ	11.59 ppb	0.048753	5.15 ppb
SQ	-26.19 ppb	-0.023833	0.60 ppb
N6	19.54 ppb	0.038734	-0.30 ppb
S6	8.96 ppb	-0.000369	0.06 ppb
N8	27.38 ppb	-0.001648	0.05 ppb
S8	45.53 ppb	-0.002467	-0.11 ppb
N10	16.2 ppb	0.004244	-0.14 ppb
S10	60.51 ppb	0.000630	0.06 ppb
N12	-727.72 ppb	0.000710	-0.51 ppb
S12	-133.95 ppb	-0.000514	0.07 ppb
N14	-1054.13 ppb	0.000241	-0.26 ppb
S14	0.00 ppb	-0.000024	0.00 ppb
Total	**61.790799 MHz**		**61.790909 MHz**

Table 5.8: Summary of muon weighting results for xyRF6 dataset, including multipole contributions and total field (reported as $\tilde{\omega}_p/2\pi$). Dataset averages are CTAG weighted across run groups and azimuthally averaged. Dipole beam projections are 1.0 by definition.

averaged beam projection.

The weighted and unweighted field per run group are shown for each RF-dataset in figs. 5.13-5.16.

5.5 Run-4/5/6 muon weighting uncertainties

Both the beam and the field contribute to the overall uncertainty of $\langle \tilde{B} \rangle_{\text{dataset}}$. Starting from eq. 5.11, the dataset uncertainty can be expressed as

$$\sigma^2_{\langle \tilde{B} \rangle} = \text{Var}(\langle \tilde{B} \rangle) = \frac{1}{N^2_{\text{dataset}}} \text{Var}\Big(\sum_{\text{run groups}} N_{rg}\ \tilde{B}_{rg} \Big). \tag{5.39}$$

This sum will be resolved differently depending on whether the run group uncertainties are correlated or independent. Many uncertainty sources have to be considered, some of which will be correlated between run groups (i.e. in time), and others will not, so the above

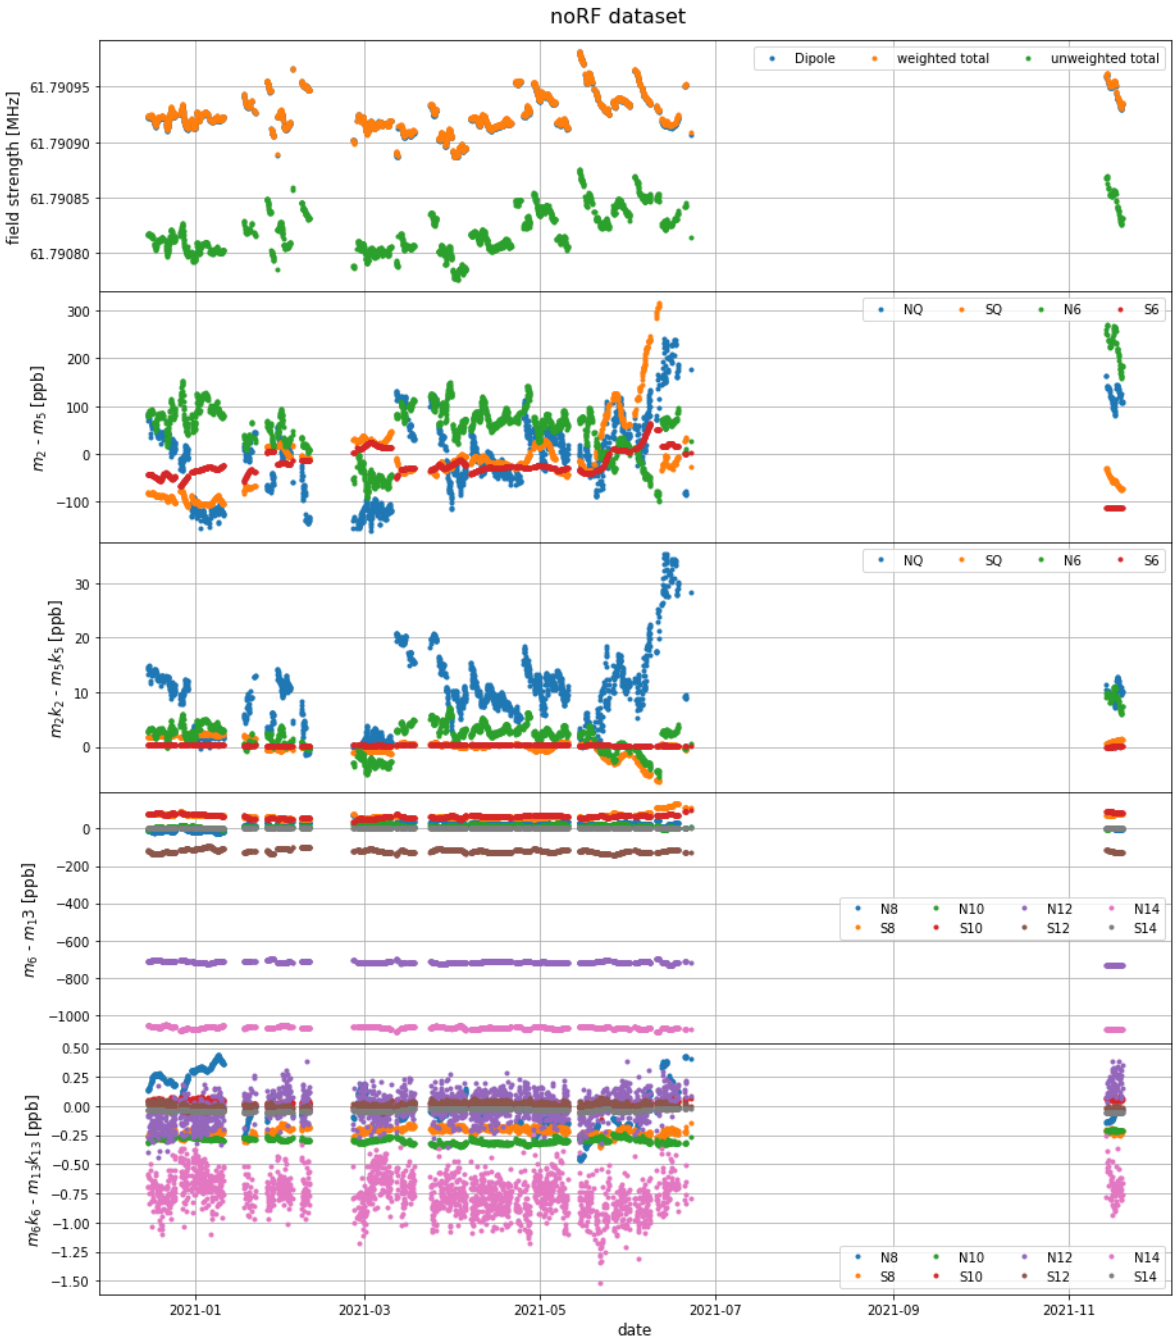


Figure 5.13: Muon weighting results for the noRF dataset. All plots show azimuthal averages per run group over time. Top: blinded dipole moment (blue), total blinded unweighted field (green), and total blinded weighted field (orange). The weighted field includes corrections from quad misalignments, detailed in section 5.5.2.4. 2nd from top: individual contributions for the four largest multipole moments (NQ, SQ, N6 and S6), unweighted. Center: individual muon-weighted contributions for the same four multipole moments (NQ, SQ, N6 and S6), which are tracked by fixed probes. 2nd from bottom: remaining unweighted multipole moments. Bottom: muon-weighted contributions of the remaining multipole moments.

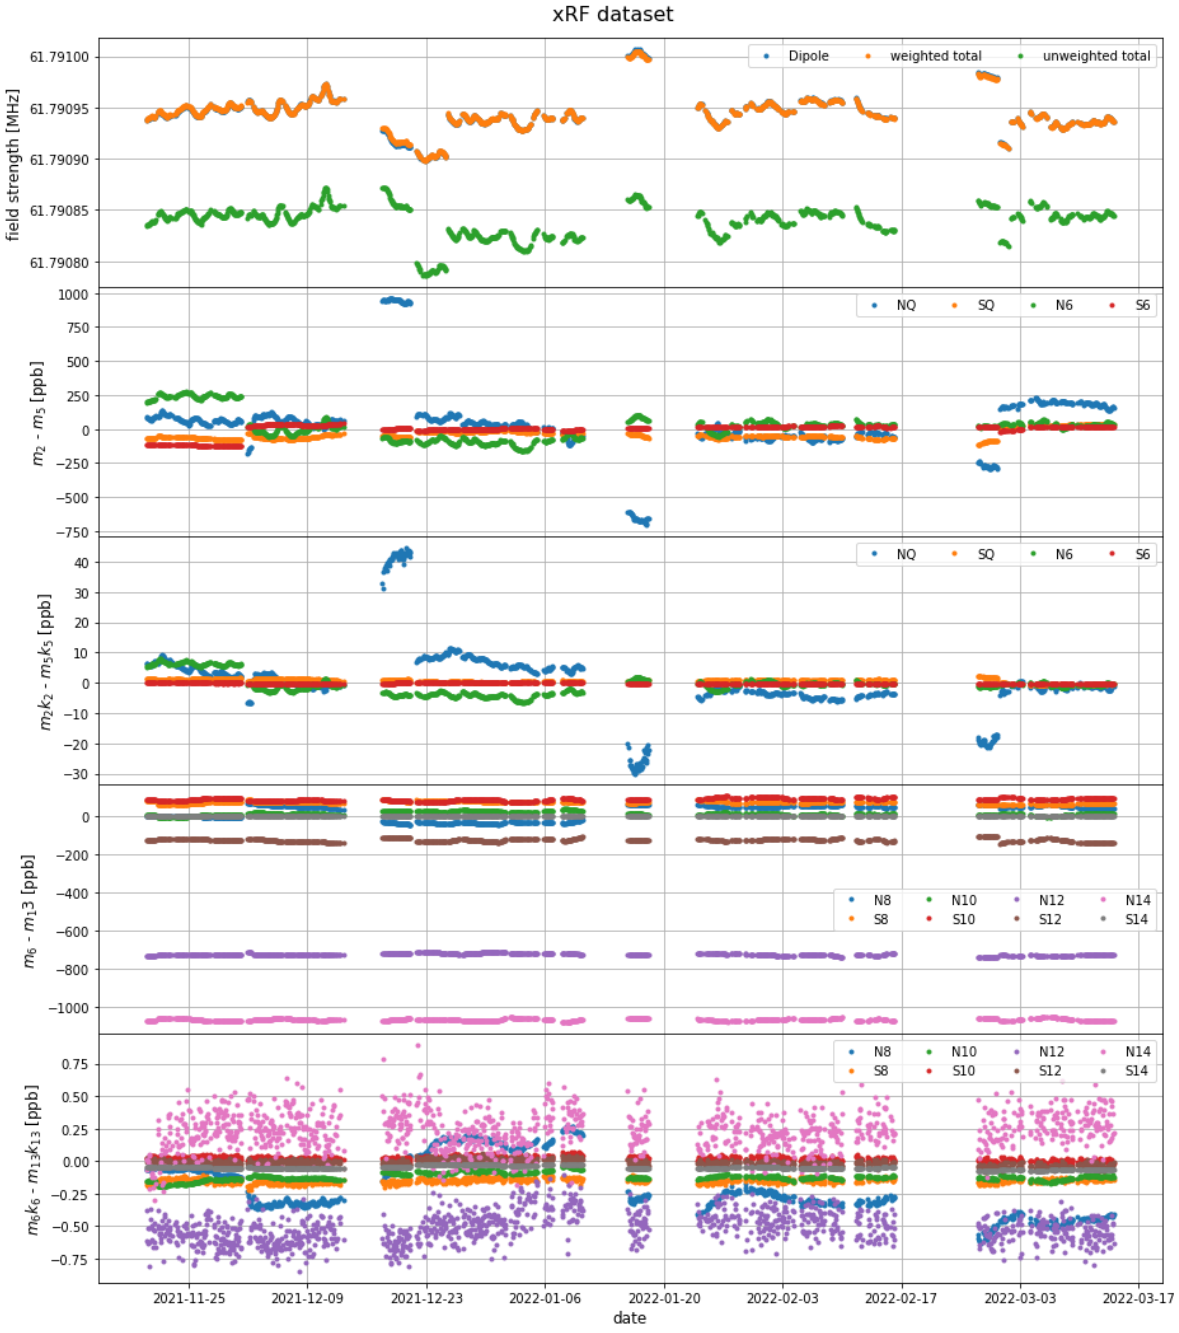


Figure 5.14: Muon weighting results for the xRF dataset. All plots show azimuthal averages per run group over time. Top: blinded dipole moment (blue), total blinded unweighted field (green), and total blinded weighted field (orange). The weighted field includes corrections from quad misalignments, detailed in section 5.5.2.4. 2nd from top: individual contributions for the four largest multipole moments (NQ, SQ, N6 and S6), unweighted. Center: individual muon-weighted contributions for the same four multipole moments (NQ, SQ, N6 and S6), which are tracked by fixed probes. 2nd from bottom: remaining unweighted multipole moments. Bottom: muon-weighted contributions of the remaining multipole moments.

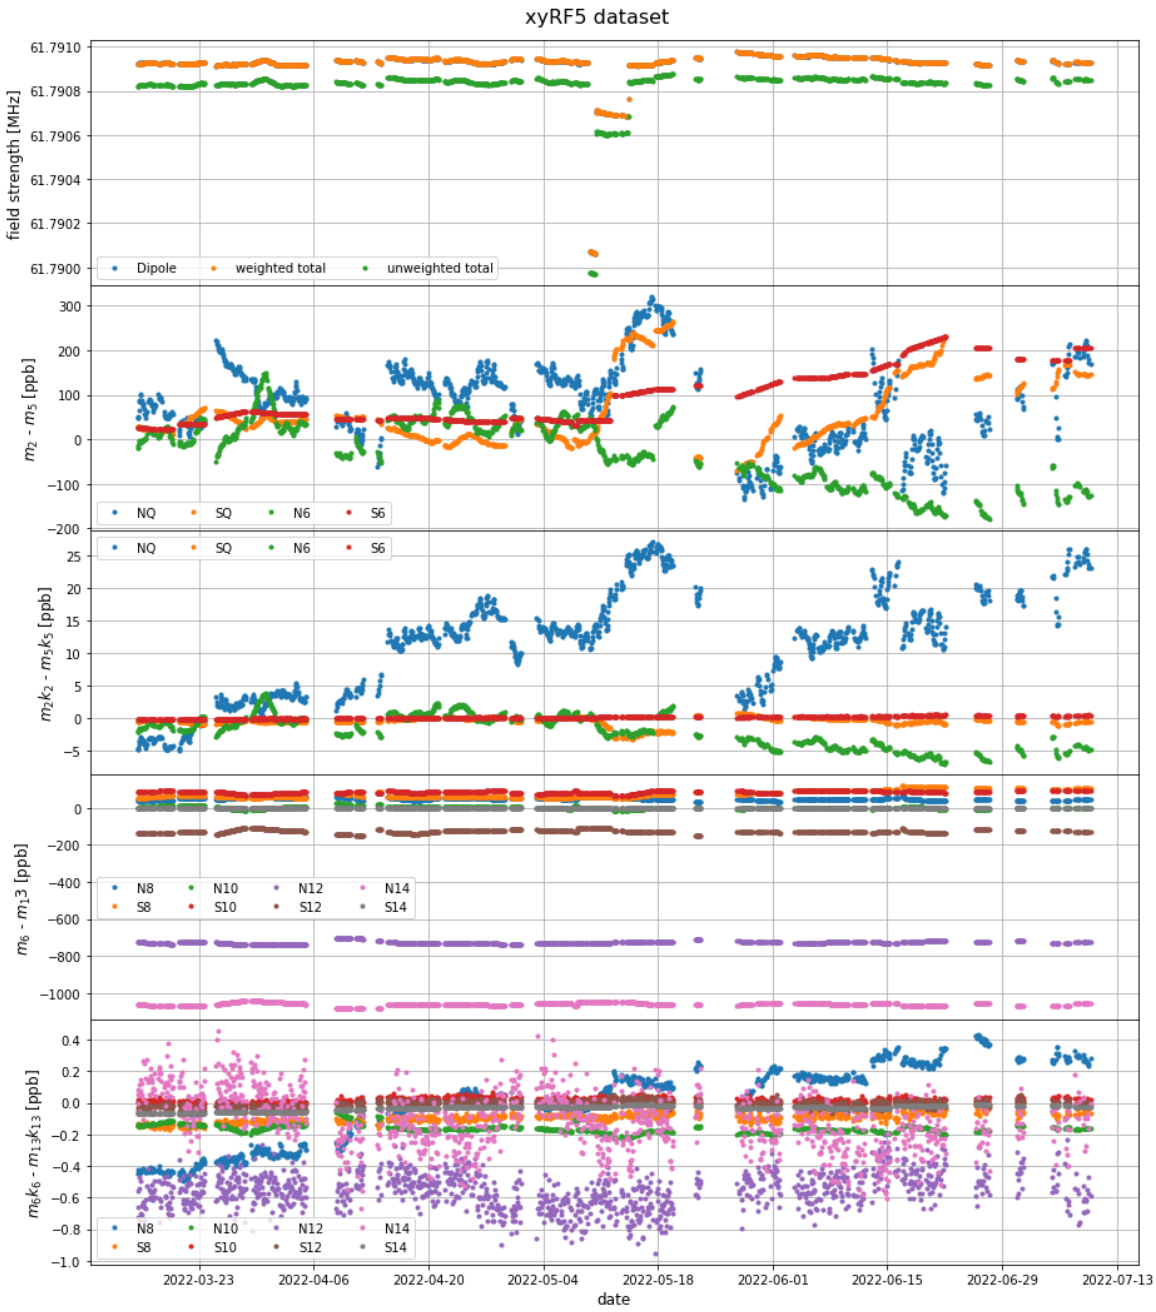


Figure 5.15: Muon weighting results for the xyRF5 dataset. All plots show azimuthal averages per run group over time. Top: blinded dipole moment (blue), total blinded unweighted field (green), and total blinded weighted field (orange). The weighted field includes corrections from quad misalignments, detailed in section 5.5.2.4. 2nd from top: individual contributions for the four largest multipole moments (NQ, SQ, N6 and S6), unweighted. Center: individual muon-weighted contributions for the same four multipole moments (NQ, SQ, N6 and S6), which are tracked by fixed probes. 2nd from bottom: remaining unweighted multipole moments. Bottom: muon-weighted contributions of the remaining multipole moments.

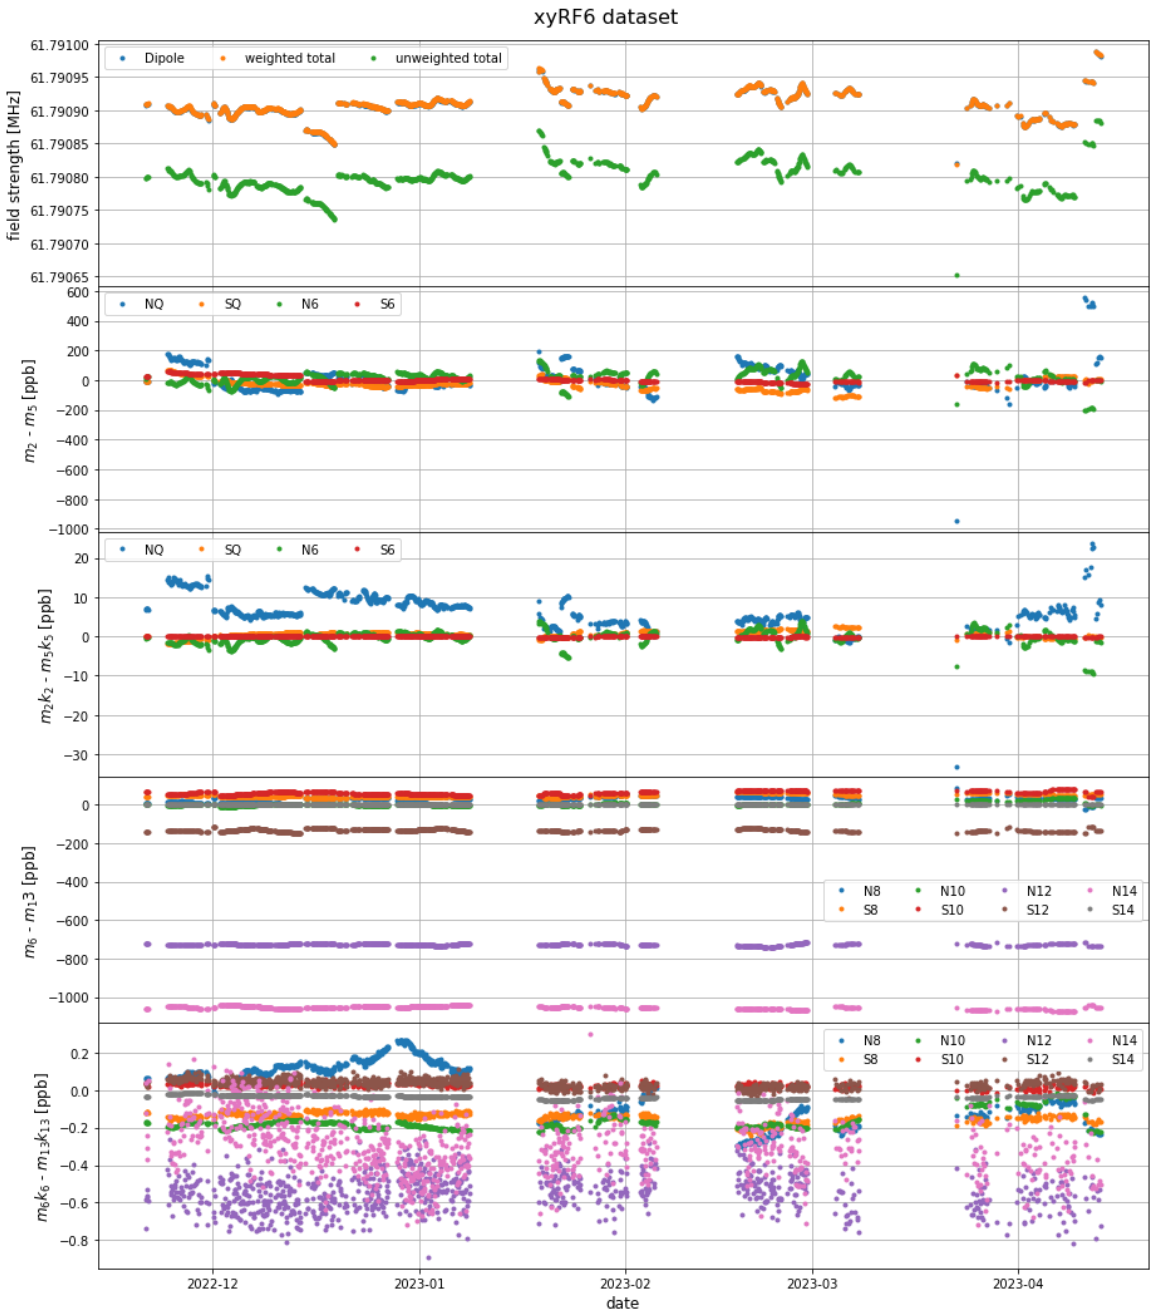


Figure 5.16: Muon weighting results for the xyRF6 dataset. All plots show azimuthal averages per run group over time. Top: blinded dipole moment (blue), total blinded unweighted field (green), and total blinded weighted field (orange). The weighted field includes corrections from quad misalignments, detailed in section 5.5.2.4. 2nd from top: individual contributions for the four largest multipole moments (NQ, SQ, N6 and S6), unweighted. Center: individual muon-weighted contributions for the same four multipole moments (NQ, SQ, N6 and S6), which are tracked by fixed probes. 2nd from bottom: remaining unweighted multipole moments. Bottom: muon-weighted contributions of the remaining multipole moments.

equation can be split into two terms:

$$\begin{aligned}\sigma^2_{\langle\tilde{B}\rangle} &= \frac{1}{N^2_{\text{dataset}}}\Big[\,\text{Var}_c(\sum_{rg} N_{rg}\ \tilde{B}_{rg}) + \text{Var}_i(\sum_{rg} N_{rg}\ \tilde{B}_{rg})\Big] \\ &= \frac{1}{N^2_{\text{dataset}}}\Big[\Big(\sum_{rg} N_{rg}\ \sigma_{c,\tilde{B}_{rg}}\Big)^2 + \sum_{rg} N_{rg}\ \sigma^2_{i,\tilde{B}_{rg}}\Big],\end{aligned} \tag{5.40}$$

where c denotes correlated uncertainty contributions and i denotes independent or uncorrelated contributions.
In either case, the run group uncertainty contribution has to be determined first. From eq. 5.13,

$$\text{Var}(\tilde{B}_{rg}) = \frac{1}{N^2_\phi}\,\text{Var}\,\Big(\sum_{i=1}^{N_\phi}\sum_{j=1}^{13} m_j(\phi_i)\ k_j(\phi_i)\Big) \tag{5.41}$$

Let $n = (i, j)$ be the combined index such that $n = 1, ..., 13 \cdot N_\phi$ and $m_n = m_{(i,j)} = m_j(\phi_i)$ and similarly for k. Then the uncertainty can be propagated via the covariance [116]

$$\text{Var}(\tilde{B}_{rg}) \propto \text{Var}\,\Big(\sum_{n=1}^{13\cdot N_\phi} m_n\ k_n\Big) = \sum_{x=m,k}\ \sum_{n,n'} \frac{\partial\tilde{B}_{rg}}{\partial x_n}\frac{\partial\tilde{B}_{rg}}{\partial x_{n'}}\,\text{Cov}(x_n, x_{n'}). \tag{5.42}$$

The field multipole strengths are generally uncorrelated with the beam projections, so the cross-term $\text{Cov}(m_n, k_{n'}) = 0$, and the field and beam uncertainties contribute independently:

$$\begin{aligned}\text{Var}(\tilde{B}_{rg}) &\propto \sum_{n,n'} k_n\ k_{n'}\ \ \text{Cov}(m_n, m_{n'}) + \sum_{n,n'} m_n\ m_{n'}\ \ \text{Cov}(k_n, k_{n'}) \\ &= \text{Var}_m(\tilde{B}_{rg}) + \text{Var}_k(\tilde{B}_{rg}).\end{aligned} \tag{5.43}$$

The field multipole moments are correlated as discussed previously, since higher orders alias into lower orders especially for the fixed probes. The beam multipole projections k_j however are independent of each other, i.e. $\text{Cov}\left(k_j(\phi_i), k_{j'}(\phi_{i'})\right) = 0$ for $j \neq j'$ [116, 69] since the number of muon decay vertices used to determine k_j is multiple orders of magnitude larger than the number of beam projections (13). Since all azimuthal beam profiles are constructed as a shifted and scaled sum of the two tracker profiles, the k_j are fully correlated across ϕ, i.e. $\text{Cov}\left(k_j(\phi_i), k_j(\phi_{i'})\right) = \sqrt{\text{Var}\left(k_j(\phi_i)\right)\text{Var}\left(k_j(\phi_{i'})\right)}$.
Thus the uncertainty contributions from the beam projections can be summarized as

$$\text{Var}_k(\tilde{B}_{rg}) = \sum_j \Big(\frac{1}{N_\phi}\sum_i m_j(\phi_i)\sigma_{k_j(\phi_i)}\Big)^2. \tag{5.44}$$

The following sections will discuss the uncertainty sources arising from the beam multipole projections. Uncertainties arising from the field measurements are not the focus of this work, with the exception of the "azimuthal averaging" uncertainty due to the different interpolation orders. This will be covered in section 5.5.4.

5.5.1 Tracking uncertainties

There are three sources of uncertainty associated with the tracker profile reconstruction: The tracker resolution (how precisely a muon decay vertex can be located), the tracker acceptance (how likely a decay positron will be registered and successfully traced back), and the tracker alignment (how well the relative position of the tracking detectors to the storage ring is known).

The strategy to evaluate these uncertainties is the same for all of them: for each run group and tracker station, the corresponding value is varied by $\pm 1\sigma$ (i.e. the resolution and acceptance correction values are modified, and a constant offset is added to the reconstructed vertices for a misalignment), and the resulting tracker profiles are propagated through the rest of the muon weighting analysis all the way to the weighted field per run group. The uncertainty of the run group weighted field is then the difference between the nominal weighted field and the weighted field based on the varied tracker profiles, and is slightly asymmetric in most cases.

Generally, the "true value" of each of these uncertainties (i.e. the change to the corresponding correction, or the change to the alignment) can be considered constant across run groups, since the uncertainty sources are hardware-based and fixed. During Run-4/5/6, a limited number of changes were made to the tracking detectors, e.g. swapping a single module in one tracker station. Such changes took place exactly four times, namely between omega-a datasets 4D/4E (mid-January 2021, trolley pairs 5R3/5R4 (late May 2022), runs 5&6 (July-November 2022), and trolley pairs 6A1/6A2 (late November 2022) [69]. This means all RF-datasets excecpt for xRF have one hardware change. These changes introduce the potential for a change in the "true" uncertainty values, but within periods of unchanged tracker hardware, the run group weighted fields can be averaged in the same way as usual. Thus, for each tracker uncertainty ("*tr.u.*") and variation ($\pm\sigma$), the weighted field can be calculated as the usual CTAG-weighted average of the run group weighted fields *within a period of unchanged tracker hardware*. Then we can take the total uncertainty on the weighted field due to a particular tracker-related uncertainty (tracker alignment, tracker resolution,

or tracker acceptance) as

$$\sigma_{tr.u.} = \sum_{\substack{\text{periods of} \\ \text{unchanged} \\ \text{hardware } T}} w_{\text{CTAG}} \cdot \max(|\langle \tilde{B} \rangle_T - \langle \tilde{B}_{+1\sigma} \rangle_T|, |\langle \tilde{B} \rangle_T - \langle \tilde{B}_{-1\sigma} \rangle_T|), \tag{5.45}$$

meaning the total weighted fields are calculated per period of unchanged hardware in the same manner they are for the RF-datasets, the uncertainty for that period is taken as the larger of the two variations compared to the nominal field in that period, and the resulting uncertainties are combined as fully correlated uncertainties. This is likely an overestimation, but since each hardware change only changes one element of the tracker (e.g. one module), the uncertainties are still highly correlated.

Table 5.9 shows the tracker uncertainties for each of the time periods of unchanged tracker hardware, as well as the CTAG fraction of those time periods to their corresponding RF-dataset. Note that since there was no hardware change in the xRF dataset, but the change of RF configuration constitutes a significant change in run conditions, it makes up a "period of unchanged tracker hardware" by itself.
The combined uncertainties per RF-dataset are listed in the uncertainty summary table 5.15.

Tracker correction	4A-4D	4E-5A	5B-5L	5M-5R3	5R4-5U	6A1	6A2-6K
Acceptance [ppb]	0.21	0.08	0.11	0.28	0.05	0.08	0.03
Resolution [ppb]	0.03	0.04	0.08	0.09	0.05	0.07	0.07
Alignment x [ppb]	0.75	0.05	0.41	1.36	0.22	0.44	0.23
Alignment y [ppb]	1.07	0.46	0.10	1.25	1.73	0.24	0.02
CTAG fraction [%]	16.4	83.6	100	69.7	30.3	0.6	99.4

Table 5.9: Uncertainty results for each tracker correction, per time period of unchanged tracker hardware. Also listed: Fraction of time period CTAGs to total CTAGs of corresponding RF-dataset.

5.5.2 Beam dynamics inputs

5.5.2.1 Beam functions

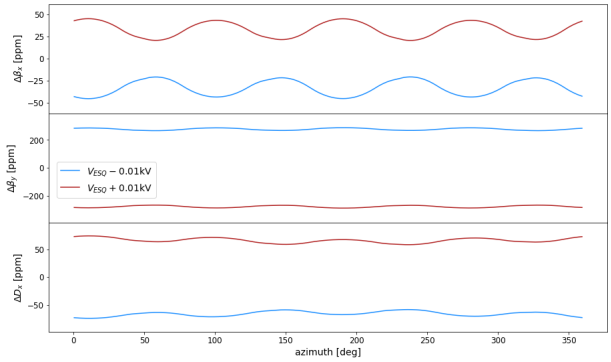
Uncertainties arising from a potential error in the beam functions are evaluated for the first time for Run-4/5/6. While less time- and computationally intensive than other simulation packages, the COSY-based simulation package still requires careful resource management, and thus variations of the beam functions were only tested on select datasets. The numerical error on the simulation results is negligible [138], and general accuracy of results has been validated such that they can be considered systematically reliable. However, an error on the inputs to the simulation could still lead to an error on the beam functions. The tested scenarios are the following:

- Dipole strength (m_1) varied by ±300ppb in dataset 5Q
- Normal quadrupole strength (m_2) varied by ±100ppb in dataset 5E
- Skew quadrupole strength (m_3) varied by ±100ppb in dataset 5R
- ESQ high voltage setting varied by ±0.01kV in dataset 5Q

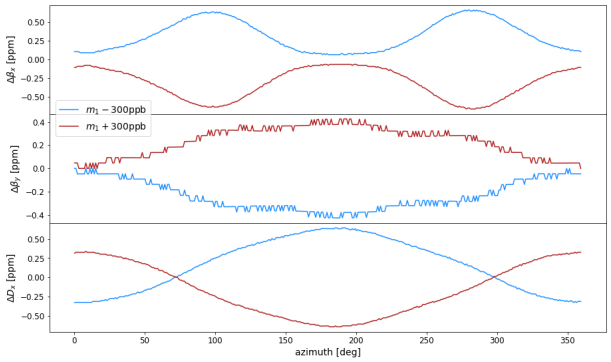
These variations are exaggerated compared to their expected uncertainties, but the effect on the weighted field is expected to be very small. Each parameter was only varied for one ω_a dataset (chosen based on which dataset showed the largest variance of the respective parameter; for the HV parameter, the dataset with the largest total field variance was chosen). Fig. 5.17 shows the change in each of the three beam functions from the above variations. Fractional changes are reported with respect to the nominal beam function values, i.e. $\Delta f_b(\phi) = \left(f_b^{varied}(\phi) - f_b^{nominal}(\phi)\right)/f_b^{nominal}(\phi)$.
The total weighted field of each dataset changed by less than 0.01 ppb due to these variations, and confirms that the COSY uncertainty is negligible.

5.5.2.2 Momentum deviation

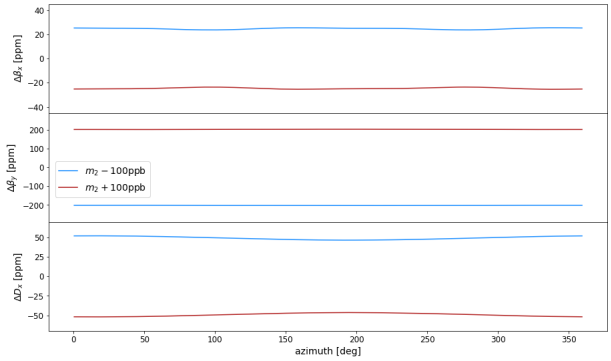
The uncertainties on each of the $\langle\delta\rangle$ and σ_δ values are propagated in the same way the tracker uncertainties are, which is by varying them independently by $\pm1\sigma$ and extrapolating the tracker profiles with the new values for all run groups and azimuth, and calculating the beam projections and weighted field from the varied extrapolated beam profiles as usual. As for the beam functions, the effect on the weighted field is negligible (< 0.01 ppb per dataset).



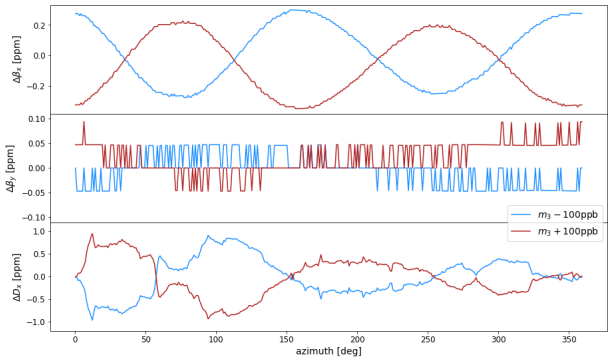
(a) HV ±0.01kV in dataset 5Q



(b) Dipole (m_1) ±300ppb in dataset 5Q



(c) m_2 ±100ppb in dataset 5E



(d) m_3 ±100ppb in dataset 5R

Figure 5.17: Varied beam function residuals (with respect to nominal function values) for uncertainty estimates.Top to bottom: radial betatron, vertical betatron, radial dispersion.

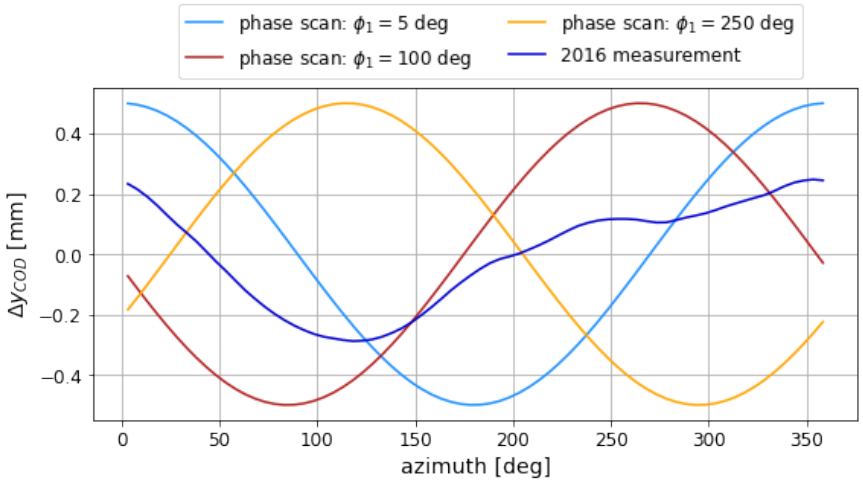


Figure 5.18: Shift of mean vertical beam position Δy_{COD}, due to radial field variations (radial field COD). Shown are three potential scenarios of a simple 0.5 mm Fourier amplitude COD, based on Fourier phase, as well as the COD that would result from the radial field measurements performed before Run-1 [126, 116].

5.5.2.3 COD due to radial field

Due to the experiment's inability to measure the radial component of the magnetic field, a vertical COD caused by azimuthal variations of the radial field can't be incorporated directly into the extrapolation of the tracker profiles. Instead, a systematic uncertainty is assessed guided by the radial field measurements performed before Run-1 [126]. A vertical COD can be expanded as a Fourier series in the same way as the radial COD (eq. 5.38. As in Run-2/3, this systematic uncertainty will be evaluated with a very conservative estimate of the $N = 1$ Fourier amplitude of 0.5 mm, about twice of what was measured before Run-1 [117]. Fig. 5.18 shows the resulting vertical shift of the beam for three different (randomly chosen) phases, as well as the COD that would correspond to the 2016 pre-Run-1 radial field measurements [126], obtained from COSY simulation [116]. Fig. 5.18 illustrates that the 0.5 mm Fourier amplitude estimate is significantly larger than even the 2016 measurement, and the radial field is expected to be smaller still due to shimming. It is also clear that the magnetic field experienced by the beam, and thus the effect of a vertical COD on the weighted field, depends strongly on the Fourier phase. To determine the largest possible effect, a range of phases will be evaluated, and the maximum change in weighted field across all phases will be taken as the uncertainty.

A COD only affects the mean of the beam profiles, and the single station extrapolation (eq. 5.34) changes to

$$\langle y\rangle(\phi) = \Delta y_{COD}(\phi) + \big(\langle y\rangle(\phi_{st}) - \Delta y_{COD}(\phi_{st})\big), \tag{5.46}$$

i.e. the COD motion is offset to match the measured tracker position. At each ϕ, the total beam profile is the weighted sum of the extrapolated tracker profiles, that is

$$\begin{aligned}\langle y\rangle(\phi) = &\frac{N_{12}}{N_{12}+N_{18}}\Big(\langle y\rangle(\phi_{12}) + \Delta y_{COD}(\phi) - \Delta y_{COD}(\phi_{12})\Big) + \\ &+ \frac{N_{18}}{N_{12}+N_{18}}\Big(\langle y\rangle(\phi_{18}) + \Delta y_{COD}(\phi) - \Delta y_{COD}(\phi_{18})\Big) \\ = &\frac{N_{12}\langle y\rangle(\phi_{12}) + N_{18}\langle y\rangle(\phi_{18})}{N_{12}+N_{18}} + \Delta y_{COD}(\phi) - \frac{N_{12}\Delta y_{COD}(\phi_{12}) + N_{18}\Delta y_{COD}(\phi_{18})}{N_{12}+N_{18}}.\end{aligned} \tag{5.47}$$

If the vertical COD is already centered between the two stations so that $N_{12}\Delta y_{COD}(\phi_{12}) + N_{18}\Delta y_{COD}(\phi_{18}) = 0$, then $\langle y\rangle(\phi)$ is simply shifted by $\Delta y_{COD}(\phi)$, and the skew quadrupole beam projection k_3 (cf. table 5.1) gains a linear dependence on $\Delta y_{COD}(\phi)$. For sufficiently small $\Delta y_{COD}(\phi)$, higher order moments also vary approximately linearly with

COD amplitude [116], such that

$$\Delta k_j(\phi) \approx \frac{\mathrm{d}k_j(\phi)}{\mathrm{d}y}\Delta y_{COD}(\phi), \tag{5.48}$$

and the change in the weighted field due to a vertical COD can be estimated as

$$\Delta\tilde{B} = \frac{1}{N_\phi}\sum_{i=1}^{N_\phi}\sum_{j=1}^{13}\Delta k_j(\phi_i)m_j(\phi_i) \approx \frac{1}{N_\phi}\sum_{i=1}^{N_\phi}\sum_{j=1}^{13}\frac{\mathrm{d}k_j(\phi_i)}{\mathrm{d}y}\Delta y_{COD}(\phi_i)m_j(\phi_i). \tag{5.49}$$

By using the already existing (nominal) extrapolated beam histograms, shifting them manually by a range of (small) test displacements, and calculating the beam functions for the shifted histograms, $\frac{\mathrm{d}k_j(\phi_i)}{\mathrm{d}y}$ can be determined in a computationally more efficient way than if the entire extrapolation step had to be completed multiple times. Additionally, since a COD from an unshimmed radial field is expected to be fairly constant over time, this estimation can be performed at the trolley pair level. For each trolley pair and each azimuthal 5° bin, 72 COD shifts in each k_j are calculated (for 72 Fourier phases between 0 and 2π). The Δk_j are combined with trolley-pair-averaged VTRs in the usual way to obtain the total $\Delta\tilde{B}_{tp}(\phi_i)$, and then azimuthally averaged. This results in 72 values for $\Delta\tilde{B}_{tp}$, depending on COD phase. The largest $\Delta\tilde{B}_{tp}$ is chosen for each trolley pair, and then the CTAG-average is taken to obtain the shift of the weighted field per RF-dataset, $\Delta\tilde{B}_{RF}$. Table 5.10 shows the results.

	noRF	**xRF**	**xyRF5**	**xyRF6**
y_{COD} from B_{rad} [ppb]	2.42	2.05	1.57	2.08

Table 5.10: COD due to B_r uncertainties. Summary of uncertainties per RF-dataset from a potential radial field COD.

5.5.2.4 COD due to ESQ misalignments

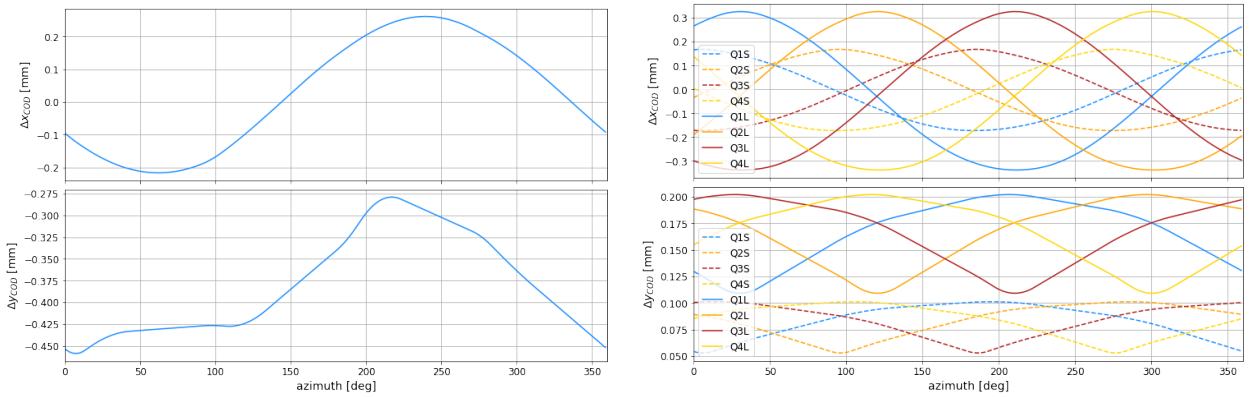
In addition to field variations, another source of radial and vertical CODs is a misalignment of the ESQ plates. During Run-1, the ESQ centroid displacements were determined from measurements of the ESQ plate misalignments [116, 137], summarized in table 5.11. Using a transport matrix technique similar to the COSY simulation, the radial and vertical CODs that arise from these centroid displacements can be calculated [137], shown in fig, 5.19a. The effects of displacements of the 8 individual ESQ sections add linearly to produce the COD, and scale linearly with displacement as well. Hence, "unit CODs", which are the

ESQ section	δx [mm]	$\sigma_{\delta x}$ [mm]	δy [mm]	$\sigma_{\delta y}$ [mm]
Q1S	−0.490	0.700	2.235	0.354
Q1L	−0.245	0.492	−0.980	0.350
Q2S	−0.995	0.430	−0.675	0.354
Q2L	−0.035	0.458	0.380	0.350
Q3S	−0.200	0.480	0.750	0.350
Q3L	0.115	0.437	−1.965	0.354
Q4S	−0.345	0.700	−1.315	0.354
Q4L	0.085	0.522	−0.220	0.350

Table 5.11: ESQ centroid displacements. Run-1 measurements and uncertainties of ESQ centroid displacements due to ESQ plate misalignments.

CODs due to a unit (1mm) displacement in x or y of a single ESQ section, can be used to construct the total COD, and allow for variation of the ESQ section displacements within their uncertainties. The 16 unit CODs are shown in fig. 5.19b.
With the radial and vertical COD determined from the ESQ misalignment, the shift of the beam projections and weighted field can be calculated in the same manner as for the vertical COD due to a radial field. The $\frac{\mathrm{d}k_j(\phi_i)}{\mathrm{d}x}$ dependence is determined through the same process of shifting the nominal beam histograms radially, calculating the new beam projections, and fitting a linear slope to the results. Similarly to the radial field COD, the ESQ misalignments are not expected to change over time, and the changes are evaluated on the trolley pair level.

Contrary to the radial field COD, where it is assumed that the radial field is zero and only an uncertainty is assessed, the measurement of the ESQ misalignments shows that a COD exists that is not accounted for in the extrapolation algorithm. Therefore a correction must be made, which corresponds to the effect of the nominal misalignments, and an additional uncertainty due to the misalignment uncertainties has to be determined. The latter task is significantly more complex than most of the other muon weighting uncertainties, because of the number of possible variations between the 8 ESQ sections. For one direction (radial or vertical), even just varying each measurement between $(-\sigma, 0, +\sigma)$ creates $3^8 = 6,561$ different CODs ("corner cases"). As in Run-2/3, a Monte Carlo (MC) method will be used to determine both the correction and uncertainty [117]. A total of $10,000$ samples are produced, by drawing from Gaussian distributions centered at the



(a) Nominal ESQ misalignment

(b) Unit displacements of individual ESQ section centroids

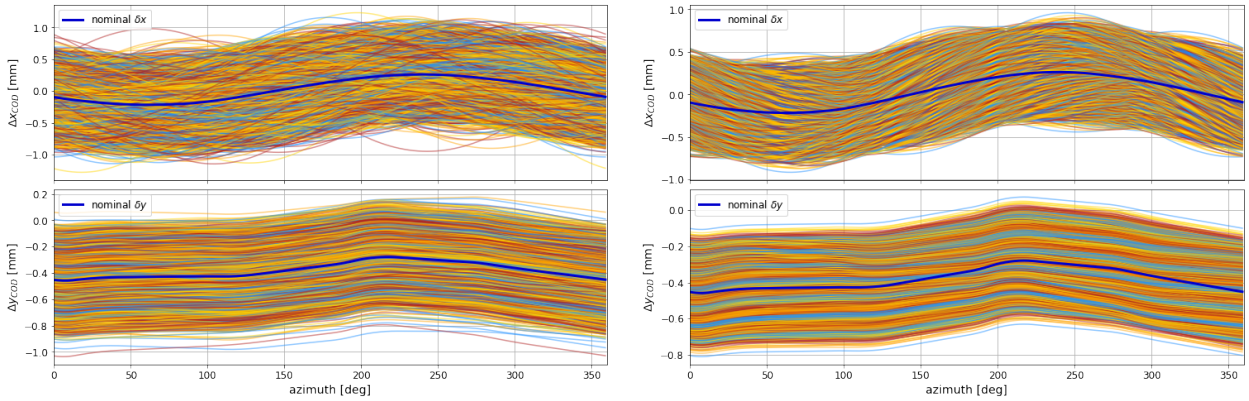
Figure 5.19: CODs from ESQ misalignments. Radial (top) and vertical (bottom) CODs due to ESQ centroid displacement. On the left, the CODs due to the nominal displacements measured during Run-1 are shown. On the right, the CODs due to a unit (1mm) displacement of a single ESQ section are displayed, which can be scaled and summed to obtain the COD of a combination of section misalignments.

nominal displacement values with the respective uncertainty as the distribution's standard deviation. Fig. 5.20a shows all MC sample CODs, plus the COD due to the nominal displacement values. Comparing the spread of the MC samples with the spread of the corner cases (fig. 5.20b) ensures no systematic effects or correlations are overlooked, and confirms the validity of the MC method.

With those $10,000$ MC sample CODs, an equal number of Δk_j and $\Delta\tilde{B}_{tp}$ can be calculated *per trolley pair* and *per transverse direction* (i.e. the radial and vertical CODs are evaluated separately, and produce independent corrections and uncertainties). For each trolley pair, the mean and standard deviation of the $10,000$ $\Delta\tilde{B}_{tp}$ values are taken as that trolley pair's correction and uncertainty due to an ESQ misalignment COD. The corrections and uncertainties per RF-dataset are the CTAG-weighted means of the trolley pair corrections and uncertainties[7]. The results are summarized in table 5.12.

As a cross-check, the entire muon weighting analysis was also performed once with the nominal misalignment COD hard-coded into the beam extrapolation code. Comparing the CTAG-weighted RF-dataset averages of the nominal weighted field with the weighted field including the hardcoded ESQ misalignment COD, the differences are 1.86ppb for the no-RF dataset, 1.02ppb for the x-RF dataset, 1.15ppb for the xyRF5 dataset, and 1.71ppb for the

[7]Since the ESQ misalignment is a hardware property and should be considered unchanged across trolley pairs, the uncertainties are treated as fully correlated.



(a) 10,000 Monte Carlo samples

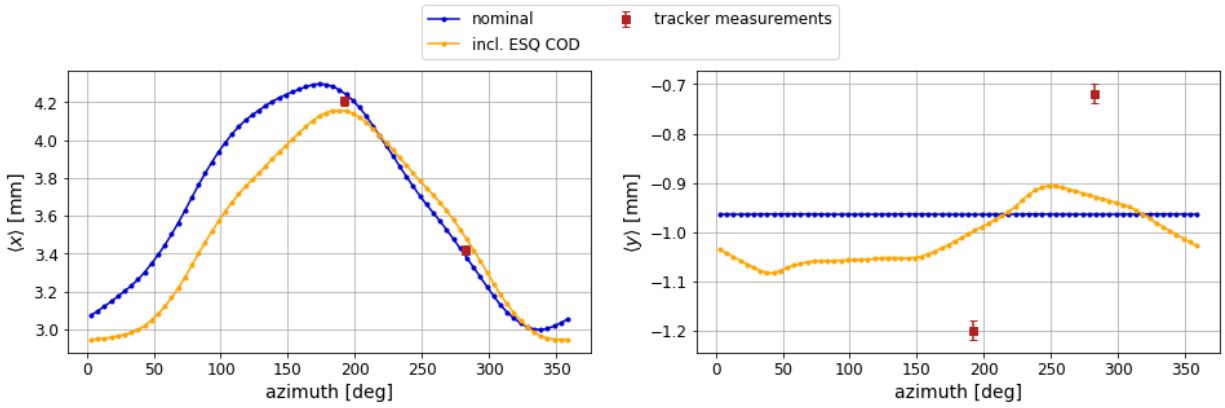
(b) All 6,561 variations of $(0, \pm\sigma)$

Figure 5.20: Monte Carlo CODs for ESQ misalignments. Radial (top) and vertical (bottom) CODs due to ESQ varying centroid displacement. On the left, the 10,000 MC sample CODs are drawn. On the right, the 6,561 combinations of nominal displacements plus $(-\sigma, 0, +\sigma)$ are shown. In both plots, the COD from nominal displacements measured during Run-1 are overlayed in dark blue.

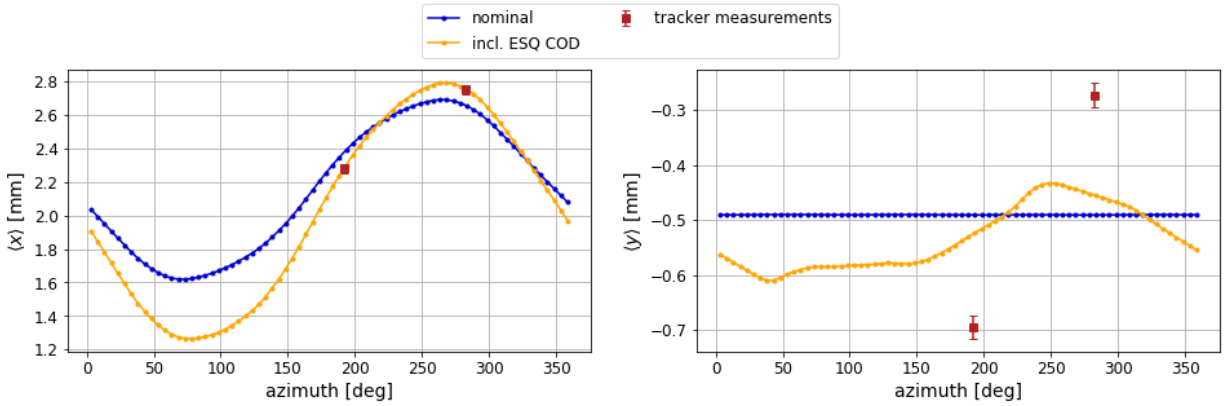
xyRF6 dataset, showing excellent agreement with the corrections reported in table 5.12. Additionally, including the ESQ misalignment COD in the extrapolation step improves the agreement between the measured radial and vertical means at the tracker locations, and the summed extrapolated histograms, for many run groups, as illustrated in fig. 5.21. In particular, the fact that the radial summed station mean now agrees well with the measurements of the individual tracker stations builds confidence in the accuracy of the misalignment values. The remaining discrepancy of the extrapolated vertical mean to the station measurements is fairly consistent across run groups, and about ±0.2mm, which would suggest that if it is due to a radial field COD, the radial field variations causing the COD are stable over time.

	no-RF		**x-RF**		**xyRF5**		**xyRF6**	
	$\Delta\tilde{B}$	$\sigma_{\tilde{B}}$	$\Delta\tilde{B}$	$\sigma_{\tilde{B}}$	$\Delta\tilde{B}$	$\sigma_{\tilde{B}}$	$\Delta\tilde{B}$	$\sigma_{\tilde{B}}$
x_{COD} from ESQ misalignment [ppb]	2.17	5.73	1.27	5.92	1.39	5.98	1.94	6.54
y_{COD} from ESQ misalignment [ppb]	-0.26	0.15	-0.17	0.12	-0.17	0.11	-0.16	0.13
total correction [ppb]	1.9		1.1		1.2		1.8	

Table 5.12: Corrections and uncertainties from CODs due to ESQ misalignment



(a) An example run group during Run-4



(b) An example run group during Run-5

Figure 5.21: Hard-coded ESQ misalignment COD extrapolation. Examples of extrapolation results with ESQ misalignment COD built into the extrapolation code, compared to nominal extrapolation. The summed station means of a randomly chosen run group are shown for both Run-4 (top) and Run-5 (bottom). The ESQ misalignment COD shifts both the radial and vertical mean, such that for most run groups the radial summed station mean now agrees well with the individual station measurements, and the discrepancy of the vertical mean (to the station measurements) is reduced.

5.5.3 Calorimeter acceptance

Like the tracking detectors, the calorimeters' acceptance is also not uniform across the storage volume. In particular, decay positrons originating near the top or bottom of the storage ring are less likely to be detected [118]. The tracker resolution and acceptance corrections are applied so that the beam profiles will represent the true muon distribution in the ring. Even muons that decay close to the center of the storage volume can still have sampled spaces close to the edge during their lifetime due to the betatron oscillations. Finding a consistent correlation between the muon ensemble's decay positions and the preceding muon probability density is a non-trivial problem.
The calorimeter acceptance effect was estimated during Run-1 for one dataset (1d) by deconvolving the tracker profiles and applying an acceptance probability weight to the resulting distribution of betatron amplitudes and equilibrium radii [118]. Calculating beam multipole projections from the acceptance-weighted reconstructed beam profiles resulted in a set of changes to the beam projections $\Delta k_j(\phi_i)$, summarized in table 5.13 (azimuthally averaged).

The beam was better centered in future run periods, and for Run-2/3 these changes $\Delta k_j(\phi_i)$ were used as a rough estimate for the size of the calorimeter acceptance effect. With more muons closer to the center, the calorimeter acceptance effect is expected to cause a smaller variation in the beam profiles and thus a smaller variation in the beam projections, hence the Run-1 values are an over-estimate. The changes $\Delta k_j(\phi_i)$ combined with the Run-2/3 multipole strengths resulted in a ≤ 1 ppb effect [117], which is below the significant level and justifies using an estimate instead of repeating the full calorimeter acceptance correction.
The same strategy is employed in Run-4/5/6, i.e. the beam projection change $\Delta k_j(\phi_i)$ determined in Run-1 is combined with the Run-4/5/6 magnetic field multipoles $m_j(\phi_i)$. For each run group, the binned VTR is multiplied by the beam projection changes, such that the respective change of that run group's muon weighted field is

$$\Delta \tilde{B}_{rg} = \sum_{i=1}^{13} \frac{1}{N_\phi} \sum_{j=1}^{N_\phi} \Delta k_i(\phi_j)\ m_i(\phi_j). \tag{5.50}$$

Then $\Delta \tilde{B}_{rg}$ is taken as the uncertainty due to the calorimeter acceptance for that run group, and the total uncertainty for an RF-dataset is calculated as the CTAG-weighted average of the $\Delta \tilde{B}_{rg}$. Treating the run group uncertainties as fully correlated is valid since the same $\Delta k_i(\phi_j)$ values are used for all run groups. Table 5.13 summarizes the changes to each

multipole component, $\Delta\tilde{B}_j$, for each of the RF-datasets, as well as the total weighted field change per RF-dataset (which is the uncertainty due to the calorimeter acceptance, and is less than 1 ppb for all RF-datasets).

	$\langle\Delta k_i\rangle_\phi$ [10^{-3}]	$\Delta\tilde{B}_i$ [ppb]			
	(Run 1d)	noRF	xRF	xyRF5	xyRF6
Dip	0.0	0.000	0.000	0.000	0.000
NQ	-1.879	0.044	-0.073	-0.136	0.015
SQ	0.002	> -0.001	> -0.001	< 0.001	> -0.001
N6	1.502	0.146	0.101	0.034	0.091
S6	-0.109	0.001	< 0.001	-0.010	-0.002
N8	0.895	0.019	0.015	0.043	0.024
S8	0.128	0.010	0.010	0.010	0.007
N10	0.715	-0.026	-0.027	-0.035	-0.034
S10	0.103	0.003	0.005	0.005	0.002
N12	0.208	-0.165	-0.167	-0.167	-0.167
S12	0.107	-0.013	-0.014	-0.014	-0.015
N14	-0.123	0.127	0.127	0.126	0.126
S14	0.037	0.000	0.000	0.000	0.000
total $\Delta\tilde{B}$		0.147	-0.022	-0.143	0.047

Table 5.13: Calorimeter acceptance uncertainty estimate. Azimuthal averages of beam projection changes $\langle\Delta k_i\rangle_\phi$ from Run-1, and resulting change in weighted field multipole components per RF-dataset $\Delta\tilde{B}_i$, as well as the total change $\Delta\tilde{B}$. Field changes are calculated as the product of azimuthally varying $\Delta k_i(\phi_j)$ and multipole strengths, azimuthally averaged and CTAG-weighted across run groups.

5.5.4 Field-related uncertainties

5.5.4.1 Azimuthal averaging

The calculation of the weighted field per run group, eq. 5.11, is based on the magnetic multipole strengths in 5° azimuthal bins. While it is assumed that the magnetic field is approximately unchanged across each azimuthal bin, small variations do exists, and practi-

cally the multipole strength in the bin with center ϕ_i is calculated as the average of 10,000 points within the bin: $m_j(\phi_i) = \frac{1}{10{,}000} \sum_{n=0}^{9{,}999} m_j(\phi_i - 2.5^\circ + n\frac{5^\circ}{10{,}000})$. As detailed in section 5.1, the values for m_j are definitively known only for the fixed probe locations, and have to be interpolated for spatial points between the probes. This interpolation is done for three different interpolation orders: Linear (cf. eq. 5.18), quadratic, and cubic. The quadratic and cubic interpolations are calculated using standard splines of order 2 and 3 respectively. The value of $m_j(\phi)$ used to calculate the weighted field is the mean of the three interpolation orders, $m_j(\phi_i) = \frac{1}{3}\sum_{n=1}^{3} m_j^{[n]}(\phi_i)$ where $[n]$ denotes the interpolation order used to calculate the azimuthal bin mean. Thus the weighted field in a run group is

$$\tilde{B}_{rg} = \frac{1}{72}\sum_{i=1}^{72}\sum_{j=1}^{13}\frac{1}{3}\sum_{n=1}^{3} m_j^{[n]}(\phi_i)\; k_j(\phi_i) = \frac{1}{3}\sum_{n=1}^{3}\frac{1}{72}\sum_{i=1}^{72}\sum_{j=1}^{13} m_j^{[n]}(\phi_i)\; k_j(\phi_i) = \frac{1}{3}\sum_{n=1}^{3}\tilde{B}_{rg}^{[n]}. \tag{5.51}$$

The uncertainty due to the azimuthal averaging ("spline uncertainty") can then simply be derived from the variance of the weighted field results from the different interpolation orders [116],

$$\sigma^2_{\tilde{B}_{rg}} = \mathrm{Var}(\tilde{B}_{rg}) = \frac{1}{3}\sum_{n=1}^{3}(\tilde{B}_{rg}^{[n]} - \tilde{B}_{rg})^2. \tag{5.52}$$

This method to evaluate a field-related uncertainty implicitly includes the correlations between multipole moments.

The spatial variability of the magnetic field is generally similar between trolley runs, and it can therefore be assumed that the "true" interpolation order (which is unknown) is consistent across run groups. This indicates that the spline uncertainty should be treated as correlated in time, i.e. the uncertainty per RF-dataset will be calculated as a straight-forward CTAG-weighted average of the run group uncertainties, which are shown in table 5.14.

	noRF	xRF	xyRF5	xyRF6
$\sigma_{\langle\tilde{B}\rangle}$ from azimuthal averaging [ppb]	1.73	1.56	2.16	1.55

Table 5.14: Muon weighted field uncertainty due to azimuthal averaging

Interestingly, if this uncertainty is evaluated per multipole moment and combined assuming

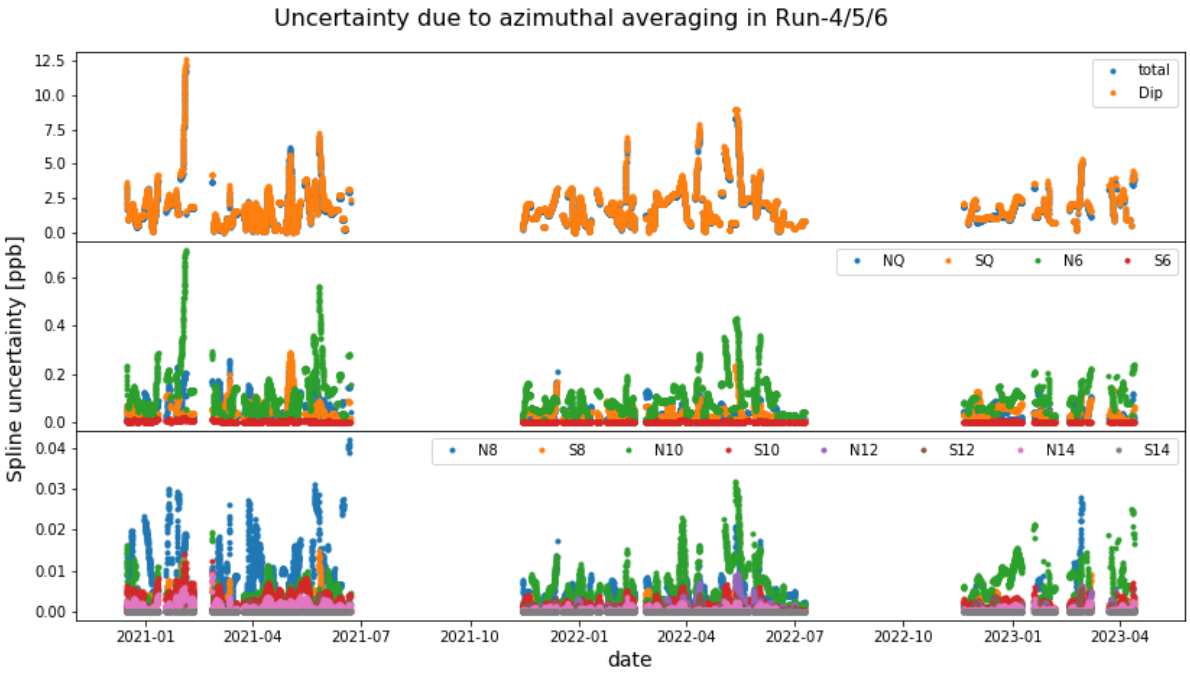


Figure 5.22: Total uncertainty on $\langle \tilde{B} \rangle$ due to averaging in azimuthal bins, and individual multipole contributions. Linear, quadratic, and cubic spatial interpolation is performed, and average within each bin calculated. Result is taken as the mean of the three interpolation orders, uncertainty as the variance. The majority of the total contribution is from the dipole (top); followed by fixed-probe-tracked multipole contributions, especially N6 (middle); multipole orders above sextupole show negligible contributions (bottom).

independence of the moments, i.e.

$$\mathrm{Var}(\tilde{B}_{rg}) = \sum_{j=1}^{13} \mathrm{Var}(\tilde{B}_{rg,j}) = \sum_{j=1}^{13} \frac{1}{3} \sum_{n=1}^{3} (\tilde{B}_{rg,j}^{[n]} - \tilde{B}_{rg,j})^2,$$

where $\tilde{B}_{rg,j}^{[n]} = \frac{1}{72} \sum_{i=1}^{72} m_j^{[n]}(\phi_i)\ k_j(\phi_i)$ and $\tilde{B}_{rg,j} = \frac{1}{3} \sum_{n=1}^{3} \tilde{B}_{rg,j}^{[n]}$, the uncertainties increase slightly to 1.85 ppb (noRF), 1.63 ppb (xRF), 2.29 ppb (xyRF5), and 1.67 ppb (xyRF6). This suggests that the multipoles are negatively correlated. Fig. 5.22 shows the total spline uncertainty per run group across all of Run-4/5/6, as well as each multipole contribution $\sigma_{\tilde{B}_{rg,j}}$ individually. The dipole is the dominant contribution and the only multipole that is consistently > 1 ppb. The next largest component is the N6, which is known to alias into the dipole as well, followed by NQ and SQ. All other multipoles have overall negligible contributions.

5.5.4.2 Other field-related sources of uncertainty

While not part of the muon weighting analysis, other field-related sources of uncertainty will be briefly summarized here for comparison. The azimuthal averaging uncertainty is grouped with other sources of uncertainty related to the spatial field maps and interpolation, e.g. uncertainties of the frequency extraction of trolley and fixed probes, motion effects of the trolley, position determination of the trolley (both in azimuth and transverse position), temperature corrections, or the uncertainties arising from the multipole parameterization, and the truncation error of the parameterization. All of the uncertainties in this category are treated as correlated uncertainties except for the "Brownian bridge" uncertainty. The Brownian bridge model assumes that the sync offset between two trolley runs behaves like a random walk, and the uncertainty on the sync offset can be evaluated using the covariance of a Brownian bridge between the two trolley runs [8]. This uncertainty is treated as independent across trolley runs and propagated as such. The total uncertainty of this category per RF-dataset is on the order of 30 ppb [51], including the contribution of the azimuthal averaging.

Another major category of field-related uncertainty sources is the series of measurements required for the absolute calibration of the trolley probes, which contributes a total of ~ 34 ppb per RF-dataset [59], about the same amount as the entire interpolation analysis.

5.5.5 Summary

Table 5.15 summarizes all uncertainties of the muon weighted field for the RF-datasets in Run-4/5/6. It includes all uncertainties that arise during the muon weighting procedure itself, including the azimuthal averaging of the virtual trolley runs, but does not include the uncertainties on the field itself.

	noRF	**xRF**	**xyRF5**	**xyRF6**
Detector effects				
Tracker acceptance [ppb]	0.10	0.11	0.21	0.04
Tracker resolution [ppb]	0.03	0.08	0.08	0.07
Tracker alignment x [ppb]	0.17	0.41	1.02	0.23
Tracker alignment y [ppb]	0.56	0.10	1.39	0.02
Calo acceptance [ppb]	0.15	0.02	0.14	0.05
Beam dynamics effects				
y_{COD} from B_{rad} [ppb]	2.42	2.05	1.57	2.08
x_{COD} from quad misalignment [ppb]	5.73	5.92	5.98	6.54
y_{COD} from quad misalignment [ppb]	0.15	0.12	0.11	0.13
Mean momentum offset $< \delta >$ [ppb]	≤ 0.01	≤ 0.01	≤ 0.01	≤ 0.01
COSY beam functions [ppb]	≤ 0.01	≤ 0.01	≤ 0.01	≤ 0.01
Field specific				
Azimuthal averaging [ppb]	1.85	1.63	2.29	1.67
total [ppb]	6.2	6.3	6.4	6.9

Table 5.15: Summary of muon weighting uncertainties, reproduced from [69].

CHAPTER 6

Learning-Based Track Reconstruction

The preceding chapter illustrates the importance of accurately measuring the beam shape in the tracking detectors to be able to determine the muon-weighted magnetic field. Additionally, as detailed in section 5.2, the number of reconstructed tracks directly influences the effective time binning in the muon weighting process. The track reconstruction algorithm, which is entirely based on geometric considerations and the numerical integration of analytic equations, has a reconstruction efficiency[1] of about 60% [136]. In particular, "short" tracks (tracks that only register in a few straws) and tracks of positrons that traversed material before entering the tracker are entirely excluded from the reconstruction algorithm. Therefore, investigating alternative or complementary approaches that could provide improved tracking efficiency is a high priority.

Machine learning (ML) algorithms and NNs in particular have been used successfully for track reconstruction in various experiments [68, 73, 74]. Due to their ability to *implicitly* learn environmental features like fringe fields or material effects, NNs are potentially capable of addressing the specific shortcomings of the analytic track reconstruction. In this chapter, a ML-based likelihood-free parameter inference approach will be explored, in which a NN is trained to approximate the likelihood of observing data $\boldsymbol{x}$ (consisting of N data points) given m parameters $\boldsymbol{\theta} = (\theta_1, ..., \theta_m)$ [34],

$$\mathcal{L}(\boldsymbol{\theta}|\boldsymbol{x}) = p(\boldsymbol{x}|\boldsymbol{\theta}) = \prod_{i=1}^{N} p(x_i|\boldsymbol{\theta}), \tag{6.1}$$

[1] The efficiency can be calculated by dividing the number of reconstructed tracks by the number of positrons registered in the calorimeter behind the respective tracker station within a certain time window. The choice of window length only affects the result within the first 30 µs of a fill, when pile-up is affecting the track reconstruction significantly.

or equivalently the log-likelihood (LLH) function

$$\ln \mathcal{L}(\boldsymbol{\theta}|\boldsymbol{x}) = \sum_{i=1}^{N} \ln p(x_i|\boldsymbol{\theta}). \tag{6.2}$$

It is often difficult or impossible to specify $\mathcal{L}$ in analytic form due to e.g. the complexity of the detector response, whereas simulated data is readily available and can be used to train a neural network. The advantage of maximizing an approximated LLH to determine the parameters is that the likelihood function inherently incorporates an uncertainty of the solution through confidence limits [34, 68].

6.1 Analytic track reconstruction algorithm

In order to put the proposed machine learning algorithm in context, as well as to produce appropriate training data, it is important to first understand the current, deterministic track reconstruction algorithm. The layout of a tracking station has been summarized in section 3.2.2 and is illustrated in more detail in fig. 6.1. Using the digitized information from the recorded straw hits, which includes the straw chamber's identifiers (station number, module number, view type[2], layer number, and wire number within the layer) and the time of the detected signal, the fundamental reconstruction steps are [141]

1. Hit pattern recognition - temporal and spatial grouping of hits into "seeds",
2. Track finding - spatial correlation of seeds into "track candidates",
3. Track fitting - positron trajectory parameter fitting within the detector area,
4. Track extrapolation - numerical extrapolation of the trajectory to the decay vertex.

The ML algorithm can combine the third and fourth steps, if the parameters θ are defined at the vertex position. The input for the NN are the track candidates defined in step two, and training data hence has to be created with the same guidelines.

6.1.1 Hit pattern recognition

Initially, hits from all modules of a station are grouped into temporal clusters (so-called "time islands") of no more than 100 µs duration [104]. Then within a time island, hits are grouped into geometrically determined spatial units [141]:

[2] A "view" is one of the two planes (angled double layers of straws) in each module. The plane where the top of the straws are closer to the muon storage volume than the bottom is called the U-view, and the plane where the bottom is closer than the top is called the V-view [104].

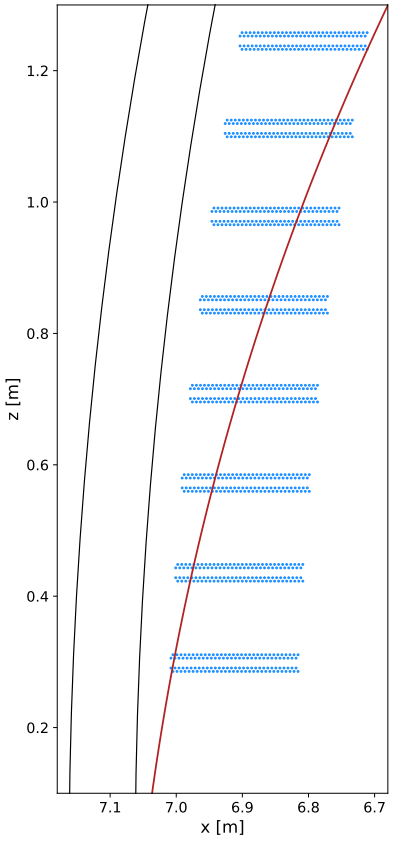
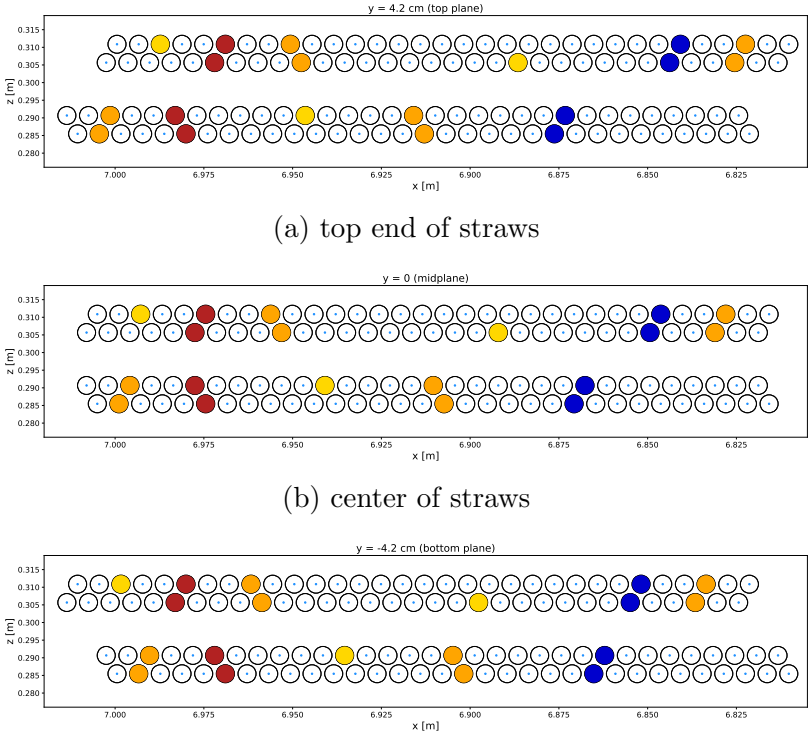


Figure 6.1: Top-down schematic view of tracker station 18. Blue dots indicate the straw locations in the mid-plane ($y = 0$) on the inside of the storage ring (indicated in black). The red line shows an exemplary positron track from a decay vertex outside the image range.



(a) top end of straws

(b) center of straws

(c) bottom end of straws

Figure 6.2: Top-down cross-sectional view of a single tracker module, at the top of the active straw region (top), mid-plane (center), and bottom of the active straw region (bottom). Each black circle is a straw tube, with the wire in the center indicated by a light blue dot. Examples of the spatial groupings necessary for track formation are shown as filled circles: Orange straws form doublet clusters, whereas yellow straws define singlet clusters; the red straws form a seed which likely passed the module close to the mid-plane, while the dark blue straws form a seed passing close to the bottom end of the straws.

- **Clusters** are groups of one ("singlet") or two ("doublet") hits in a view/plane. A doublet cluster consists of a hit in each layer of the plane, which are adjacent to each other. This is equivalent to enforcing that the wire numbers of the hits differ by ≤ 1. If only one layer registers a hit, this hit is defined as a singlet cluster.

- **Seeds** are similarly made up of one or two adjacent clusters, but due to the relative shift between the two views outside of the vertical midway plane, the wire numbers can differ by ≤ 3. Additionally, the hits in the two planes are required to be no more than 50 ns apart.

Fig. 6.2 shows the different spatial units (orange/yellow for examples of doublet/singlet clusters, and two seed examples in red and dark blue) within a module. Three horizontal cross-sections of the module are drawn (top and bottom of the module as well as the $y = 0$ plane), illustrating the distance of clusters within a seed at different heights, depending on the vertical hit position.

6.1.2 Track finding

Track candidates are built from seeds via an iterative process. Starting from individual seeds at either end (first or last module in a station), the algorithm steps through each subsequent module and adds the seed that best matches the existing partial track according to the following general strategy:

- If the track candidate already contains at least two seeds, a fit to a simple circular trajectory ("circle fit") is performed (assuming a uniform magnetic field), and the closest spatially matching seed on the neighboring module is identified.

- If the neighboring module only contains a single seed, it will be merged with the track candidate, provided it lies within the tolerance window of the circle fit (if applicable) and/or the wire number of the seed is moving in the correct direction (depending on which module initiated the track candidate) [141].

- If there are multiple seeds on the neighboring module, the "best match" criterion depends on whether any of them are close to each other in time (less than 15 ns apart) or space (wire number differs by ≤ 4).

  - Only seeds whose wire number is moving in the right direction can be added to the track candidate.

  - If none of the seeds on the neighboring module are close to each other in time, the seed with the shortest time gap to the track candidate is selected.

- If any seeds are close in time but none are close in space, the seed with the smallest wire difference to the track candidate is selected.
- If there are both seeds close in time and seeds close in space to each other, and a circle fit was performed, the seed identified by the fit is added to the track candidate.
- Otherwise (if no circular fit exists), the decisions falls back to the shortest time gap.

Once a track candidate has been assigned all compatible seeds, the hit times of all hits in the track, together with an estimate of the momentum obtained from the circle fit, permit the determination of the time the positron entered the tracking station t_0 [136]. This step is crucial, since the difference between t_0 and a straw's hit time, accounting for time of flight between the detector entrance and the respective straw, is the drift time[3] between the point of ionization within the straw and the central wire[4]. Using the average drift velocity of the Ar:C_2H_6 gas mixture (about 50 mm/µs [136]), the drift time can be converted to a distance between the "true hit" and the straw wire. Due to the geometry of the straws, there is some ambiguity as it pertains to the exact location of the "true hit" within the straw - the drift time/distance only defines a cylinder around the wire, and the point of ionization could lie anywhere on this cylinder. This problem is commonly referred to as the "left-right ambiguity", which is non-trivial to resolve. Tremendous effort has gone into developing a reliable way to determine the left-right index of a hit (the details of which will not be documented here), and even a single incorrect assignment within a track will likely cause the track fitting to fail [136].

After the initial track candidates are built, the following conditions have to be met for a candidate to advance to the track fitting stage [141, 136]:

- Can't share seeds, clusters, or hits with other track candidates.
- Must contain hits in at least five planes in a row.
- Must have at least 12 straw hits.

[3]For comparison, a positron traverses the entire tracking station in only about 3 ns, while drift times are typically in the range of 10-60 ns.

[4]While there exist potentially many points of ionization along the positron trajectory within a single straw, the point closest to the wire will have the shortest drift time and initiate the signal (which is registered as the hit time). Thus from the drift time, the distance to the point of ionization closest to the straw can be calculated.

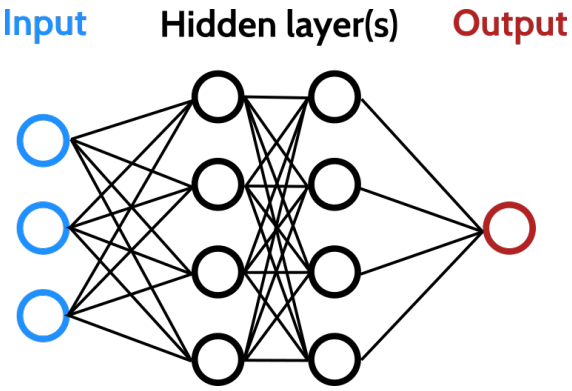


Figure 6.3: Basic structure of a fully connected NN: Each layer consists of a number of neurons which are connected to all neurons in the preceding and following layer. The input layer has one neuron per data feature (e.g. x position, z position, slope, etc.), while the output layer consists of a single neuron making the final prediction. Layers between input and output are called hidden layers.

- Number of hits in U-views and number of hits in V-views can't differ by more than four.
- Must have hits in at least 70% of all layers between start and end of track.
- Each hit's drift time must be no longer than 70 ns.

6.2 A brief neural network primer

Fundamentally, a neural network is a function consisting of a set of nested nonlinear transformations that aims to approximate a (typically complex) relationship between its inputs by gradually adjusting ("learning") its parameters. The concept was developed independently in statistics and artificial intelligence [93], and is generally used for one of two tasks - regression (predicting a continuous value determined from the data) or classification (predicting data to belong to one of a finite number of distinct groups or "classes"). The application relevant to the problem at hand is a binary classifier, i.e. a neural net approximating a function that assigns a probability to data being in one of two classes. The process of finding the optimal parameters ("training") requires a quantity by which to measure prediction accuracy. Then the objective will be to optimize (usually minimize) this quantity, called the loss function L, which involves some form of gradient-based iterative update of the parameters [93, 87]. The simplest such optimization algorithm is gradient descent, which updates the

parameters $\boldsymbol{p} = (p_1, ..., p_n)$ according to

$$\boldsymbol{p}_{i+1} = \boldsymbol{p}_i - \epsilon \nabla_{\boldsymbol{p}} L(\boldsymbol{p}_i), \tag{6.3}$$

where ϵ is the so-called learning rate. However, calculating the gradients for large networks and/or large sets of data can be very computationally expensive. Instead, several more sophisticated optimization algorithms exist that approximate the gradients from smaller sets of samples, as well as, in some cases, have adaptive learning rates [87]. One popular choice is the Adam (adaptive moments) optimizer, which will be relevant in this model.
The basic structure of a classifier NN is illustrated in fig. 6.3: The smallest computational units, called neurons, are organized in stages ("layers") forming the "net". In this case, the layers are fully connected (or "dense"), meaning each layer's neurons are connected to all neurons in the preceding layer[5]. Each neuron independently computes a weighted sum of the inputs x_i it receives [93],

$$z = b + \sum_i w_i x_i, \tag{6.4}$$

where the weights w_i and bias b of each neuron are trainable parameters. To introduce nonlinearity, an activation function a is applied to the weighted sum before the neuron's output y is passed on to the next layer,

$$y = a(z). \tag{6.5}$$

The activation function is chosen per layer (and is often the same for all hidden layers) and should be differentiable almost everywhere [87] (otherwise gradient descent will not be possible). A standard activation function is the rectified linear unit $\mathrm{ReLU}(z) = \max\{0, z\}$, but other choices are possible and must be evaluated by comparing model performance for different activations.
The network architecture, i.e. the number of hidden layers and hidden units per layer, is one the so-called *hyperparameters* (along with e.g. the initial learning rate), which will affect the model performance but are not directly trained. Instead, they are often tuned in a trial-and-error manor by manually comparing the model performance with different choices of the hyperparameters. If the network has too few hidden units to learn the complexity of the features to be approximated, it will not be able to predict well. On the other hand, larger networks require more computing resources and are prone to overfitting [87, 93]. A mitigation strategy that can be employed is weight decay, a type of regularization that adds

[5]Other layer types exist, e.g. convolutional layers, which only connect to a subset of the previous layer's neurons. Only dense layers will be used in this model.

a penalty term to the loss function, proportional to the square of the weights [93]. The decay parameter λ (the proportionality constant) then becomes another hyperparameter of the model. This allows for a larger network architecture to ensure enough neurons to capture the underlying behavior, while simultaneously shrinking the weights of insignificant neurons and thus limiting overfitting. A version of the Adam optimizer with weight decay is implemented in `tensorflow.keras.optimizers.AdamW`, which will be the optimizer of choice in this model.

6.3 Machine learning algorithm for Muon g-2 tracking

The goal is to develop a ML algorithm that can take a track candidate's information - a variable number of hits, each with straw identity information and drift distance (or equivalently drift time), as well as optionally left-right index and vertical position - and obtain the parameters with the highest likelihood to produce this track candidate. A naive attempt to directly train a NN for multivariate regression would require the derivation and implementation of an appropriate customized loss function. Instead, the problem will be solved in two phases with established strategies: First, a NN will be used to approximate the negative LLH function, followed by a multivariate minimization algorithm to find the parameters $\boldsymbol{\theta}$ that maximize the (positive) LLH.

A classifier $d(\boldsymbol{x}, \boldsymbol{\theta})$, trained with binary cross-entropy loss[6], can learn to distinguish between correlated data/parameter pairs, i.e. samples $(\boldsymbol{x}, \boldsymbol{\theta})$ drawn from the joint probability distribution $p(\boldsymbol{x}, \boldsymbol{\theta})$ (class label $y = 1$), and uncorrelated/independent samples, i.e. pairs where the data and parameters are each drawn from their individual probability distributions such that $(\boldsymbol{x}, \boldsymbol{\theta}) \sim p(\boldsymbol{x})p(\boldsymbol{\theta})$ (class label $y = 0$) [68, 94]. In practice, correlated samples are events produced by a MC simulation (or any other statistically representative collection of realistic events), while uncorrelated samples are constructed as random combinations of independent data and parameter values from the same simulation. The optimal classifier,

$$d'(\boldsymbol{x}, \boldsymbol{\theta}) = p(y = 1|\boldsymbol{x}, \boldsymbol{\theta}) = \frac{p(\boldsymbol{x}|\boldsymbol{\theta})}{p(\boldsymbol{x}|\boldsymbol{\theta}) + p(\boldsymbol{x})p(\boldsymbol{\theta})}, \tag{6.6}$$

[6]For a two-class model with class label $y \in 0, 1$ (hence "binary"), the cross-entropy loss function is defined as [93]

$$L(y, d) = -y \ln(d) + (y - 1) \ln(1 - d),$$

where $d = d(\boldsymbol{x}, \boldsymbol{\theta}) \in [0, 1]$ is the prediction of the classifier.

relates to the likelihood-to-evidence ratio $r(\boldsymbol{x}|\boldsymbol{\theta})$ via [94]

$$\frac{d'(\boldsymbol{x},\boldsymbol{\theta})}{1-d'(\boldsymbol{x},\boldsymbol{\theta})} = \frac{p(\boldsymbol{x}|\boldsymbol{\theta})}{p(\boldsymbol{x})} = \frac{\mathcal{L}(\boldsymbol{\theta}|\boldsymbol{x})}{p(\boldsymbol{x})} = r(\boldsymbol{x}|\boldsymbol{\theta}), \tag{6.7}$$

or equivalently

$$d'(\boldsymbol{x},\boldsymbol{\theta}) = \frac{r(\boldsymbol{x}|\boldsymbol{\theta})}{1+r(\boldsymbol{x}|\boldsymbol{\theta})}. \tag{6.8}$$

Therefore, if the classifier's output layer uses a sigmoid activation function,

$$s(z) = \frac{1}{1+e^{-z}} = \frac{e^z}{1+e^z}, \tag{6.9}$$

the LLH can be extracted from the trained model by replacing the final sigmoid function with a linear activation ($l(z) = z$), since

$$d'(\boldsymbol{x},\boldsymbol{\theta}) = s(z) = \frac{e^z}{1+e^z} = \frac{r(\boldsymbol{x}|\boldsymbol{\theta})}{1+r(\boldsymbol{x}|\boldsymbol{\theta})} \quad \Leftrightarrow \quad z = \ln r(\boldsymbol{x}|\boldsymbol{\theta}). \tag{6.10}$$

For a given track candidate, $\boldsymbol{x}$ is fixed, and thus maximizing the LLH is equivalent to maximizing

$$\ln r(\boldsymbol{x}|\boldsymbol{\theta}) = \ln \mathcal{L}(\boldsymbol{\theta}|\boldsymbol{x}) - \ln p(\boldsymbol{x}). \tag{6.11}$$

Optimization algorithms are included in many standard computation packages, e.g. `scipy.optimize.minimize`, and can be used out-of-the-box. This entire workflow is summarized in fig. 6.4.

6.4 Toy model

To assess whether the proposed ML model can successfully approximate the LLH function, a simplified toy model[7] of a single tracker station (station 18) will be used. Since the extrapolation to the vertex positions by definition stops at the point of radial tangency (i.e. the point inside the storage ring where the positron momentum is exactly tangential to the ring curvature) [141], the training data will consist solely of positrons emitted with no radial momentum component. Additionally, an idealized environment with no material effects (i.e. all positrons have an unobstructed path between the decay vertex and the tracker), no background in the detectors (and thus no hits aside from real tracks), and significantly simplified magnetic fields is used. Training data is produced with a Monte Carlo simulation,

[7] A toy model is an idealized representation of a given setup or system, which allows for the study of specific effects with simple/small-scale simulations.

sampling from the following distributions for the track parameters:

- Positron momentum magnitude p between 1.7 GeV (a typical threshold energy for some ω_a analyses [11]) and 3.0 GeV, sampled from the lab frame number distribution shown in fig. 3.6b. The upper limit is chosen for practical reasons, as positron tracks with higher momenta are barely bent more than the stable orbit.
- Radial vertex position R between $R_0 \pm r_0 = 7.112 \pm 0.045$ m, drawn from a Gaussian

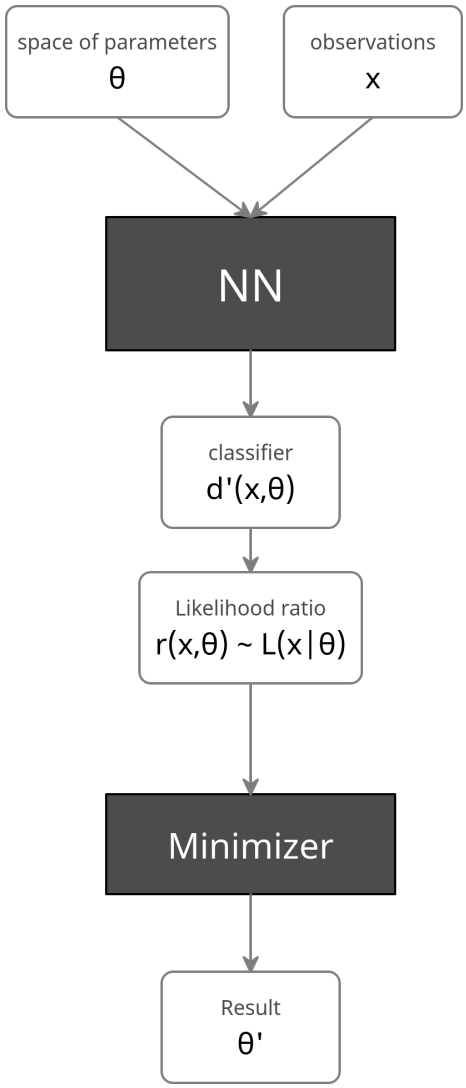


Figure 6.4: LLH estimator flow chart. Flow chart of the ML algorithm to extrapolate a track candidate. A NN classifier is trained to estimate the (negative) likelihood ratio, requiring both the observed track ($\boldsymbol{x}$) and proposed parameters ($\boldsymbol{\theta}$) as input. For a given track, a minimization algorithm can then determine the most likely set of parameters.

distribution centered at R_0 with $\sigma_R = 18\,\text{mm}$ as an approximation of a typical measured beam shape.

- Azimuthal vertex position φ measured from the laboratory x-axis in the *upstream* direction (i.e. $\varphi = 0$ will be close to the tracking station entrance, and $\varphi = 90°$ is a quarter turn in front of the station), also sampled from a uniform distribution.
- Only tracks that hit a minimum of five modules are retained for the NN training.

One particular challenge for the NN is the fact that tracks consist of a variable number of hits. The number of input nodes for the NN must be fixed, which suggests two possible avenues for mitigation: Either determine the LLH for each hit independently and perform the sum over all hits in a track manually outside of the NN; or parameterize the "no hit" state such that a track can be padded to a fixed maximal length. Initial studies of the simple model detailed below show that the second option does not produce an accurate approximation of the likelihood function, in particular for the radial vertex position (which is the most important of the parameters, alongside the vertical position). However, not registering a hit in a detector layer *can* be useful information about the track as well, if the positron passed between straws. Thus the chosen strategy here will be to pad the track to a total length of 32 hits (since there can be no more than one hit per detector layer) by using stand-in values outside the physical range for no-hit layers, yet train the model on individual hits. This means that one of the things the NN has to learn implicitly, is the detection efficiency of a straw.

6.4.1 Idealized detector setup

The setup of the first toy model assumes

- a uniform vertical magnetic field across the entire tracking station (equal to the field inside the storage ring B_0),
- positrons are only emitted perfectly horizontal in the $y = 0$ plane, i.e. the positron momentum has no vertical component and the vertex is located at $y = 0$, and
- each tracker layer is a single, one-dimensional detector with perfect efficiency, which registers the x position of a hit with $\sigma_x = 150\,\mu\text{m}$ (the same precision as the drift distance in the straw trackers [136]).

In this model, the input parameters that uniquely define a track are its momentum magnitude p, the radial vertex position R, and the azimuthal vertex position φ as described above.

6.4.2 Analytic likelihood

Since the detectors in this setup have perfect detection efficiency and are 1-dimensional (meaning the z position of a hit is always exactly the detector position), the only deviation of a registered hit from its true, deterministic position is caused by the uncertainty of the measured x position on the detector. Therefore the likelihood function can be written analytically as

$$\mathcal{L}(\boldsymbol{x}_{\text{meas}}, \boldsymbol{\theta}) = I(\boldsymbol{x}_{\text{meas}}, \boldsymbol{\theta}) \; \Pi_{i=1}^{32} \; \frac{1}{\sqrt{2\pi}\sigma_x} \exp(-\frac{(x_{i,\text{meas}} - x_{i,\text{exp}}(\boldsymbol{\theta}))^2}{2\sigma_x^2}), \tag{6.12}$$

or equivalently as the log-likelihood

$$\ln \mathcal{L}(\boldsymbol{x}_{\text{meas}}, \boldsymbol{\theta}) = \ln I(\boldsymbol{x}_{\text{meas}}, \boldsymbol{\theta}) - 32 \ln(\sqrt{2\pi}\sigma_x) - \sum_{i=1}^{32} \frac{(x_{i,\text{meas}} - x_{i,\text{exp}}(\boldsymbol{\theta}))^2}{2\sigma_x^2}, \tag{6.13}$$

where $x_{i,\text{meas}}$ is the measured position, $x_{i,\text{exp}}$ is the expected/true position, and $I(\boldsymbol{x}_{\text{meas}}, \boldsymbol{\theta})$ is an indicator function defined as

$$I(\boldsymbol{x}_{\text{meas}}, \boldsymbol{\theta}) = \begin{cases} 1 & \text{if measured hit pattern matches expected hit pattern,} \\ 0 & \text{if a hit is observed in a layer with no expected hit, or vice versa.} \end{cases}$$

In the background-free environment of this toy model, there is no possibility of observing a hit in a detector which the track does not intersect, and the perfect efficiency means all intersected detector layers will register a hit. Thus the likelihood of a mismatch in the observed hit pattern (which detectors measure a hit and which ones don't) and the expected hit pattern based on the true parameters $\boldsymbol{\theta}$ is zero.

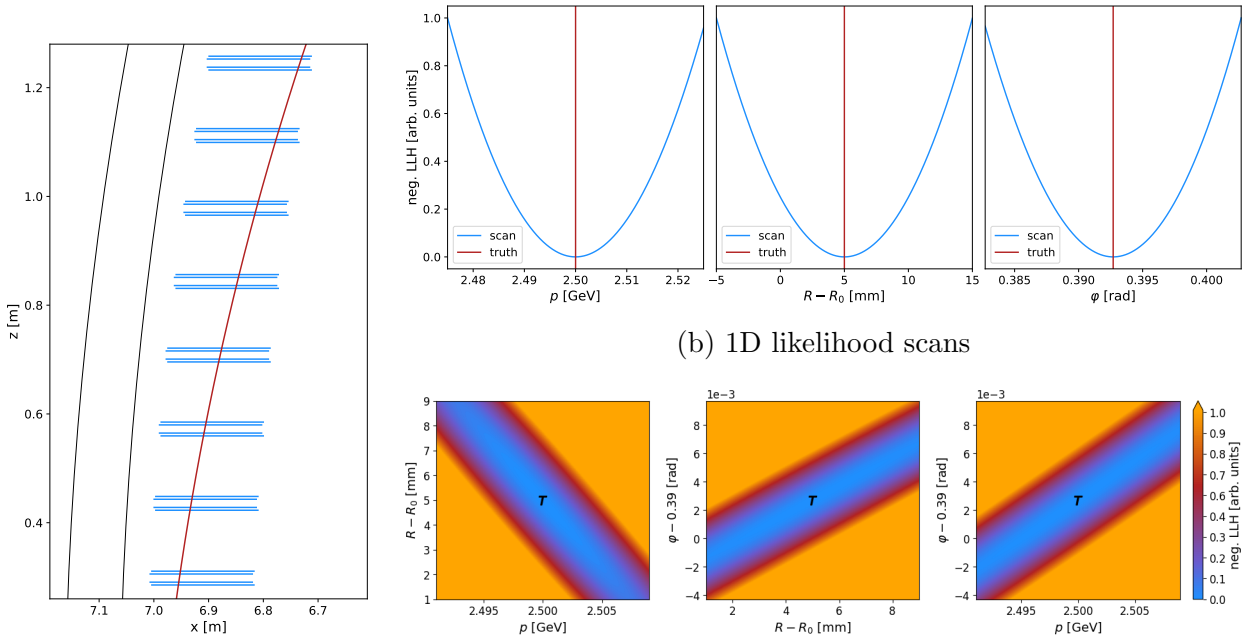
Three exemplary sets of track parameters are summarized in table 6.1 and illustrated in figs. 6.5-6.7. Event A is a mid-range momentum track that hits all detector layers, whereas events B and C are on the low and high end of the momentum spectrum, respectively, and each miss some layers. The figures show each event's trajectory, as well as 1-dimensional and 2-dimensional scans of the analytic negative LLH function. Only finite values are plotted, which means only parameter ranges where $I(\boldsymbol{x}_{\text{meas}}, \boldsymbol{\theta}) \neq 0$ are visible[8]. A notable change in the shape of both the 1D and 2D scans is visible between the different events. Event C in particular has much narrower ranges of valid likelihood values, and larger gradients near the minima, due to the smaller parameter range that can feasibly

[8]To be specific, $I(\boldsymbol{x}_{\text{meas}}, \boldsymbol{\theta}) = 0$ ranges are white in the 2D scans, and in the 1D scan the line graph is only drawn where $I(\boldsymbol{x}_{\text{meas}}, \boldsymbol{\theta}) = 1$.

create a track with more missed detector layers.

	p [GeV]	$R - R_0$ [mm]	φ [rad]
Event A	2.5	+5	$\pi/8$
Event B	1.9	−2	$\pi/14$
Event C	2.8	−7	$\pi/7$

Table 6.1: Simple toy model example events. True track parameters of three example events in the simplest toy model, shown in figs. 6.5-6.7.



(a) Trajectory

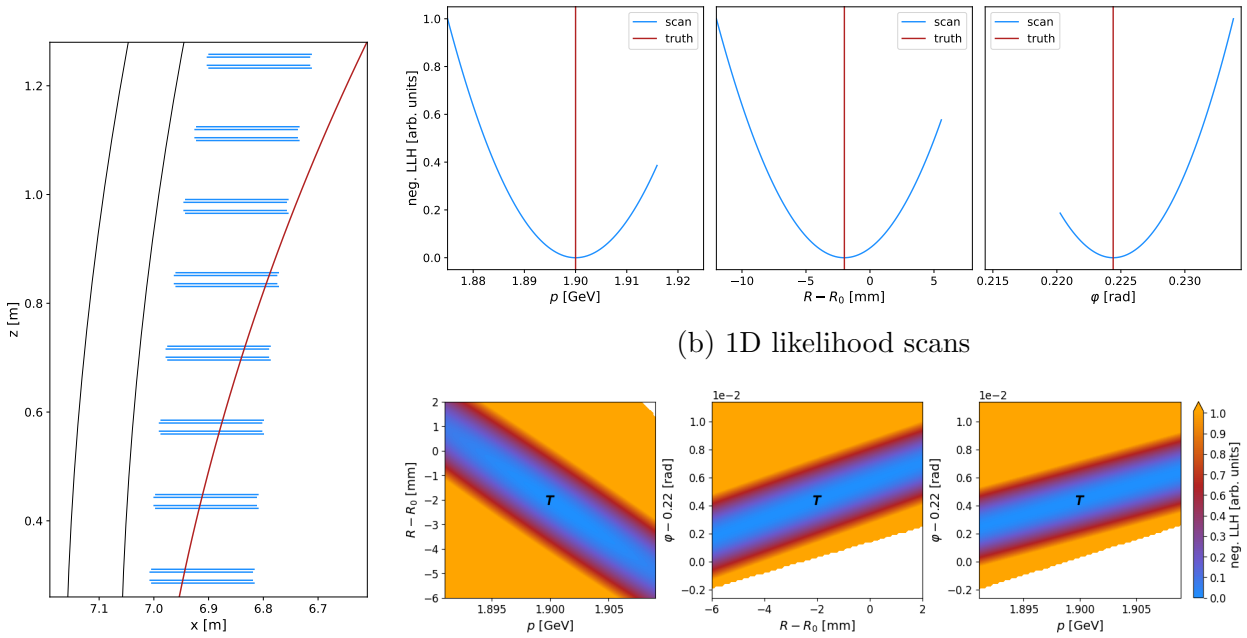
(b) 1D likelihood scans

(c) 2D likelihood scans

Figure 6.5: Analytic likelihood for toy event A: (a) Trajectory through the idealized tracking detectors, (b) 1-dimensional scans of the analytic (negative) log-likelihood, where the two non-varying parameters are set to their true values, and (c) 2-dimensional scans of the analytic (negative) log-likelihood, with the third parameter fixed to its true value.

6.4.3 Neural net

From the measured x positions and the known z positions of the detectors, some helpful derived quantities can be constructed, namely estimates of the slope and curvature of the track at a given detector layer, and an indicator function h that signals whether there



(a) Trajectory

(b) 1D likelihood scans

(c) 2D likelihood scans

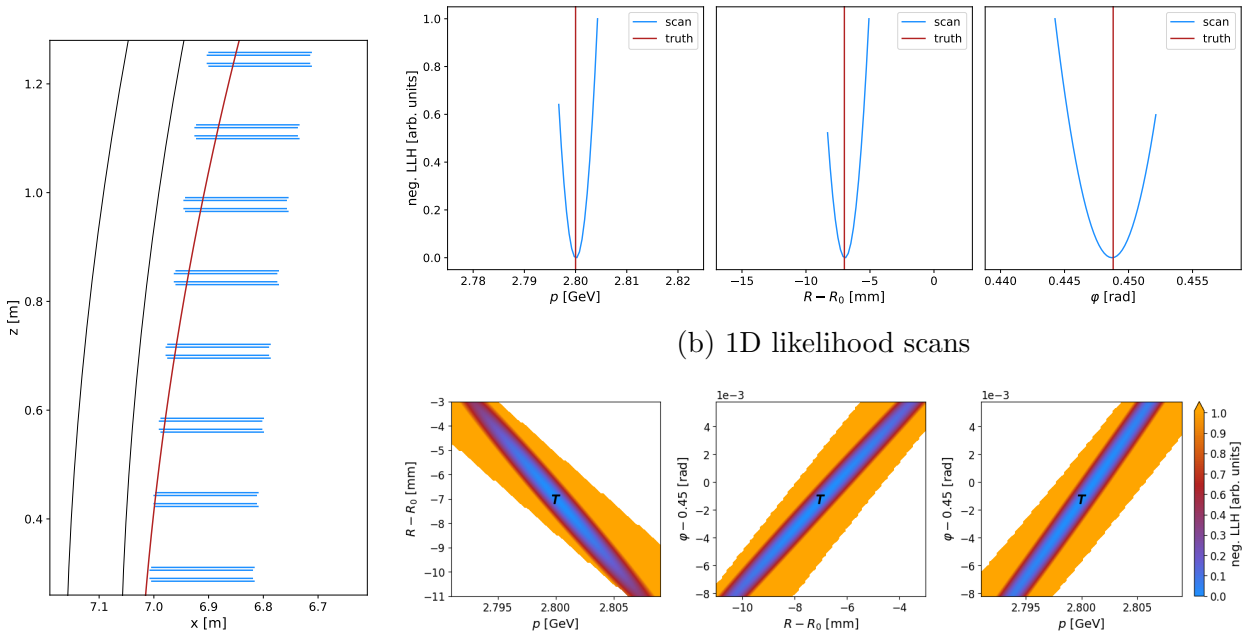
Figure 6.6: Analytic likelihood for toy event B: (a) Trajectory through the idealized tracking detectors, (b) 1-dimensional scans of the analytic (negative) log-likelihood, where the two non-varying parameters are set to their true values, and (c) 2-dimensional scans of the analytic (negative) log-likelihood, with the third parameter fixed to its true value.

was a hit registered in the layer[9]. Before splitting the set of 32 measurements that make up a track into individual hits, the slope s is approximated as the central finite difference (forward/backward difference for the first/last hit of the track), and the curvature k is calculated in the same way from the estimated slope. Since the trajectory is a curved path, it is also helpful to supply the NN with the squares of the x and z positions, as well as the distance from the origin ($R_i = \sqrt{x_i^2 + z_i^2}$)[10]. Moreover, the radial distance from the ideal orbit, $R - R_0$, will be added to the parameter input.

Table 6.2 describes the network architecture of the NN classifier for this toy model. The choice of implementing two separate input layers for the parameters $\boldsymbol{\theta} = (p, R, R - R_0, \varphi)$ and the hit data (of detector layer i) $\boldsymbol{x}_i = (x_i, x_i^2, z_i, z_i^2, R_i, s_i, k_i, h_i)$ was made for clarity rather than functionality, since they are simply concatenated next. The main body of

[9] A dedicated binary input for hit vs. no hit instead of just a non-physical x value allows the NN to explicitly include an "on/off" factor, rather than having to rely on learning weights that can handle both valid and invalid input.

[10] While the NN can learn features like squared distances, it will often train more easily if these are provided with the input.



(a) Trajectory

(b) 1D likelihood scans

(c) 2D likelihood scans

Figure 6.7: Analytic likelihood for toy event C: (a) Trajectory through the idealized tracking detectors, (b) 1-dimensional scans of the analytic (negative) log-likelihood, where the two non-varying parameters are set to their true values, and (c) 2-dimensional scans of the analytic (negative) log-likelihood, with the third parameter fixed to its true value.

the NN consists of three fully connected hidden layers with a Leaky ReLU activation [87] (`tensorflow.keras.layers.LeakyReLU`)

$$\text{L-ReLU}(x) = \begin{cases} x & \text{for } x \geq 0, \\ \alpha x & \text{for } x < 0 \end{cases} \tag{6.14}$$

followed by a dense output layer with sigmoid activation. The advantage of the Leaky ReLU is that all neurons are able to be trained via gradient-based methods, even if their output is negative. The Adam optimizer with weight decay, `tensorflow.keras.optimizers.AdamW`, is used with an initial learning rate $\epsilon = 0.001$ and a weight decay parameter $\lambda = 0.01$. $100,000$ MC sample events are created, resulting in a total of $3,200,000$ individual hit samples, which are split 80:20 into a training dataset and a validation dataset[11]. In addition to the binary cross-entropy loss function, the binary accuracy (which is just the percentage

[11]This is a common practice when training NNs. The training dataset is used to calculate the gradients and update the parameters, while the validation dataset is only used after each epoch to evaluate the loss function and any other metrics. Monitoring the validation loss can indicate when the model starts overfitting and inform when to end training.

of correctly classified samples) is recorded during training as well. The evolution of both of these metrics over the training period are shown in fig. 6.8. Accuracy levels out at about 95%, and both the loss and accuracy show no divergence between training and validation data in the later epochs, indicating a good fit. Once the classifier is sufficiently trained, the output activation is changed to a linear/identity function for LLH prediction [68].

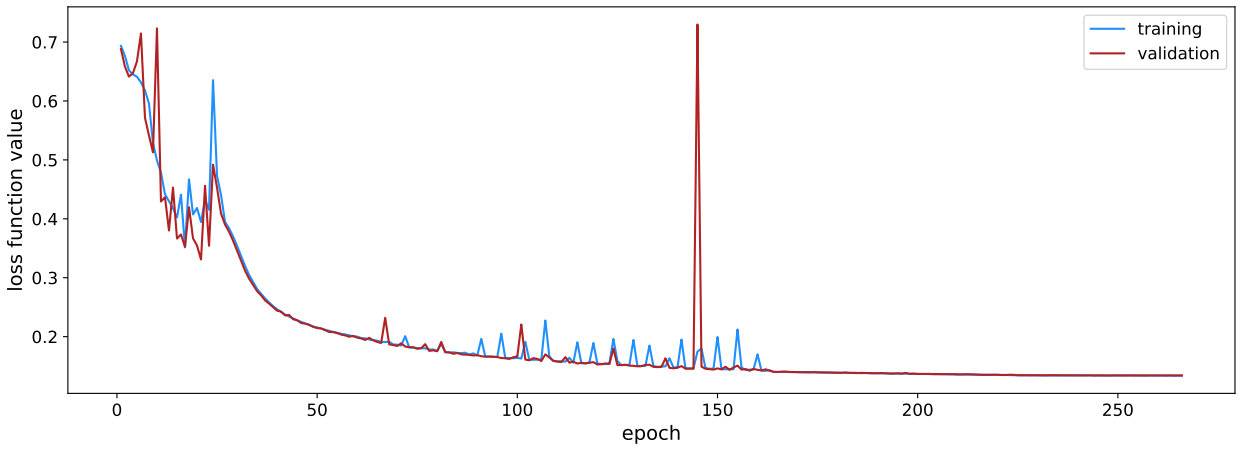
Layer type	**# Neurons**	**# Parameters**	**Activation**
Input $\boldsymbol{\theta}$	4	0	none
Input $\boldsymbol{x}_i$	8	0	none
Concatenate	12	0	none
Dense	32	416	Leaky ReLU $\alpha = 0.3$
Dense	64	2112	Leaky ReLU $\alpha = 0.3$
Dense	128	8320	Leaky ReLU $\alpha = 0.3$
Dense	64	8256	Leaky ReLU $\alpha = 0.3$
Dense	32	2080	Leaky ReLU $\alpha = 0.3$
Output (dense)	1	33	sigmoid (eq. 6.9)

Table 6.2: Network architecture of the simple toy model classifier. The total number of trainable parameters is $21,217$.

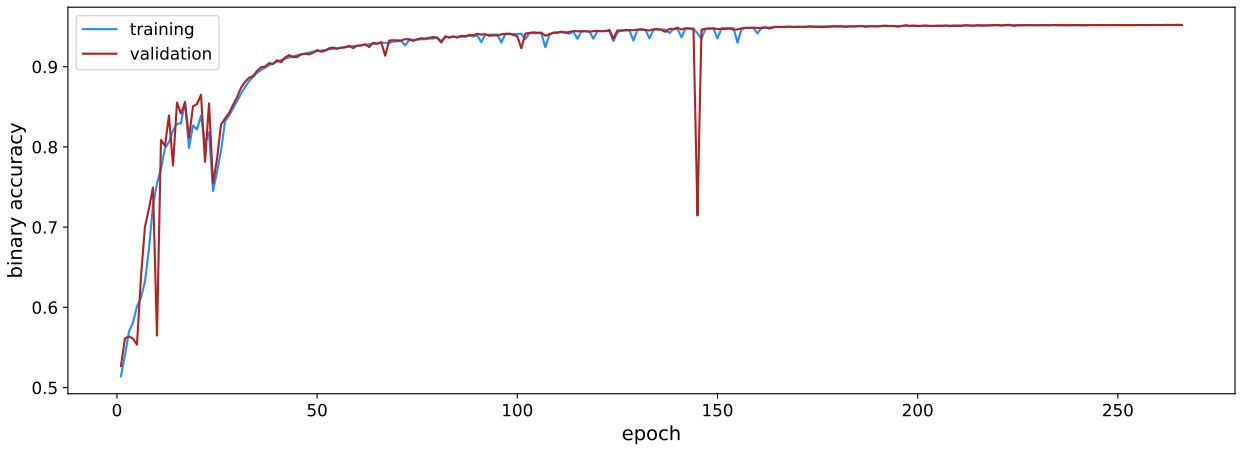
Figs. 6.9-6.11 show scans of the approximated negative LLH function similarly to figs. 6.5-6.7. Note that the NN cannot learn infinite weights, therefore the output will be finite even for "impossible" events for which the analytic likelihood is zero. The 1D likelihood scans now show the full range of possible parameter values, whereas the analytic 1D scans showed only a subset to avoid most of the invalid parameter ranges. The 2D scan ranges are the same as for the analytic scans however. Commonly seen in all scans is a broadening of the likelihood distributions and shallower minima in particular for the radial vertex position, as well as small shifts in the 2D minimum regions. The general shape of the LLH function seems to be well captured, although the definitive test of the accuracy of the approximation has to be tested by comparing minimization results, which will be presented in the following section.

6.4.4 Minimization

For the minimization step, `scipy.optimize.minimize(method='Nelder-Mead')` is used. The initial guess for the parameters is determined, analogously to the deterministic track

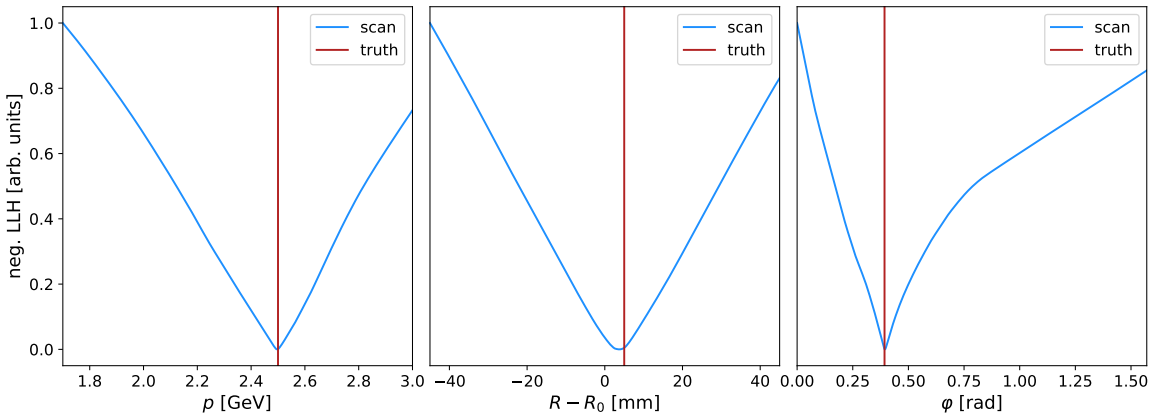


(a) Loss (binary cross-entropy)

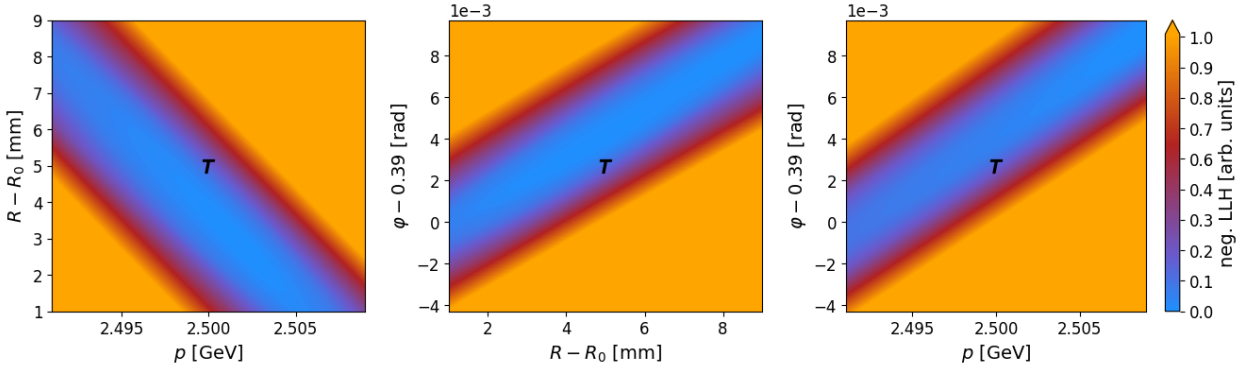


(b) Binary accuracy

Figure 6.8: Loss and accuracy during NN training. Development of quality measures during NN training. The top shows the loss function value per epoch for training and validation dataset. The fact that the validation loss is consistently about the same as the training loss suggests likely no overfitting. The bottom shows the binary accuracy per epoch, which is a more intuitive quantity to understand. Again, training and validation accuracy agree extremely well and level out at about 95%.

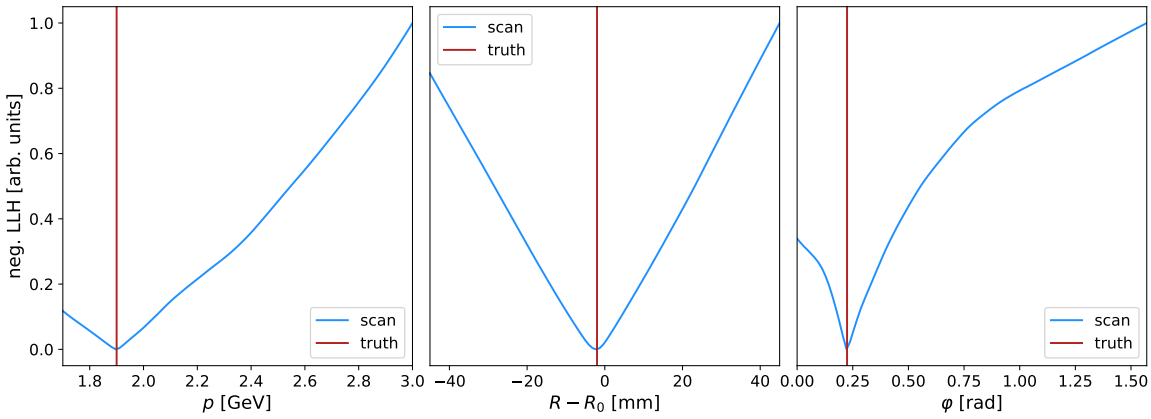


(a) 1D likelihood scans

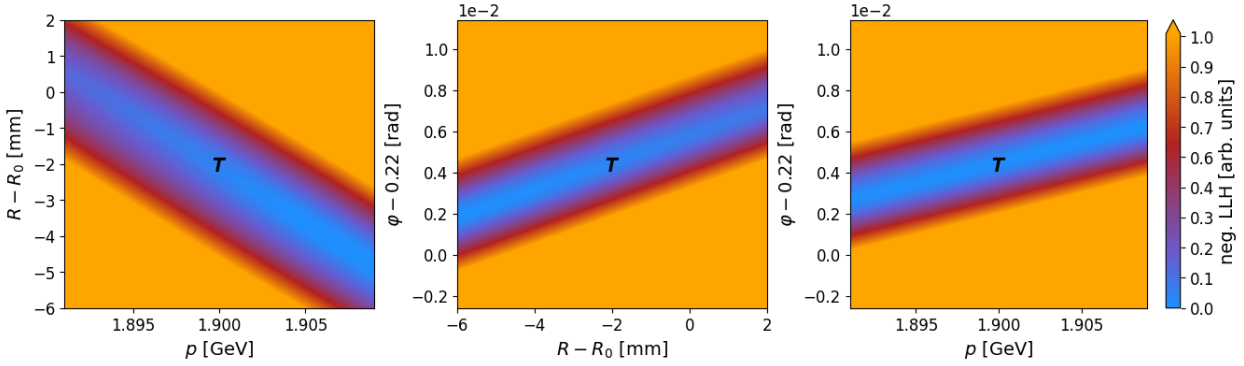


(b) 2D likelihood scans

Figure 6.9: NN approximated likelihood for toy event A: (a) 1-dimensional scans of the approximated (negative) log-likelihood, where the two non-varying parameters are set to their true values, and (b) 2-dimensional scans of the approximated (negative) log-likelihood, with the third parameter fixed to its true value.

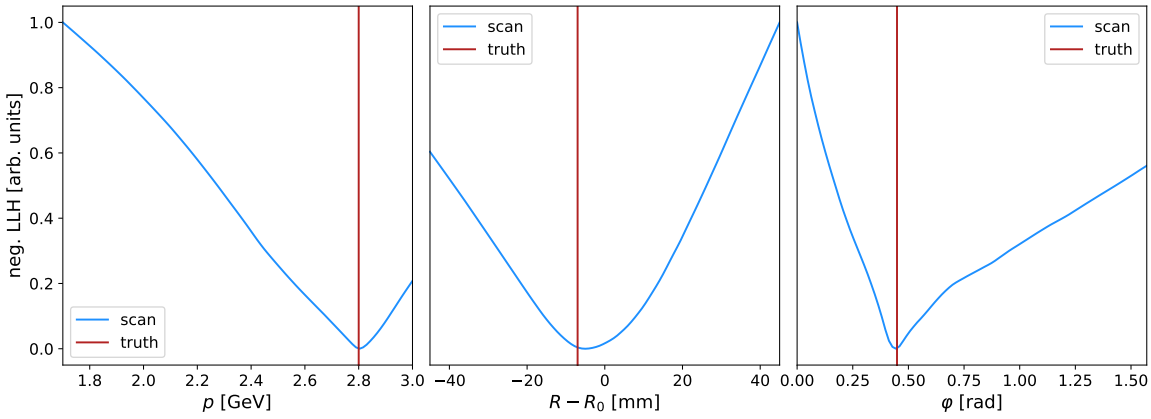


(a) 1D likelihood scans

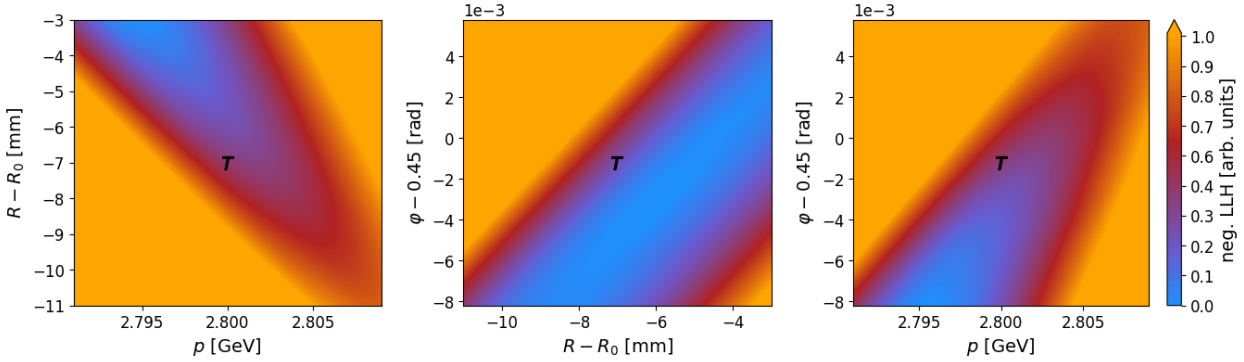


(b) 2D likelihood scans

Figure 6.10: NN approximated likelihood for toy event B: (a) 1-dimensional scans of the approximated (negative) log-likelihood, where the two non-varying parameters are set to their true values, and (b) 2-dimensional scans of the approximated (negative) log-likelihood, with the third parameter fixed to its true value.



(a) 1D likelihood scans



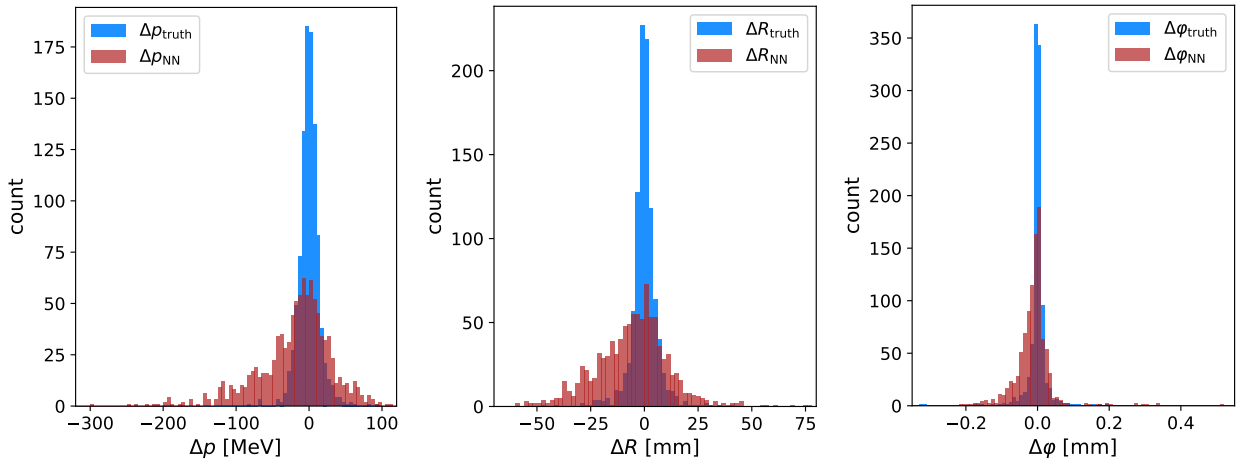
(b) 2D likelihood scans

Figure 6.11: NN approximated likelihood for toy event C: (a) 1-dimensional scans of the approximated (negative) log-likelihood, where the two non-varying parameters are set to their true values, and (b) 2-dimensional scans of the approximated (negative) log-likelihood, with the third parameter fixed to its true value.

finding algorithm, via a circle fit of all hits in the track. Given that this toy model uses a uniform magnetic field, this guess should be very close to the truth already. The minimization is performed for the neural net's approximated LLH, as well as the analytic LLH function. The results for the three example events are summarized in table 6.3, where Δ values with subscript "truth" indicate the difference between analytic LLH minimzation results and the true initial parameters, and Δ values with subscript "NN" indicate the difference between the analytic LLH minimization and the NN minimization. Notably, minima of the *analytic* likelihood do not necessarily recover the true initial parameters. The underlying assumption of the LLH minimization approach to track reconstruction is that the parameters that minimize the negative likelihood are the "correct" ones. Standard fits typically aim to minimize the mean-squared error, which is equivalent to maximizing the likelihood if the measurement errors are independent and Gaussian [34], thus it is expected but has not been confirmed whether the deterministic track reconstruction algorithm obtains similar results. Nonetheless, the quantity that determines whether the NN is approximating the LLH well is the distance from the analytic LLH min, not from the truth, as that is the quantity it is being trained to estimate. Based on the three example events, the NN's reconstruction accuracy is within acceptable bounds for Event A (hitting all detector layers), but not for the other two events. A second MC set with 1000 samples was used to evaluate the reconstruction performance on a more statistically relevant level, presented in fig. 6.12, which shows that while the NN tends to have minimum values close to the analytic LLH minima (as demonstrated by the modes of the distributions being approximately zero), the distributions are relatively flat with long tails. The radial vertex position in particular is expected to have an uncertainty of less than 6 mm [136] for a valid track reconstruction, which is unfortunately not compatible with the distribution of ΔR_{NN} shown in fig. 6.12b.

6.5 Towards a full-scale model

The previous sections have shown that while the toy model NN in the current state does not approximate the LLH well enough to reconstruct the track parameters to the required precision, it does learn the general features and overall shape, and in particular is able to reconstruct tracks with some missing detector layers. It is likely that further tuning of the network architecture, hyperparameters, and derived inputs could produce reasonable improvements in the reconstruction accuracy. If further studies are of interest for future work, the following list provides suggestions and considerations for the next steps towards a full-scale model and potential production-level implementation (roughly in order of proposed execution):



(a) Momentum (b) Radial vertex position (c) Azimuthal vertex position

Figure 6.12: NN minimization results: Distributions of track parameters from minimization of the analytic LLH and its NN approximation based on 1000 random samples. Subscript "truth" indicates the difference between analytic LLH minimization and the true initial parameters, and shows that while very close in most cases, the minimization of the negative likelihood does not inherently guarantee to recover the initial parameters. Subscript "NN" is the difference between the two minimization results. Based on the location of the distribution peaks, the NN is generally learning the correct function, but the width of the tails show it is not very precise in the current implementation.

	Event A	Event B	Event C
$p_{\min}$ (analytic) [GeV]	2.494	1.902	2.790
$p_{\min}$ (NN) [GeV]	2.497	1.807	2.776
$\lvert\Delta p_{\text{truth}}\rvert$ (analytic) [MeV]	6	2	10
$\lvert\Delta p_{\text{NN}}\rvert$ (NN vs. analytic) [MeV]	3	95	14
$R_{\min} - R_0$ (analytic) [mm]	3	-2	-12
$R_{\min} - R_0$ (NN) [mm]	6	-19	-20
$\lvert\Delta R_{\text{truth}}\rvert$ (analytic) [mm]	2	0	5
$\lvert\Delta R_{\text{NN}}\rvert$ (NN vs. analytic) [mm]	3	17	8
$\varphi_{\min}$ (analytic) [rad]	0.387	0.225	0.428
$\varphi_{\min}$ (NN) [rad]	0.394	0.195	0.394
$\lvert\Delta\varphi_{\text{truth}}\rvert$ (analytic) [rad]	0.006	0.001	0.021
$\lvert\Delta\varphi_{\text{NN}}\rvert$ (NN vs. analytic) [rad]	0.007	0.030	0.034

Table 6.3: Minimization results for analytic and NN LLH of toy events: Results of the minimization of the negative LLH, both analytic and NN approximated. Additionally, the absolute difference between the analytic minimum and the true initial parameters ($|\Delta p_{\text{truth}}|$, $|\Delta R_{\text{truth}}|$, and $|\Delta\varphi_{\text{truth}}|$), and the absolute difference between minima of the analytic and approximated LLHs ($|\Delta p_{\text{NN}}|$, $|\Delta R_{\text{NN}}|$, and $|\Delta\varphi_{\text{NN}}|$) are shown.

- Replace 1D detectors with 2D straws, i.e. continue to work only in the horizontal plane, but implement the true straw geometry in the cross-section (which includes gaps between detectors). Even with maintaining 100% detector efficiency, this will lead to occasional intermediate "missing" layers as tracks pass between straws. Additionally, instead of having a specified x and z position of the measured hit, the drift distance must not be an input in addition to the straw wire's x and z position.

- With the inclusion of the drift distance, studies regarding the importance of left/right information for the NN must be performed, weighing potential improvements in performance when included against the negative consequences if a left/right label is assigned incorrectly.

- Implement (more) realistic magnetic fields, in particular varying vertical field strength along the tracker station, but also a radial field component. Since simulations of a trajectory in a continuously changing magnetic field will have to move to numerical

integration to determine the path rather than calculating positions from an analytic equation, a first step towards this could consist of piecewise constant magnetic fields instead.

- Include the vertical dimension. This will require an additional track parameter, e.g. the angle of the positron momentum to the horizontal plane, or simply the vertical momentum component (while the momentum component in the $x - z$ plane remains tangential to the ideal muon orbit). An additional consideration is whether to include a z position for the hit - for seeds that contain hits in both views, this can be estimated, but for seeds with only one cluster there would be no z value, or it would have to have a very large uncertainty, and the uncertainty would have to be supplied to the NN as well. It is possible that implicit learning of vertical information just from the wire combinations that form the track is sufficient.

- Move to less than 100% detector efficiency and/or include background hits. This stage will have to be strongly informed by the real detector performance, as well as the track finding step, which can apply some quality cuts already.

- Eventually, a realistic simulation of tracker data will have to be used, which includes true distributions of the initial parameters as well as detector acceptance and resolution effects, material effects for trajectories that intersect other components of the experiment, etc. Any NN is limited by the quality of the data it is trained on, thus the more realistic the simulated data can be, the better the model will be able to extrapolate measured tracks.

- Potentially investigate improvements of the minimization step for production implementation. At present, the minimization of a single (simple) track takes around 8 s. Any gain in this step will significantly decrease the reconstruction time for the amount of tracking data produced by the experiment.

- If a model can be developed that performs consistently and reliably well on track candidates that are within the quality cuts defined in section 6.1.2, it should be tested with track candidates that would ordinarily be discarded, e.g. shorter tracks. The ultimate goal of exploring alternative reconstruction methods is to increase the amount of reconstructed tracks, thus being able to have less stringent quality cuts would be a desired outcome.

CHAPTER 7

Summary

In this work, the muon-weighted average of the magnetic field during Run-4/5/6 of the Fermilab Muon g-2 experiment has been calculated, which is a crucial step in the determination of the anomalous magnetic moment of the muon. The results are presented in chapter 5, tables 5.5-5.8, and illustrate the significant impact of the muon weighting in reducing the contribution of higher order magnetic field multipoles. A new algorithm for constructing run groups has been implemented, and effects of outlier run groups carefully assessed. Additionally, an extensive evaluation of systematic uncertainty sources was performed, including testing the sensitivity of the muon weighting procedure to variations in any of the necessary supplemental inputs like the simulated beam functions, and a cross-check of the correction due to ESQ misalignments causing a COD. The uncertainties are summarized in tables 5.15.

Furthermore, chapter 4 explored a different field parameterization than the 2D multipole expansion that is typically used in the Muon g-2 experiment. The toroidal geometry allows for the construction of simulated source-free fields that inherently satisfy Laplace's equation, in order to be able to validate fit results in either parameterization against truth values, as well as investigate underlying assumptions about the magnetic field behavior. A fit algorithm for the toroidal field description was developed, taking into account numerical stability, memory requirements due to data set size, and robustness against perturbations due to round-off error.

Finally, a ML approach to track reconstruction based on a likelihood-free inference technique was examined in chapter 6, with the help of an idealized toy model of a tracker station. A NN was trained on simplified simulated tracks, with parameters drawn from realistic distributions, to approximate the negative LLH function of observing a measured track given a set of parameters. This allowed an optimization algorithm to estimate the initial parameters by minimizing the approximated LLH function.

7.1 Conclusions

Neither the results of the muon weighting, nor the determination of the uncertainties have produced any unexpected results. The new strategy to define run groups has been shown to work more efficiently and require less manual intervention than the previous algorithm, as shown in section 5.2.1. All uncertainties were either on the same order of magnitude as in previous run periods, or smaller due to a more stable and better positioned beam. The additional cross-check of the ESQ misalignment COD being included in the beam profile reconstruction shows excellent agreement with the corrections calculated after the fact.

The STFs were extremely beneficial in being able to validate the critical assumption that $|\vec{B}| \approx B_y$ for both the 2D multipole and toroidal fits of the trolley maps, as well as showing that the interpolation of the trolley maps to a common time grid does not significantly impact the fit's ability to reconstruct the field correctly. Moreover, it was shown that the 2D and 3D fits differ by only a few ppb, either to real or simulated data, which demonstrated that the 2D multipole parameterization is a reliable approach for the purpose of the experiment. The toroidal framework has proven to be a useful tool for simulations and systematic studies, as well as for performing cross-checks of the conventional fit algorithm.

The ML algorithm has been shown to successfully learn to approximate the LLH function, providing a proof of principle of using this approach in the Muon g-2 tracking detectors. The quantitative results are not yet competitive with the currently implemented, deterministic track reconstruction algorithm, but promising. The general framework has been established, and next steps for refinement and further development have been discussed at length in chapter 6.

7.2 Outlook

While there is currently no direct succession planned for a storage-ring based muon g-2 experiment, any future effort will require reduced systematic uncertainties, many of which are addressed in this dissertation. The methods in this work are applicable to other types of problems as well. The muon-weighting is effectively a type of phase-space weighting, which is relevant in any setting where a non-uniform distribution of particles is affected by fields. ML-based track extrapolation algorithms are already in use in several other experiments, and opportunities for similar methods continue to grow. On the other hand, the toroidal

parameterization is specific to this geometry and magnetic fields and can't be easily generalized to other experimental setups. In the context of an appropriate problem however, it shows significant potential and there are several avenues for further study. For one, STFs enable the separation of effects due to different toroidal components, and more tests to correlate artifacts in the fit residuals with higher order components would be warranted. Secondly, the STFs in this work did not simulate any uncertainties or noise on the data, so a natural extension would be to investigate at which level of distortion either of the fit algorithms start to become unreliable. And finally, there are still other systematic checks that could be performed, e.g. the effect of different azimuthal bin sizes in the muon weighting.

APPENDIX A

Experiment Parameters and Properties

Major radius R_0	7.112m
Minor radius (vacuum chamber radius)	5cm
Radius of active muon storage volume r_0	4.5cm
Nominal magnetic field	1.4513T
High voltage	18.2kV
Nominal muon momentum	3.094GeV
Cyclotron frequency	6.71MHz
Anomalous precession frequency	0.23MHz
Horizontal betatron frequency	6.23MHz
Vertical betatron frequency	2.48MHz

Table A.1: Parameters and operational conditions of the muon storage ring

	noRF	xRF	xyRF-5	xyRF-6
included ω_a-datasets	4A-U, 5A	5B-L	5M-U	6A-K
ω_a run range	37310-46429	46431-50335	50399-54150	55131-59143
field run range	21019-24413	24413-25668	25668-27125	28100-29744

Table A.2: Definition of RF datasets

	Field	Tracking	COSY
Viewpoint	transverse cross-section	3D cylindrical	either
Radial coordinate	$x = R - R_0$ or polar: $r = \sqrt{x^2 + y^2}$, $\theta = \arctan(y/x)$	R within $R_0 \pm r_0$	x or R
Azimuthal ($\phi = 0$) reference point	ring's 12 o'clock (top down view)	exit of inflector field + 18.35 deg	upstream edge of ESQ1 field + 33.35 deg

Table A.3: Overview of different coordinate systems in the Muon g-2 experiment. In any coordinate system, y indicates the vertical direction.

APPENDIX B

Derivations and Detailed Calculations

B.1 2D transverse coordinate system

B.1.1 Magnetic field components from polar Laplace's equation

The magnetic scalar potential Φ for a source free region in polar coordinates (r, θ), with coefficients a_n and b_n, is (quoting eq. 4.3)

$$\Phi(r,\theta) = a_0 + \sum_{n=1}^{\infty} r^n \Big(a_n \sin(n\theta) + b_n \cos(n\theta) \Big). \tag{4.3}$$

The transformation between polar and (2D) cartesian coordinates, as well as the partial derivatives of each coordinate, are given in table B.1.[1]

[1]The expression given in table B.1 for θ is technically only correct for $x > 0$, since arctan only takes values between $-\pi/2$ and $\pi/2$ (excluding). It can be extended to the entire interval $[-\pi, \pi]$ through piecewise definition,

$$\theta = \begin{cases} \arctan(y/x) & \text{if } x > 0, \\ \arctan(y/x) + \pi & \text{if } x < 0, y \geq 0, \\ \arctan(y/x) - \pi & \text{if } x < 0, y < 0, \\ \pi/2 & \text{if } x = 0, y > 0, \\ -\pi/2 & \text{if } x = 0, y < 0. \end{cases}$$

Alternatively, one can make use of the tangent half-angle identity,

$$\tan\frac{\theta}{2} = \frac{1-\cos\theta}{\sin\theta} = \frac{\sin\theta}{1+\cos\theta},$$

to obtain an expression for θ that is continuous, defined on the entire interval, and has the same partial derivatives as $\arctan(y/x)$:

$$\theta = 2\arctan(\frac{r-x}{y}) = 2\arctan(\frac{y}{r+x}).$$

Polar $(r,\,\theta)$		Cartesian $(x,\,y)$	
$r = \sqrt{x^2 + y^2}$	(B.1)	$x = r\cos\theta$	(B.2)
$\theta = \arctan\left(\frac{y}{x}\right)$	(B.3)	$y = r\sin\theta$	(B.4)
$\frac{\partial r}{\partial x} = \frac{x}{r} = \cos\theta$	(B.5)	$\frac{\partial x}{\partial r} = \cos\theta = \frac{x}{r}$	(B.6)
$\frac{\partial r}{\partial y} = \frac{y}{r} = \sin\theta$	(B.7)	$\frac{\partial x}{\partial \theta} = -r\sin\theta = -y$	(B.8)
$\frac{\partial \theta}{\partial x} = \frac{1}{1+\frac{y^2}{x^2}}\frac{-y}{x^2} = -\frac{y}{r^2}$	(B.9)	$\frac{\partial y}{\partial r} = \sin\theta = \frac{y}{r}$	(B.10)
$\frac{\partial \theta}{\partial y} = \frac{1}{1+\frac{y^2}{x^2}}\frac{1}{x} = \frac{x}{r^2}$	(B.11)	$\frac{\partial y}{\partial \theta} = r\cos\theta = x$	(B.12)

Table B.1: Transformation between polar and cartesian coordinates

The magnetic field $\vec{B} = (B_x, B_y)$ can be obtained from the scalar potential via [91]

$$\vec{B} = -\nabla\Phi = (-\frac{\partial\Phi}{\partial x}, -\frac{\partial\Phi}{\partial y}), \tag{B.13}$$

and therefore the individual magnetic field component in the y direction can be written as

$$\begin{aligned}
B_y &= -\frac{\partial\Phi}{\partial y} = -\frac{\partial\Phi}{\partial r}\frac{\partial r}{\partial y} - \frac{\partial\Phi}{\partial\theta}\frac{\partial\theta}{\partial y} && \text{(B.14)}\\
&= -\frac{y}{r}\sum_{n=1}^{\infty} nr^{n-1}\Big(a_n\sin(n\theta) + b_n\cos(n\theta)\Big) - \frac{x}{r^2}\sum_{n=1}^{\infty} nr^n\Big(a_n\cos(n\theta) - b_n\sin(n\theta)\Big)\\
&= -\sin(\theta)\sum_{n=1}^{\infty} nr^{n-1}\Big(a_n\sin(n\theta) + b_n\cos(n\theta)\Big) - \cos(\theta)\sum_{n=1}^{\infty} nr^{n-1}\Big(a_n\cos(n\theta) - b_n\sin(n\theta)\Big)\\
&= \sum_{n=1}^{\infty} nr^{n-1}\Big[-a_n\Big(\cos(\theta)\cos(n\theta) + \sin(\theta)\sin(n\theta)\Big) + b_n\Big(\cos(\theta)\sin(n\theta) - \sin(\theta)\cos(n\theta)\Big)\Big]\\
&= \sum_{n=1}^{\infty} nr^{n-1}\Big[-a_n\cos((n-1)\theta) + b_n\sin((n-1)\theta)\Big]. && \text{(B.15)}
\end{aligned}$$

Similarly, the magnetic field component in the x direction is

$$B_x = -\frac{\partial \Phi}{\partial x} = -\frac{\partial \Phi}{\partial r}\frac{\partial r}{\partial x} - \frac{\partial \Phi}{\partial \theta}\frac{\partial \theta}{\partial x} \tag{B.16}$$

$$= -\frac{x}{r}\sum_{n=1}^{\infty} nr^{n-1}\Big(a_n \sin(n\theta) + b_n \cos(n\theta)\Big) + \frac{y}{r^2}\sum_{n=1}^{\infty} nr^{n}\Big(a_n \cos(n\theta) - b_n \sin(n\theta)\Big)$$

$$= -\cos(\theta)\sum_{n=1}^{\infty} nr^{n-1}\Big(a_n \sin(n\theta) + b_n \cos(n\theta)\Big) + \sin(\theta)\sum_{n=1}^{\infty} nr^{n-1}\Big(a_n \cos(n\theta) - b_n \sin(n\theta)\Big)$$

$$= \sum_{n=1}^{\infty} nr^{n-1}\Big[-a_n\Big(\cos(\theta)\sin(n\theta) - \sin(\theta)\cos(n\theta)\Big) - b_n\Big(\cos(\theta)\cos(n\theta) + \sin(\theta)\sin(n\theta)\Big)\Big]$$

$$= \sum_{n=1}^{\infty} nr^{n-1}\Big[-a_n \sin((n-1)\theta) - b_n \cos((n-1)\theta)\Big], \tag{B.17}$$

where the following trigonometric identities have been used to arrive at eq. B.15 and B.17:

$$\cos(\theta)\cos(n\theta) = 0.5\Big(\cos((n-1)\theta) + \cos((n+1)\theta)\Big) \tag{B.18}$$

$$\sin(\theta)\sin(n\theta) = 0.5\Big(\cos((n-1)\theta) - \cos((n+1)\theta)\Big) \tag{B.19}$$

$$\cos(\theta)\sin(n\theta) = 0.5\Big(\sin((n+1)\theta) + \sin((n-1)\theta)\Big) \tag{B.20}$$

$$\sin(\theta)\cos(n\theta) = 0.5\Big(\sin((n+1)\theta) - \sin((n-1)\theta)\Big) \tag{B.21}$$

Now let $m = n - 1$. Then

$$B_x = -b_1 + \sum_{m=1}^{\infty}(m+1)r^m\Big[-a_{m+1}\sin(m\theta) - b_{m+1}\cos(m\theta)\Big]$$

$$= -\tilde{b}_0 + \sum_{m=1}^{\infty}\left(\frac{r}{r_0}\right)^m\Big[\tilde{a}_m \sin(m\theta) - \tilde{b}_m \cos(m\theta)\Big], \tag{B.22}$$

$$B_y = -a_1 + \sum_{m=1}^{\infty}(m+1)r^m\Big[-a_{m+1}\cos(m\theta) + b_{m+1}\sin(m\theta)\Big]$$

$$= \tilde{a}_0 + \sum_{m=1}^{\infty}\left(\frac{r}{r_0}\right)^m\Big[\tilde{a}_m \cos(m\theta) + \tilde{b}_m \sin(m\theta)\Big], \tag{B.23}$$

where A_0 corresponds to a constant magnetic field in the y direction, B_0 corresponds to a constant magnetic field in the x direction, and

$$\tilde{a}_m = -(m+1)\ r_0^m\ a_{m+1}, \tag{B.24}$$

$$\tilde{b}_m = (m+1)\ r_0^m\ b_{m+1}. \tag{B.25}$$

B.2 Toroidal coordinate system

B.2.1 Calculating focal points

In section 4.3.1 it was stated that any 2-dimensional circle at $(R_0, 0)$ with radius r_0, corresponds to a unique combination of focal point position a_0 and coordinate value $\eta = \eta_0$. To determine a_0 and η_0, consider that all lines of constant η are circles centered on the x-axis (see fig. 4.2b), and containing one of the focal points within them. Any point on the circle must have the same ratio of d_1 to d_2, where we define d_1 as the distance to the left focus at $(-a_0, 0)$ and d_2 the distance to the right focus at $(a_0, 0)$. Thus this is also true for the two points at $(R_0 - r_0, 0)$ and $(R_0 + r_0, 0)$. Since all relevant points lie on the x-axis, the distances are just the differences in the x-coordinate. This means that

$$\frac{d_1}{d_2} = \frac{(R_0 - r_0) + a_0}{a_0 - (R_0 - r_0)} = \frac{(R_0 + r_0) + a_0}{(R_0 + r_0) - a_0}, \tag{B.26}$$

$$\Leftrightarrow (R_0 - r_0 + a_0)(R_0 + r_0 - a_0) = (R_0 + r_0 + a_0)(-R_0 + r_0 + a_0),$$

$$\Leftrightarrow R_0^2 - r_0^2 + 2a_0r_0 - a_0^2 = r_0^2 + 2a_0r_0 + a_0^2 - R_0^2,$$

$$\Leftrightarrow a_0^2 = R_0^2 - r_0^2. \tag{B.27}$$

Let $d = d_1/d_2$ such that $\eta = \ln d$ as in eq. 4.19. With $\pm a_0$ uniquely determined from R_0 and r_0, we can then calculate d for the circle around the *positive* focal point ($\eta > 0 \leftrightarrow d > 1$) directly from eq. B.26, and the corresponding η_0 as

$$\begin{aligned}
\eta_0 &= \ln \frac{d_1}{d_2} = \ln \left(\frac{R_0 + r_0 + a_0}{R_0 + r_0 - a_0} \right) = \ln \left(\frac{(R_0 + r_0 + a_0)^2}{(R_0 + r_0)^2 - a_0^2} \right) \\
&= \ln \left(\frac{(R_0 + r_0)^2 + 2(R_0 + r_0)a_0 + (R_0^2 - r_0^2)}{(R_0 + r_0)^2 - (R_0^2 - r_0^2)} \right) \\
&= \ln \left(\frac{2R_0(R_0 + r_0) + 2(R_0 + r_0)a_0}{2r_0(R_0 + r_0)} \right) = \ln \left(\frac{R_0 + a_0}{r_0} \right) = \ln \left(\frac{R_0 + \sqrt{R_0^2 - r_0^2}}{r_0} \right) \\
&= \ln \left(\frac{R_0}{r_0} + \sqrt{\left(\frac{R_0}{r_0} \right)^2 - 1} \right) = \operatorname{arcosh} \left(\frac{R_0}{r_0} \right).
\end{aligned} \tag{B.28}$$

Our region of interest, i.e. the source-free muon storage volume in which Laplace's equation is valid, is then defined by

$$\eta \geq \operatorname{arcosh} \left(\frac{R_0}{r_0} \right), \qquad \vartheta \in [-\pi, \pi], \qquad \psi \in [0, 2\pi]. \tag{B.29}$$

B.2.2 Derivation of Laplace's equation in toroidal coordinates

To determine the expression for the Laplacian in toroidal coordinates, we use the following transformation identities between standard cartesian and toroidal coordinates:

$$x = a\frac{\sinh\eta}{\cosh\eta - \cos\vartheta}\cos\psi = \frac{a}{C}\sinh\eta\cos\psi \tag{B.30}$$

$$y = a\frac{\sinh\eta}{\cosh\eta - \cos\vartheta}\sin\psi = \frac{a}{C}\sinh\eta\sin\psi \tag{B.31}$$

$$z = a\frac{\sin\vartheta}{\cosh\eta - \cos\vartheta} = \frac{a}{C}\sin\vartheta \tag{B.32}$$

where we have defined $C = \cosh\eta - \cos\vartheta$ as a useful shorthand[2].
For general curvilinear coordinates (u,v,w), the Laplacian is [91]:

$$\nabla^2 f = \frac{1}{h_u h_v h_w}\left[\partial_u\left(\frac{h_v h_w}{h_u}\partial_u f\right) + \partial_v\left(\frac{h_u h_w}{h_v}\partial_v f\right) + \partial_w\left(\frac{h_u h_v}{h_w}\partial_w f\right)\right] \tag{B.33}$$

with $h_j = |(\partial_j x, \partial_j y, \partial_j z)^T|$.
For the toroidal coordinate η, we calculate

$$\partial_\eta x = a\cos\psi\frac{\cosh\eta(\cosh\eta - \cos\vartheta) - \sinh^2\eta}{(\cosh\eta - \cos\vartheta)^2} = \frac{a}{C^2}(1 - \cosh\eta\cos\vartheta)\cos\psi, \tag{B.34}$$

$$\partial_\eta y = \frac{a}{C^2}(1 - \cosh\eta\cos\vartheta)\sin\psi, \tag{B.35}$$

$$\partial_\eta z = -\frac{a}{C^2}\sin\vartheta\sinh\eta, \tag{B.36}$$

and thus

$$\begin{aligned} h_\eta &= \frac{a}{C^2}\sqrt{(1 - \cosh\eta\cos\vartheta)^2 + \sinh^2\eta\sin^2\vartheta} \\ &= \frac{a}{C^2}\sqrt{(1 - 2\cosh\eta\cos\vartheta + \cosh^2\eta\cos^2\vartheta) + (\cosh^2\eta - 1)\sin^2\vartheta} \\ &= \frac{a}{C^2}\sqrt{1 - 2\cosh\eta\cos\vartheta + \cosh^2\eta(\cos^2\vartheta + \sin^2\vartheta) - \sin^2\vartheta} \\ &= \frac{a}{C^2}\sqrt{1 - 2\cosh\eta\cos\vartheta + \cosh^2\eta - (1 - \cos^2\vartheta)} \\ &= \frac{a}{C^2}\sqrt{-2\cosh\eta\cos\vartheta + \cosh^2\eta + \cos^2\vartheta} \\ &= \frac{a}{C^2}\sqrt{(\cosh\eta - \cos\vartheta)^2} = \frac{a}{C}. \end{aligned} \tag{B.37}$$

[2]Since $\eta \geq \operatorname{arcosh}\left(\frac{R_0}{r_0}\right) = \operatorname{arcosh}\left(\frac{7112}{45}\right) > 1$, C is in our case always positive.

Similarly, we calculate for ϑ

$$\partial_\vartheta x = -a\sinh\eta\frac{\sin\vartheta}{(\cosh\eta-\cos\vartheta)^2}\cos\psi = -\frac{a}{C^2}\sinh\eta\sin\vartheta\cos\psi, \tag{B.38}$$

$$\partial_\vartheta y = -\frac{a}{C^2}\sinh\eta\sin\vartheta\sin\psi, \tag{B.39}$$

$$\partial_\vartheta z = a\frac{(\cosh\eta-\cos\vartheta)\cos\vartheta-\sin^2\vartheta}{(\cosh\eta-\cos\vartheta)^2} = \frac{a}{C^2}(\cosh\eta\cos\vartheta-1), \tag{B.40}$$

and

$$h_\vartheta = \frac{a}{C^2}\sqrt{\sinh^2\eta\,\sin^2\vartheta+(\cosh\eta\cos\vartheta-1)^2} = \frac{a}{C}. \tag{B.41}$$

And finally, for ψ,

$$\partial_\psi x = -\frac{a}{C}\sinh\eta\sin\psi, \tag{B.42}$$

$$\partial_\psi y = \frac{a}{C}\sinh\eta\cos\psi, \tag{B.43}$$

$$\partial_\psi z = 0, \tag{B.44}$$

which leads to

$$h_\psi = \frac{a}{C}\sinh\eta, \tag{B.45}$$

and we obtain the Laplacian in toroidal coordinates from eq. (B.33) as

$$\nabla^2\Phi = \frac{C^3}{a^2\sinh\eta}\left[\partial_\eta\left(\frac{\sinh\eta}{C}\,\partial_\eta\Phi\right)+\sinh\eta\,\partial_\vartheta\left(\frac{1}{C}\,\partial_\vartheta\Phi\right)+\frac{1}{C\sinh\eta}\,\partial_\psi^2\Phi\right]. \tag{4.23}$$

B.2.3 Radial, longitudinal, and vertical field components

The magnetic field components can be obtained as the negative gradients of the scalar magnetic potential

$$\Phi = \sqrt{\cosh\eta-\cos\vartheta}\,\sum_{n,m}U_{nm}(\eta)V_m(\vartheta)W_n(\psi). \tag{4.24}$$

In the cylindrical g-2 lab coordinates (R,ϕ,y), the expressions are [91]

$$B_R = -\frac{\partial\Phi}{\partial R} = -\frac{\partial\Phi}{\partial\eta}\,\frac{\partial\eta}{\partial R}-\frac{\partial\Phi}{\partial\vartheta}\,\frac{\partial\vartheta}{\partial R}-\frac{\partial\Phi}{\partial\psi}\,\frac{\partial\psi}{\partial R}, \tag{B.46}$$

$$B_\phi = -\frac{1}{R}\frac{\partial\Phi}{\partial\phi} = -\frac{1}{R}\left(\frac{\partial\Phi}{\partial\eta}\,\frac{\partial\eta}{\partial\phi}+\frac{\partial\Phi}{\partial\vartheta}\,\frac{\partial\vartheta}{\partial\phi}+\frac{\partial\Phi}{\partial\psi}\,\frac{\partial\psi}{\partial\phi}\right), \tag{B.47}$$

$$B_y = -\frac{\partial\Phi}{\partial y} = -\frac{\partial\Phi}{\partial\eta}\,\frac{\partial\eta}{\partial y}-\frac{\partial\Phi}{\partial\vartheta}\,\frac{\partial\vartheta}{\partial y}-\frac{\partial\Phi}{\partial\psi}\,\frac{\partial\psi}{\partial y}. \tag{B.48}$$

The toroidal coordinates can be determined[3] from (R, ϕ, y) via [105, 3]

$$\eta = \operatorname{artanh}\left(\frac{2aR}{R^2 + a^2 + y^2}\right), \tag{4.20}$$

$$\vartheta = \arctan\left(\frac{2ay}{R^2 + y^2 - a^2}\right), \tag{4.21}$$

$$\psi = \phi. \tag{4.22}$$

Using (for $\alpha = \arctan f$ and $\beta = \operatorname{artanh}\, g$)

$$\frac{\mathrm{d}}{\mathrm{d}f}\arctan f = \frac{1}{1+f^2} = \frac{1}{1+\tan^2\alpha} = \cos^2\alpha, \tag{B.49}$$

$$\frac{\mathrm{d}}{\mathrm{d}g}\operatorname{artanh}\, g = \frac{1}{1-g^2} = \frac{1}{1-\tanh^2\beta} = \cosh^2\beta, \tag{B.50}$$

and the coordinate transformation identities

$$R = a\frac{\sinh\eta}{\cosh\eta - \cos\vartheta} = \frac{a}{C}\sinh\eta, \tag{B.51}$$

$$y = a\frac{\sin\vartheta}{\cosh\eta - \cos\vartheta} = \frac{a}{C}\sin\vartheta, \tag{B.52}$$

the partial derivatives of the toroidal coordinates can be calculated as

$$\begin{aligned}\frac{\partial\eta}{\partial R} &= \cosh^2\eta\left(\frac{2a}{R^2+a^2+y^2} - \frac{(2aR)(2R)}{(R^2+a^2+y^2)^2}\right) = \cosh^2\eta\left(\frac{\tanh\eta}{R} - \frac{\tanh^2\eta}{a}\right)\\ &= \frac{\cosh\eta\,(\cosh\eta - \cos\vartheta)}{a} - \frac{\sinh^2\eta}{a} = \frac{1}{a}\,(1 - \cosh\eta\,\cos\vartheta),\end{aligned} \tag{B.53}$$

$$\frac{\partial\eta}{\partial\phi} = 0, \tag{B.54}$$

$$\begin{aligned}\frac{\partial\eta}{\partial y} &= \cosh^2\eta\,(-1)\frac{(2aR)(2y)}{(R^2+a^2+y^2)^2} = -\cosh^2\eta\,\tanh^2\eta\,\frac{y}{aR} = -\sinh^2\eta\,\frac{\sin\vartheta}{a\sinh\eta}\\ &= -\frac{1}{a}\,\sinh\eta\,\sin\vartheta,\end{aligned} \tag{B.55}$$

$$\begin{aligned}\frac{\partial\vartheta}{\partial R} &= \cos^2\vartheta\,(-1)\frac{(2ay)(2R)}{(R^2+y^2-a^2)^2} = -\cos^2\vartheta\,\tan^2\vartheta\frac{R}{ay} = -\sin^2\vartheta\frac{\sinh\eta}{a\sin\vartheta}\\ &= -\frac{1}{a}\,\sinh\eta\,\sin\vartheta,\end{aligned} \tag{B.56}$$

$$\frac{\partial\vartheta}{\partial\phi} = 0, \tag{B.57}$$

[3] for $R \neq a$, which would correspond to $\eta = \infty$

$$\frac{\partial \vartheta}{\partial y} = \cos^2\vartheta \left(\frac{2a}{R^2+y^2-a^2} - \frac{(2ay)(2y)}{(R^2+y^2-a^2)^2} \right) = \cos^2\vartheta \left(\frac{\tan\vartheta}{y} - \frac{\tan^2\vartheta}{a} \right)$$
$$= \frac{\cos\vartheta\ (\cosh\eta - \cos\vartheta)}{a} - \frac{\sin^2\vartheta}{a} = -\frac{1}{a}\ (1 - \cosh\eta\ \cos\vartheta), \tag{B.58}$$
$$\frac{\partial \psi}{\partial R} = 0, \tag{B.59}$$
$$\frac{\partial \psi}{\partial \phi} = 1, \tag{B.60}$$
$$\frac{\partial \psi}{\partial y} = 0. \tag{B.61}$$

Then the magnetic field components become

$$B_R = -\frac{1}{a}\ (1 - \cosh\eta\ \cos\vartheta)\ \frac{\partial \Phi}{\partial \eta} + \frac{1}{a}\ \sinh\eta\ \sin\vartheta\ \frac{\partial \Phi}{\partial \vartheta}, \tag{4.35}$$
$$B_\phi = -\frac{1}{R}\ \frac{\partial \Phi}{\partial \psi} = -\frac{1}{a}\ \frac{\cosh\eta - \cos\vartheta}{\sinh\eta}\ \frac{\partial \Phi}{\partial \psi}, \tag{4.36}$$
$$B_y = \frac{1}{a}\ \sinh\eta\ \sin\vartheta\ \frac{\partial \Phi}{\partial \eta} + \frac{1}{a}\ (1 - \cosh\eta\ \cos\vartheta)\ \frac{\partial \Phi}{\partial \vartheta}. \tag{4.34}$$

The partial derivatives of the scalar potential with respect to the toroidal coordinates (neglecting the explicit coordinate dependencies) are

$$\frac{\partial \Phi}{\partial \eta} = \frac{\sinh\eta}{2\sqrt{\cosh\eta - \cos\vartheta}}\ \sum_{n,m} U_{nm} V_m W_n + \sqrt{\cosh\eta - \cos\vartheta}\ \sum_{n,m} \frac{\partial U_{nm}}{\partial \eta}\ V_m W_n, \tag{B.62}$$
$$\frac{\partial \Phi}{\partial \vartheta} = \frac{\sin\vartheta}{2\sqrt{\cosh\eta - \cos\vartheta}}\ \sum_{n,m} U_{nm} V_m W_n + \sqrt{\cosh\eta - \cos\vartheta}\ \sum_{n,m} U_{nm} \frac{\partial V_m}{\partial \vartheta} W_n, \tag{B.63}$$
$$\frac{\partial \Phi}{\partial \psi} = \sqrt{\cosh\eta - \cos\vartheta}\ \sum_{n,m} U_{nm} V_m \frac{\partial W_n}{\partial \psi}, \tag{B.64}$$

which allows us to write the magnetic field components in terms of derivatives of the separated solutions as

$$B_R = \frac{1}{a}\ \sinh\eta\ \sin\vartheta\ \left(\frac{\sin\vartheta}{2\sqrt{C}}\ \sum_{n,m} U_{nm} V_m W_n + \sqrt{C}\ \sum_{n,m} U_{nm} \frac{\partial V_m}{\partial \vartheta} W_n \right)$$
$$- \frac{1}{a}\ (1 - \cosh\eta\ \cos\vartheta)\ \left(\frac{\sinh\eta}{2\sqrt{C}}\ \sum_{n,m} U_{nm} V_m W_n + \sqrt{C}\ \sum_{n,m} \frac{\partial U_{nm}}{\partial \eta}\ V_m W_n \right)$$

$$
\begin{aligned}
&= \frac{1}{a}\sum_{n,m} W_n \left[\sqrt{C}\sinh\eta\sin\vartheta\ U_{nm}\ \frac{\partial V_m}{\partial\vartheta} + \frac{\sinh\eta}{2\sqrt{C}}\ U_{nm}\ V_m\left(\sin^2\vartheta - 1 + \cosh\eta\cos\vartheta\right)\right.\\
&\qquad \left. -\sqrt{C}(1-\cosh\eta\cos\vartheta)\frac{\partial U_{mn}}{\partial\eta}\ V_m\right]\\
&= \frac{1}{a}\sum_{n,m} W_n \left[\sqrt{C}\sinh\eta\sin\vartheta\ U_{nm}\ \frac{\partial V_m}{\partial\vartheta} + \frac{\sinh\eta}{2\sqrt{C}}\ U_{nm}\ V_m\left(\cosh\eta\cos\vartheta - \cos^2\vartheta\right)\right.\\
&\qquad \left. -\sqrt{C}(1-\cosh\eta\cos\vartheta)\frac{\partial U_{mn}}{\partial\eta}\ V_m\right]\\
&= \frac{\sqrt{C}}{a}\sum_{n,m} W_n \left[\sinh\eta\sin\vartheta\ U_{nm}\ \frac{\partial V_m}{\partial\vartheta} + \frac{1}{2}\sinh\eta\cos\vartheta\ U_{nm}\ V_m\right.\\
&\qquad \left. -(1-\cosh\eta\cos\vartheta)\frac{\partial U_{mn}}{\partial\eta}\ V_m\right],
\end{aligned}
\tag{B.65}
$$

$$
B_\phi = -\frac{\sqrt{C^3}}{a\sinh\eta}\ \sum_{n,m} U_{nm}V_m\frac{\partial W_n}{\partial\psi},
\tag{B.66}
$$

$$
\begin{aligned}
B_y &= \frac{1}{a}\ \sinh\eta\ \sin\vartheta\ \left(\frac{\sinh\eta}{2\sqrt{C}}\ \sum_{n,m} U_{nm}V_mW_n + \sqrt{C}\ \sum_{n,m}\frac{\partial U_{nm}}{\partial\eta}\ V_mW_n\right)\\
&\quad + \frac{1}{a}\ (1-\cosh\eta\ \cos\vartheta)\ \left(\frac{\sin\vartheta}{2\sqrt{C}}\ \sum_{n,m} U_{nm}V_mW_n + \sqrt{C}\ \sum_{n,m} U_{nm}\frac{\partial V_m}{\partial\vartheta}W_n\right)\\
&= \frac{1}{a}\sum_{n,m} W_n \left[\sqrt{C}\sinh\eta\sin\vartheta\ \frac{\partial U_{nm}}{\partial\eta}\ V_m + \frac{\sin\vartheta}{2\sqrt{C}}\ U_{nm}\ V_m\left(\sinh^2\eta + 1 - \cosh\eta\cos\vartheta\right)\right.\\
&\qquad \left. +\sqrt{C}(1-\cosh\eta\cos\vartheta)U_{mn}\ \frac{\partial V_m}{\partial\vartheta}\right]\\
&= \frac{1}{a}\sum_{n,m} W_n \left[\sqrt{C}\sinh\eta\sin\vartheta\ \frac{\partial U_{nm}}{\partial\eta}\ V_m + \frac{\sin\vartheta}{2\sqrt{C}}\ U_{nm}\ V_m\left(\cosh^2\eta - \cosh\eta\cos\vartheta\right)\right.\\
&\qquad \left. +\sqrt{C}(1-\cosh\eta\cos\vartheta)U_{mn}\ \frac{\partial V_m}{\partial\vartheta}\right]\\
&= \frac{\sqrt{C}}{a}\sum_{n,m} W_n \left[\sinh\eta\sin\vartheta\ \frac{\partial U_{nm}}{\partial\eta}\ V_m + \frac{1}{2}\cosh\eta\sin\vartheta\ U_{nm}\ V_m\right.\\
&\qquad \left. +(1-\cosh\eta\cos\vartheta)U_{mn}\ \frac{\partial V_m}{\partial\vartheta}\right].
\end{aligned}
\tag{B.67}
$$

B.2.4 Exclusion of diverging terms

The explicit forms of the separated solutions are

$$U_{mn}(\eta) = A_{nm} P^n_{m-1/2}(\cosh\eta) + B_{nm} Q^n_{m-1/2}(\cosh\eta), \tag{4.27}$$

$$V_m(\vartheta) = a_m \sin(m\vartheta) + b_m \cos(m\vartheta), \tag{4.25}$$

$$W_n(\psi) = c_n \sin(n\psi) + d_n \cos(n\psi). \tag{4.26}$$

Both $V(\vartheta)$ and $W(\psi)$ as well as their derivatives are finite as long as the coefficients are finite. However, the Legendre functions (toroidal harmonics) $P^n_{m-1/2}$ and $Q^n_{m-1/2}$ can diverge for certain values of η. The Legendre functions can be expressed in terms of the hypergeometric function ${}_2F_1$ and gamma function Γ as [42, 3]

$$\begin{aligned} P^n_{m-1/2}(\cosh\eta) = {} & \frac{\Gamma(m+n+\frac{1}{2})}{2^n\ \Gamma(n+1)\ \Gamma(m-n+\frac{1}{2})}\ (\tanh\eta)^n (\cosh\eta)^{m-\frac{1}{2}} \\ & \times\ {}_2F_1\left(\frac{1}{2}(n-m+\frac{1}{2}), \frac{1}{2}(n-m+\frac{3}{2}); n+1; (\tanh\eta)^2\right), \end{aligned} \tag{B.68}$$

$$\begin{aligned} Q^n_{m-1/2}(\cosh\eta) = {} & (-1)^n \sqrt{\frac{\pi}{2}}\ \frac{\Gamma(m+n+\frac{1}{2})}{2^m\ \Gamma(m+1)}\ \frac{(\tanh\eta)^n}{(\cosh\eta)^{m+\frac{1}{2}}} \\ & \times\ {}_2F_1\left(\frac{1}{2}(n+m+\frac{1}{2}), \frac{1}{2}(n+m+\frac{3}{2}); m+1; \frac{1}{(\cosh\eta)^2}\right). \end{aligned} \tag{4.40}$$

$P^n_{m-1/2}(\cosh\eta)$ is singular for $\eta = \infty$ and $m > 0$ due to the $(\cosh\eta)^{m+\frac{1}{2}}$ factor, since $\cosh\eta \to \infty$ as $\eta \to \infty$. $Q^n_{m-1/2}(\cosh\eta)$ is singular for $\eta = 0$ and $n > 0$ because its hypergeometric function diverges when $\cosh^{-2}\eta = 1$. The hypergeometric function ${}_2F_1(a, b; c; z)$ with $a, b, c \in \mathbb{R}$ and $z \in \mathbb{C}$ can be expressed as a series for $|z| \leq 1$ as [3]

$${}_2F_1(a, b; c; z) = \frac{\Gamma(c)}{\Gamma(a)\Gamma(b)} \sum_{n=0}^{\infty} \frac{\Gamma(a+n)\Gamma(b+n)}{\Gamma(c+n)} \frac{z^n}{n!}. \tag{B.69}$$

This series always converges for $|z| < 1$, but for $|z| = 1$ the value of $(c - a - b)$ determines convergence or divergence. In the case of $Q^n_{m-1/2}(\cosh\eta)$, where $z = \cosh^{-2}\eta \in (0, 1]$, the series always converges for $\eta \neq 0$. For $\eta = 0$ and thus $\cosh^{-2}\eta = 1$, we examine

$$c - a - b = m + 1 - 0.5(n + m + 0.5) - 0.5(n + m + 1.5) = -n, \tag{B.70}$$

which for $n \geq 0$ leads to divergence [3].
Similarly, the hypergeometric function in the remaining $P^n_{-1/2}(\cosh\eta)$ terms (the $m = 0$ case)

diverges in the $\eta \to \infty$ limit, because $\tanh^2 \eta \to 1$ and

$$c - a - b = n + 1 - 0.5(n + 0.5) - 0.5(n + 1.5) = 0. \tag{B.71}$$

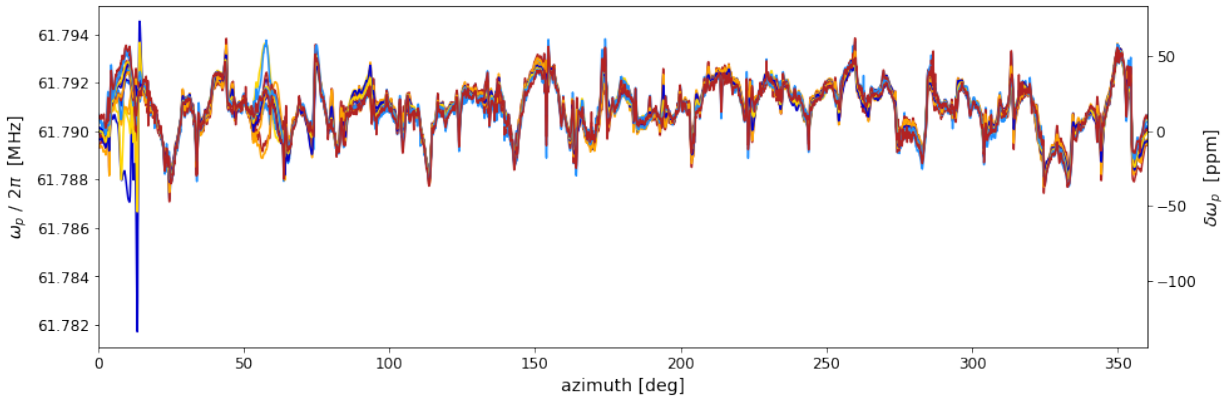
Since the relevant region of interest includes the focal point (where $\eta = \infty$) and excludes the vertical axis (where $\eta = 0$), it is immediately clear that **the potential is restricted to terms including only the toroidal harmonics of the second kind, $\boldsymbol{Q^n_{m-1/2}}$.**

APPENDIX C

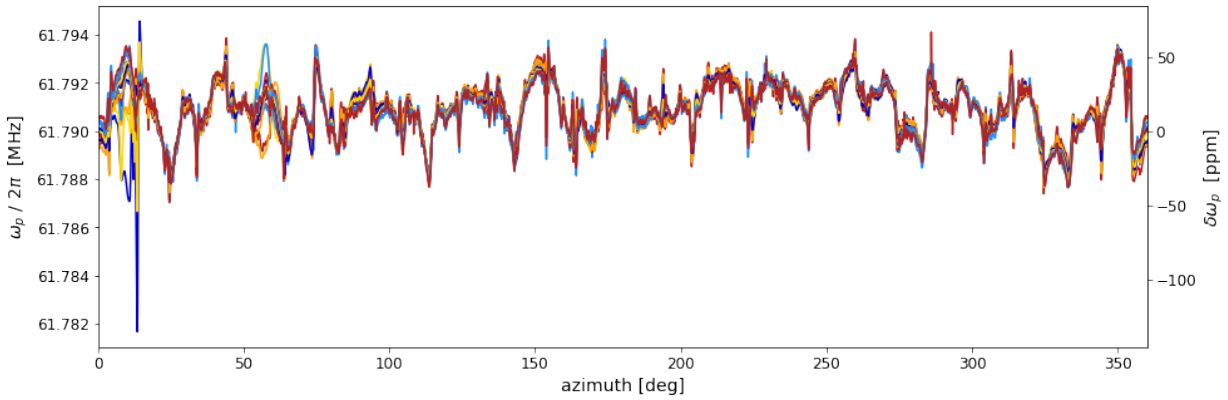
Simulated Toroidal Field Plots

C.1 Input trolley maps

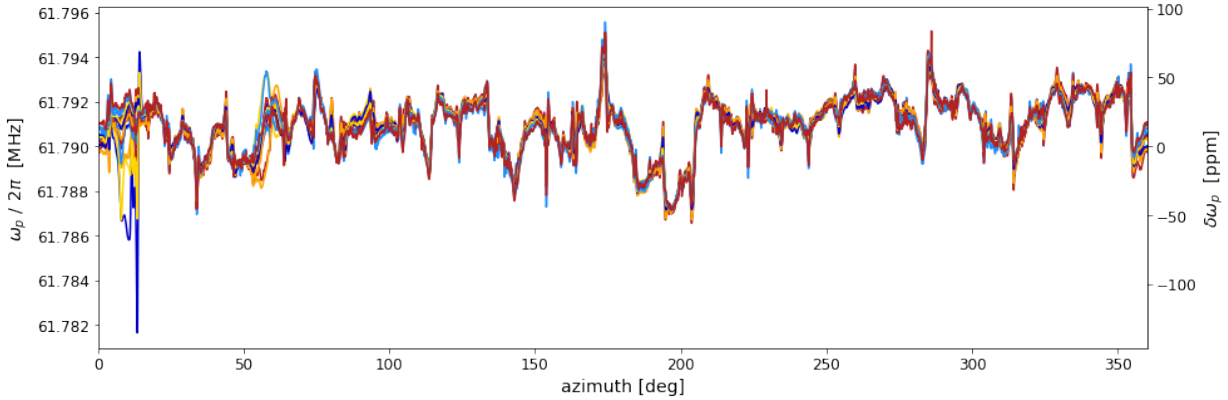
C.1.1 Azimuthal distribution



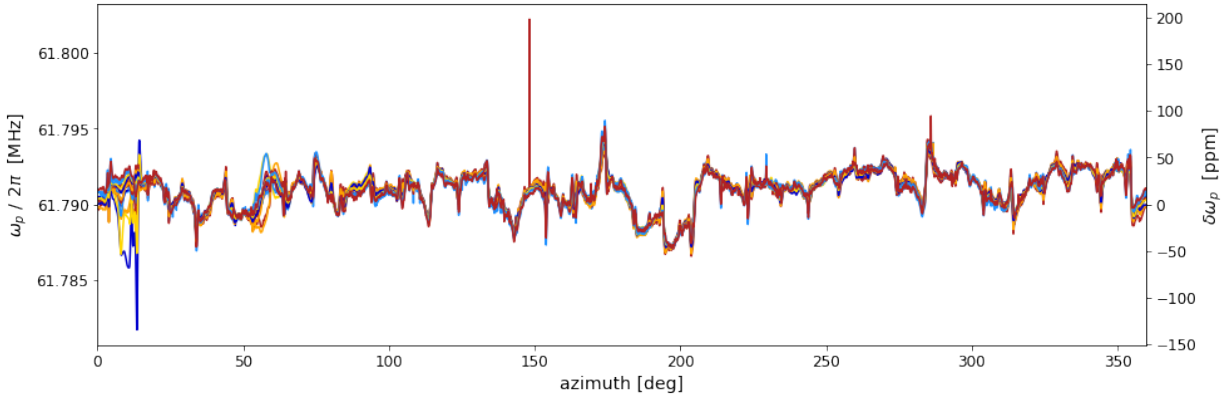
(a) Trolley run 21118 (all probes)



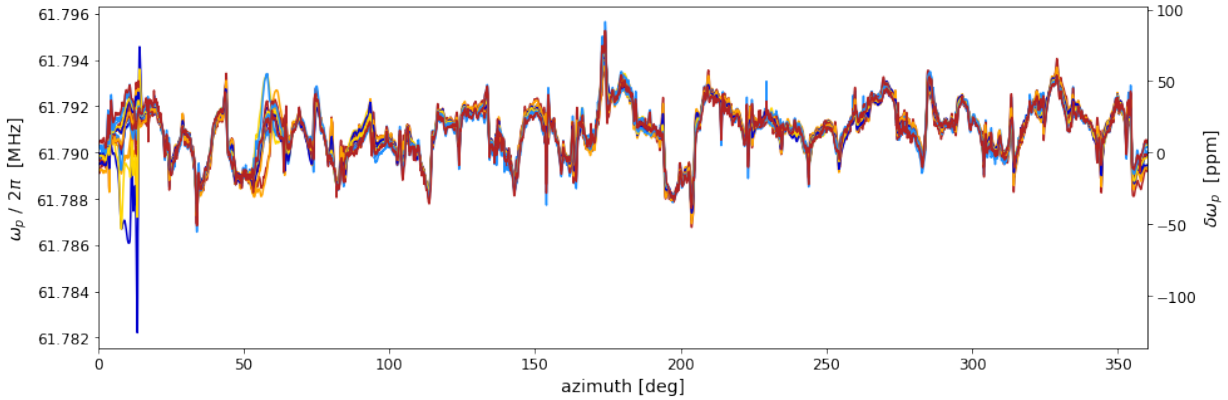
(b) Trolley run 21132 (all probes)



(c) Trolley run 25392 (all probes)



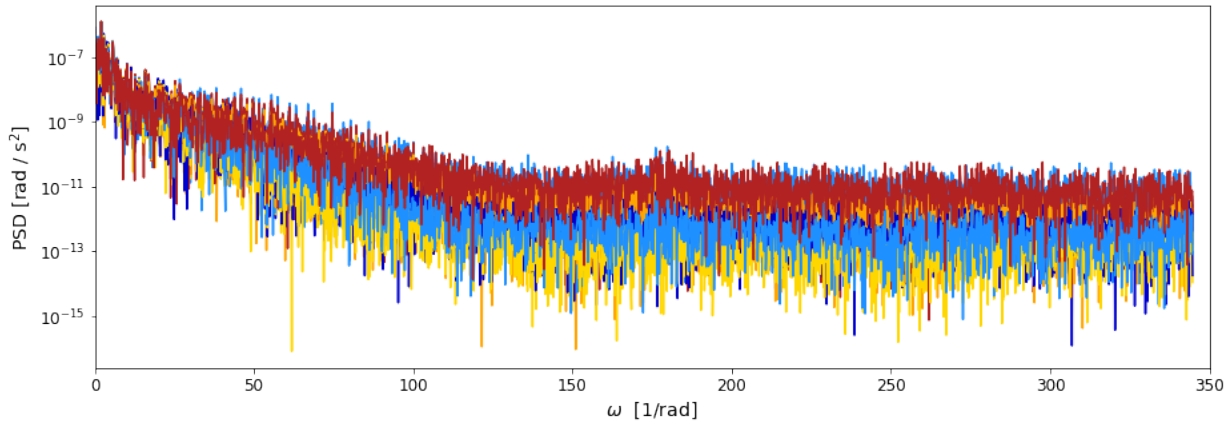
(d) Trolley run 25405 (all probes)



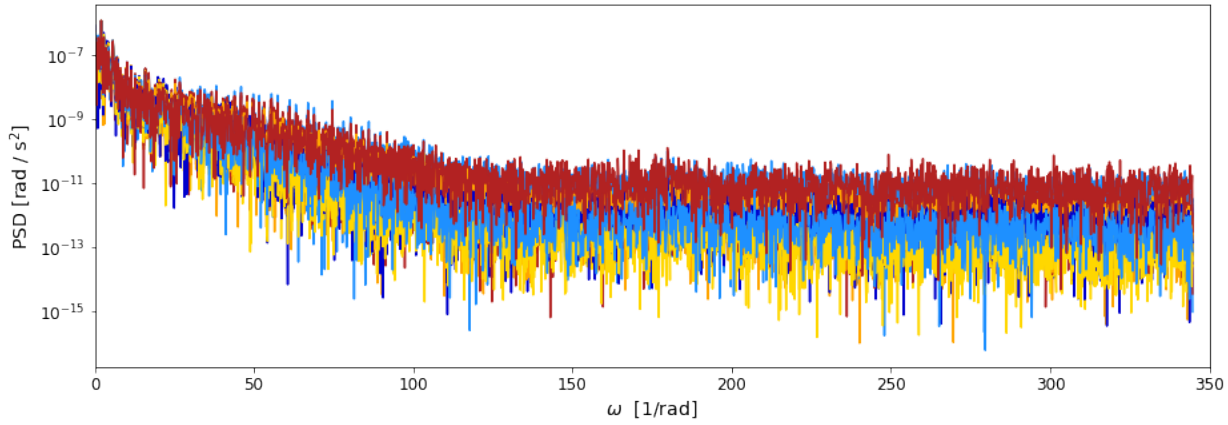
(e) Trolley run 28100 (all probes)

Figure C.1: Trolley runs that are used as input to create the simulated toroidal fields.

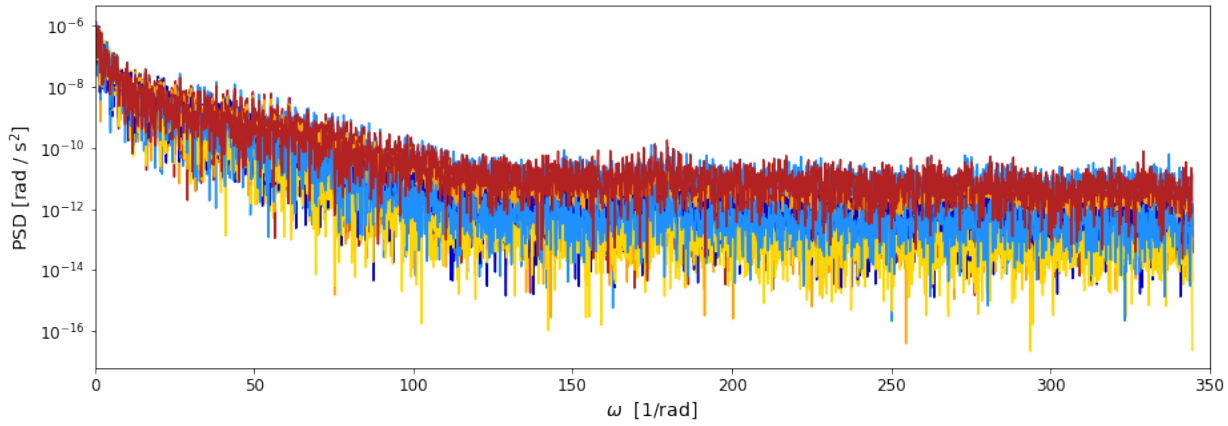
C.1.2 FFTs



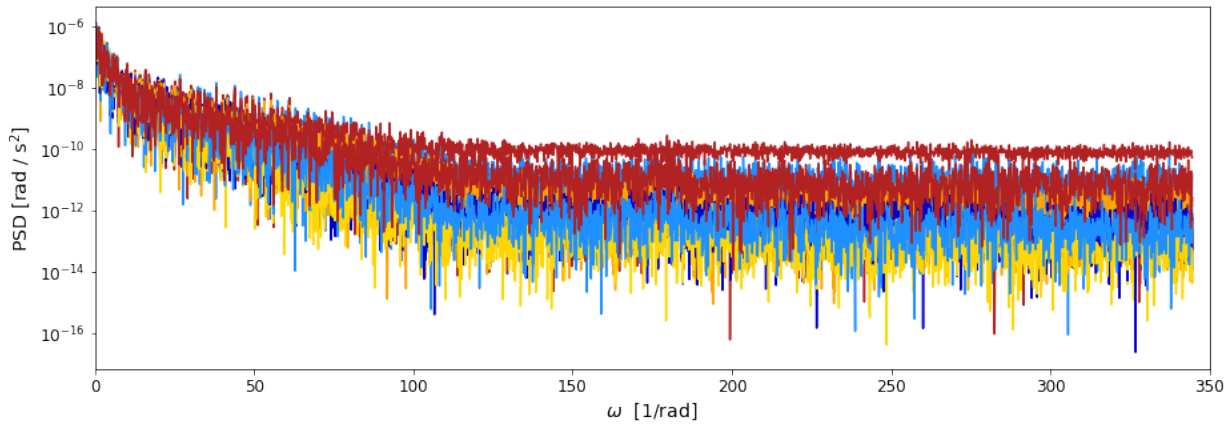
(a) Trolley run 21118 (all probes)



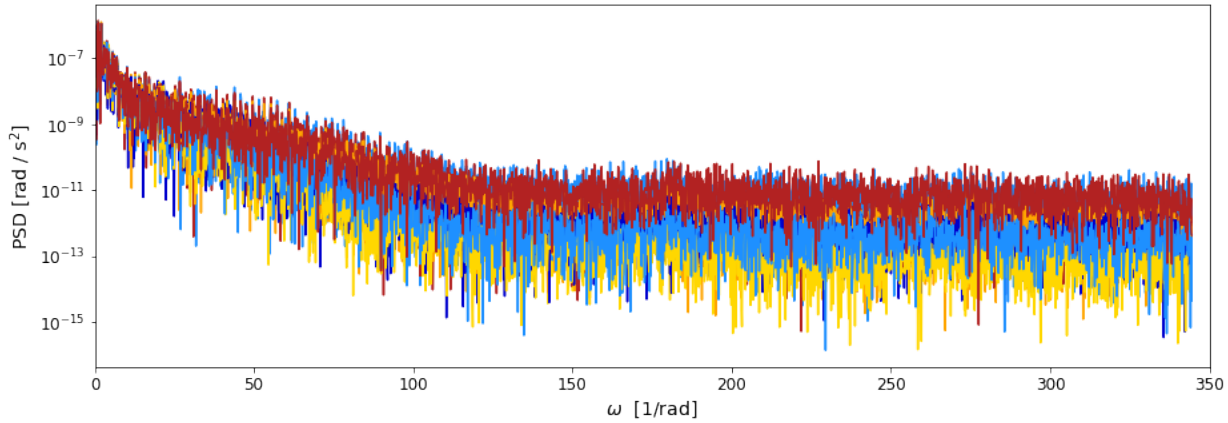
(b) Trolley run 21132 (all probes)



(c) Trolley run 25392 (all probes)



(d) Trolley run 25405 (all probes)

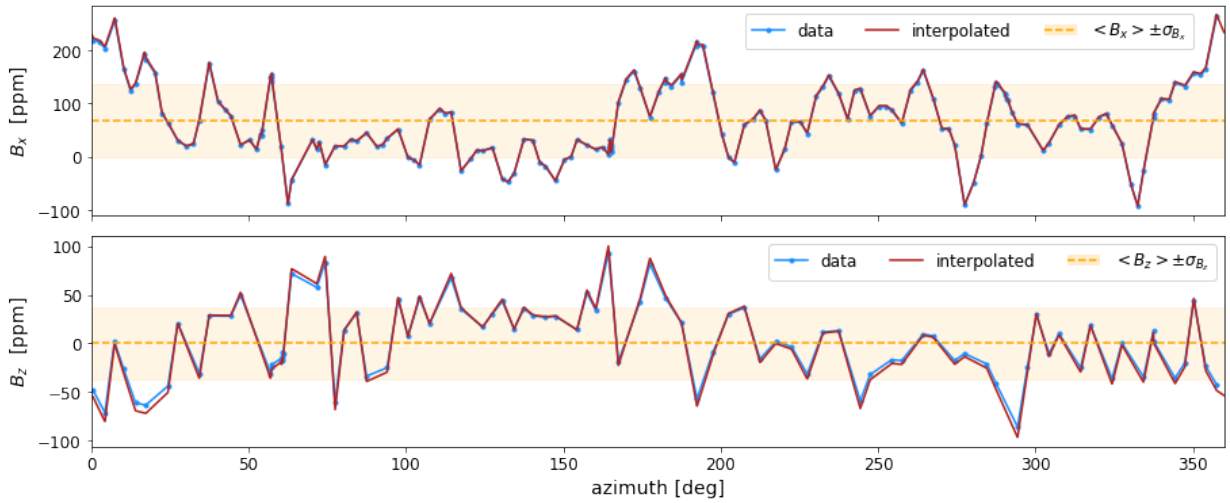


(e) Trolley run 28100 (all probes)

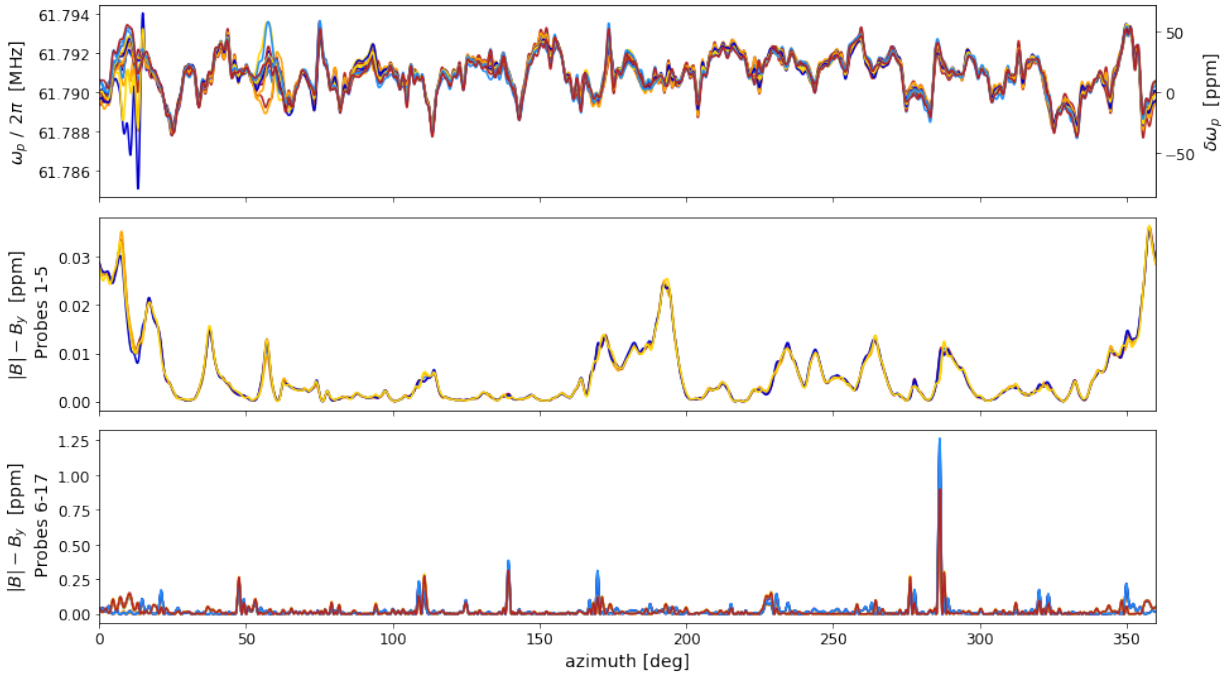
Figure C.2: FFTs of trolley runs, overlaying each of the 17 probes.

C.2 STF1

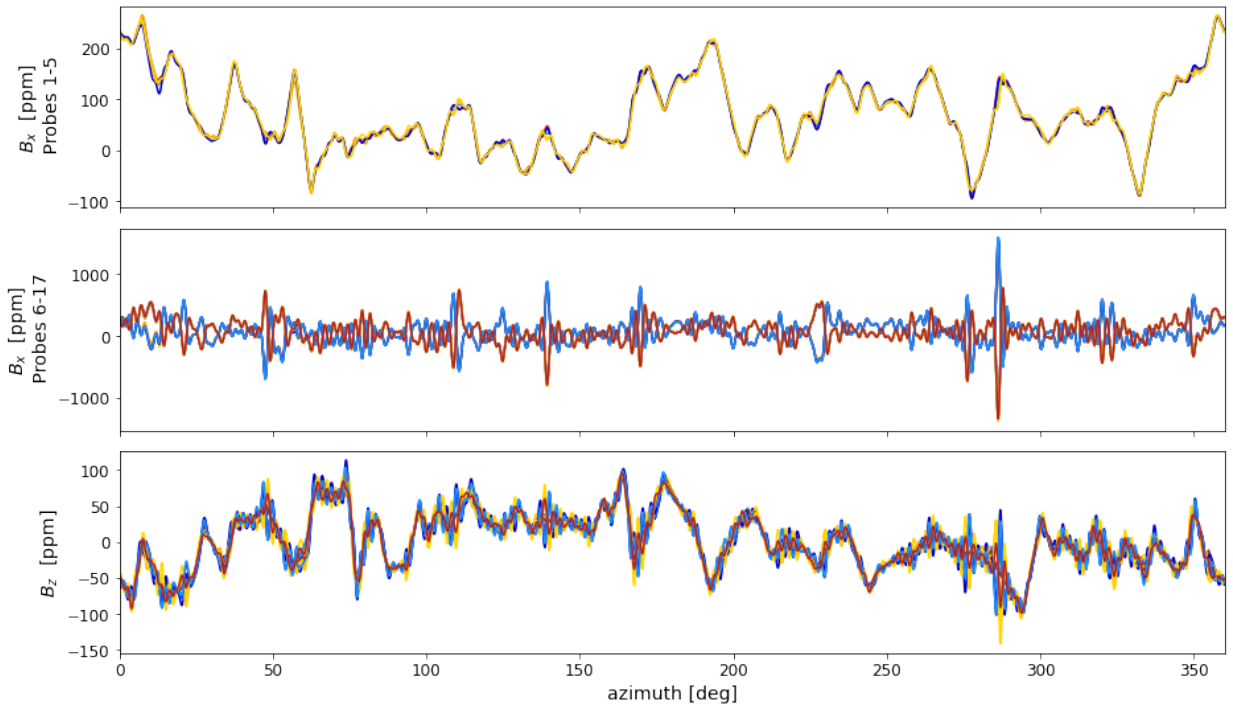
STF1 is based on trolley run 21118 (fig. C.1a).



(a) Radial and longitudinal field input



(b) Simulated field magnitude and difference to vertical field

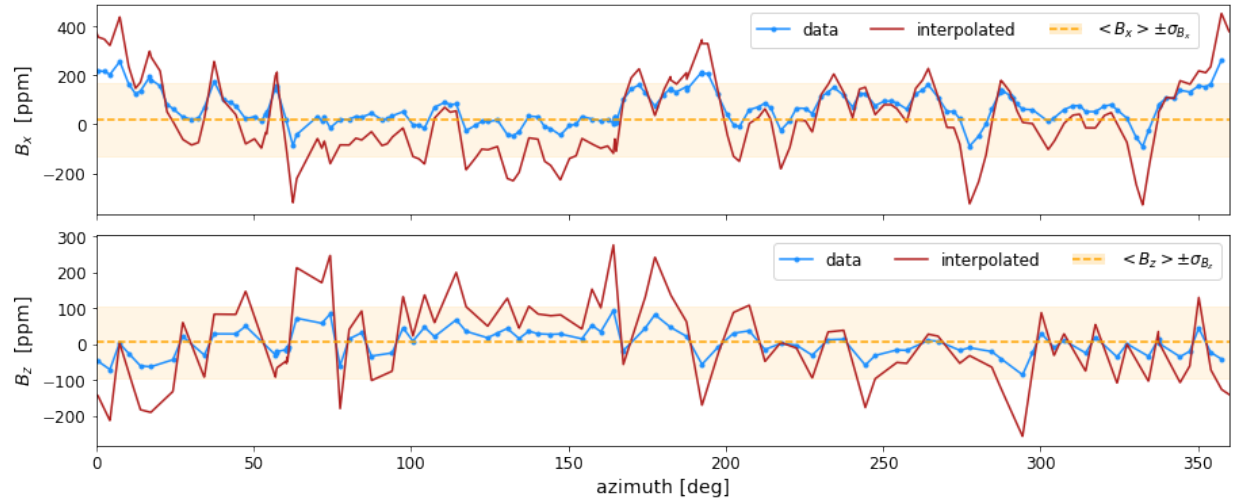


(c) Simulated radial and longitudinal field

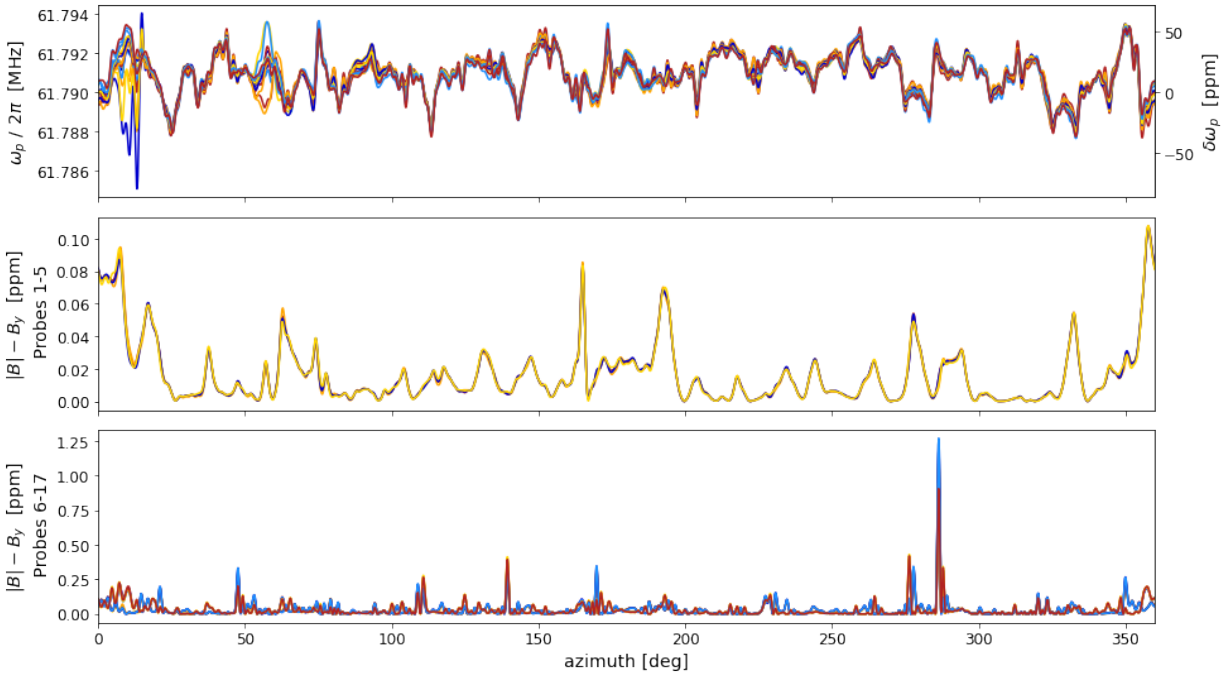
Figure C.3: Graphs of STF1

C.3 STF2

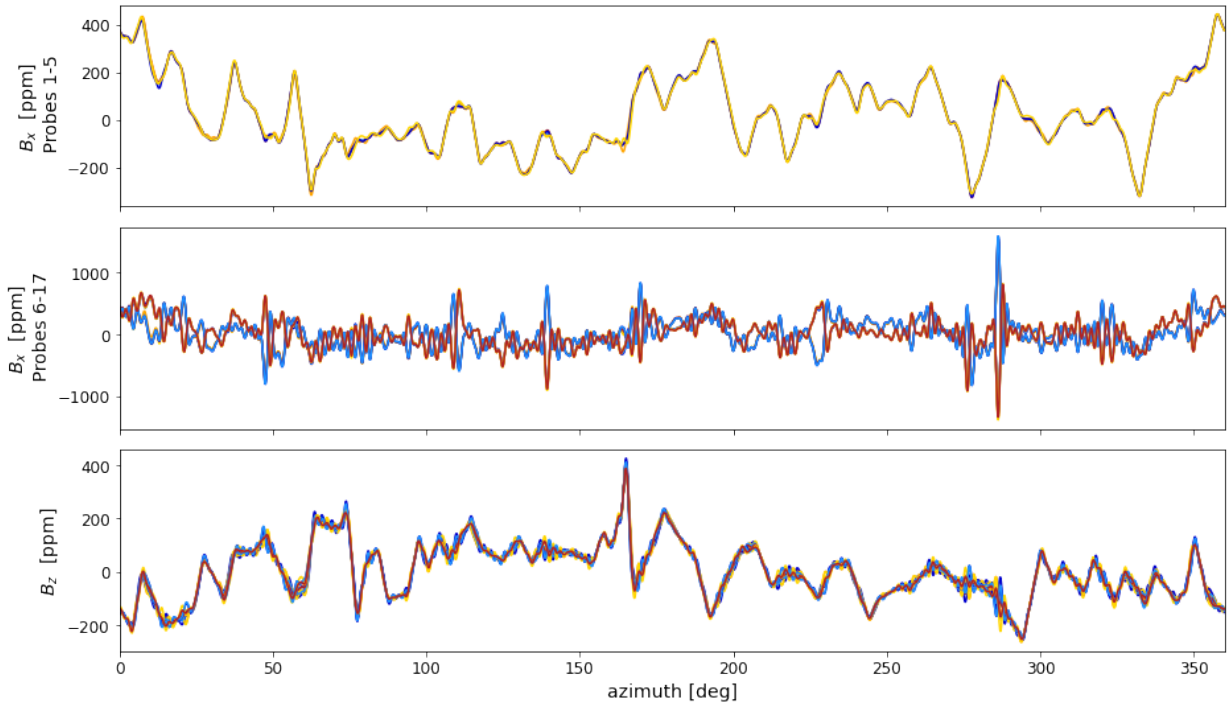
STF2 is based on trolley run 21118 (fig. C.1a).



(a) Radial and longitudinal field input



(b) Simulated field magnitude and difference to vertical field

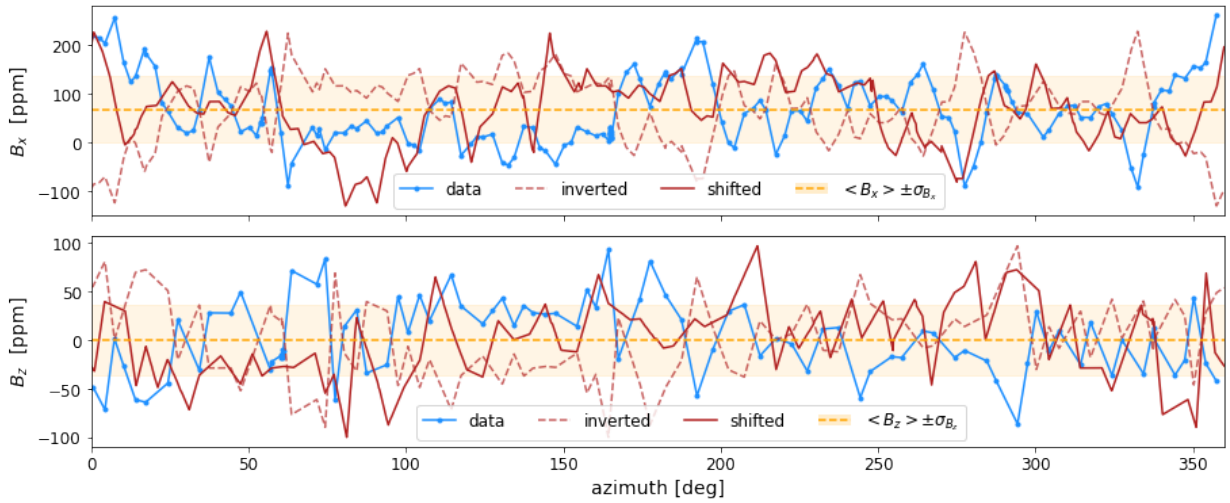


(c) Simulated radial and longitudinal field

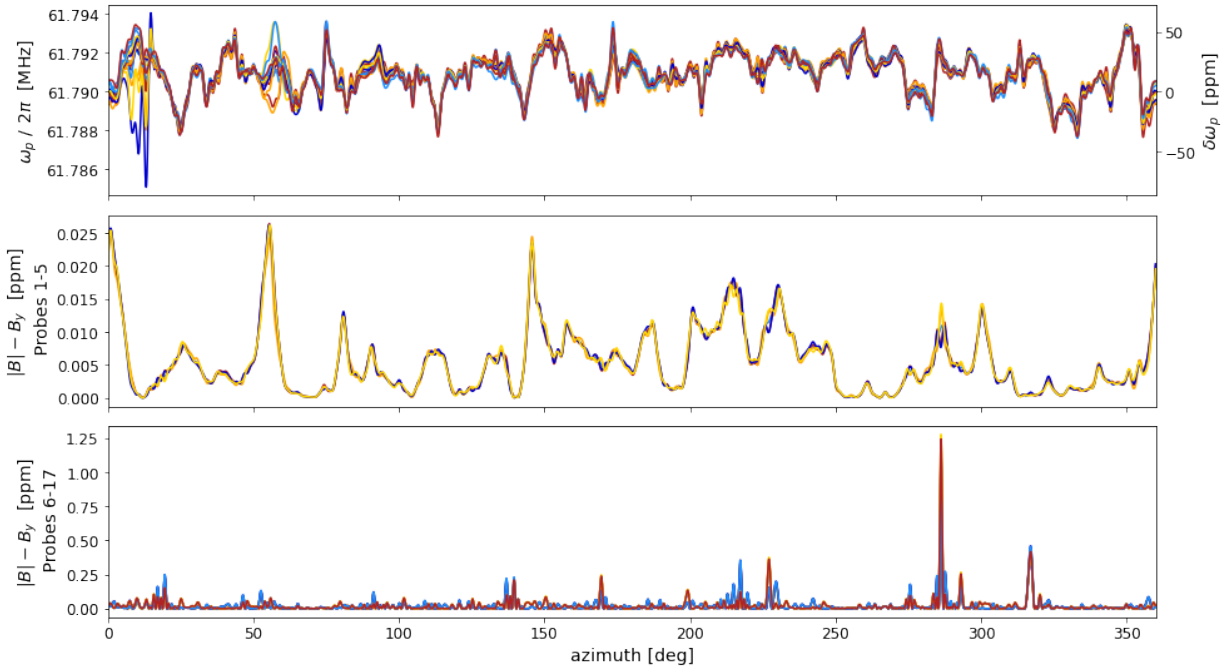
Figure C.4: Graphs of STF2

C.4 STF3

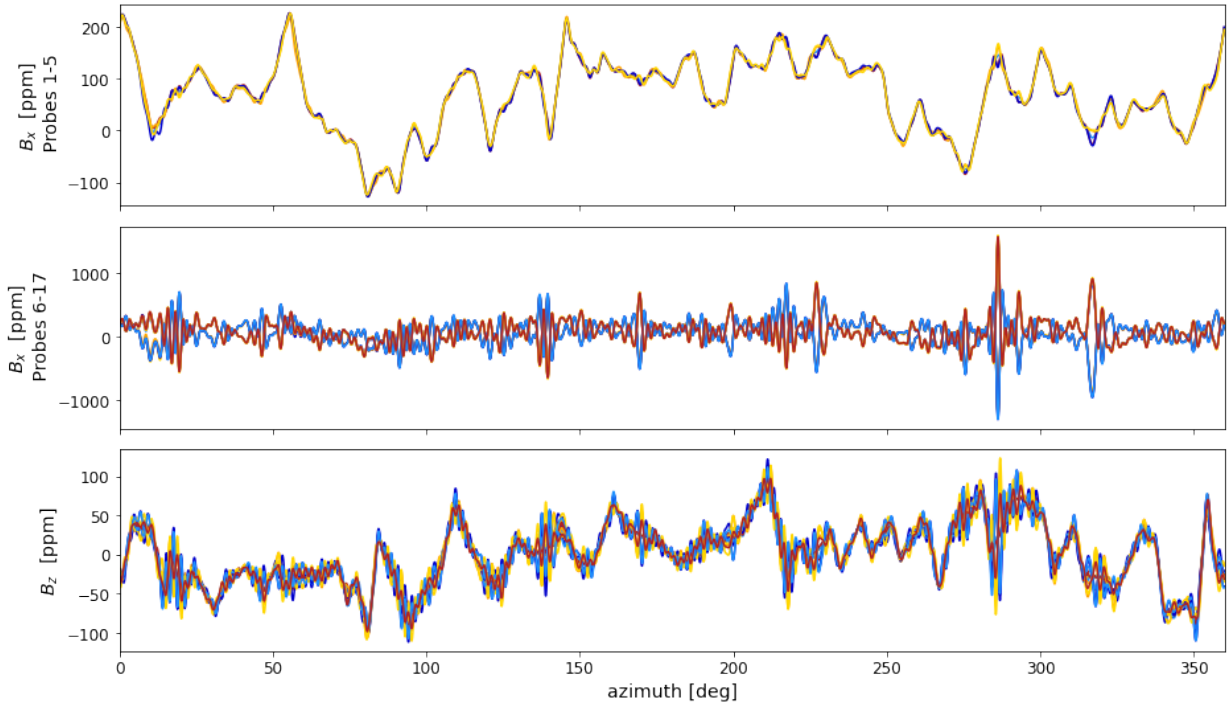
STF3 is based on trolley run 21132 (fig. C.1b).



(a) Radial and longitudinal field input



(b) Simulated field magnitude and difference to vertical field

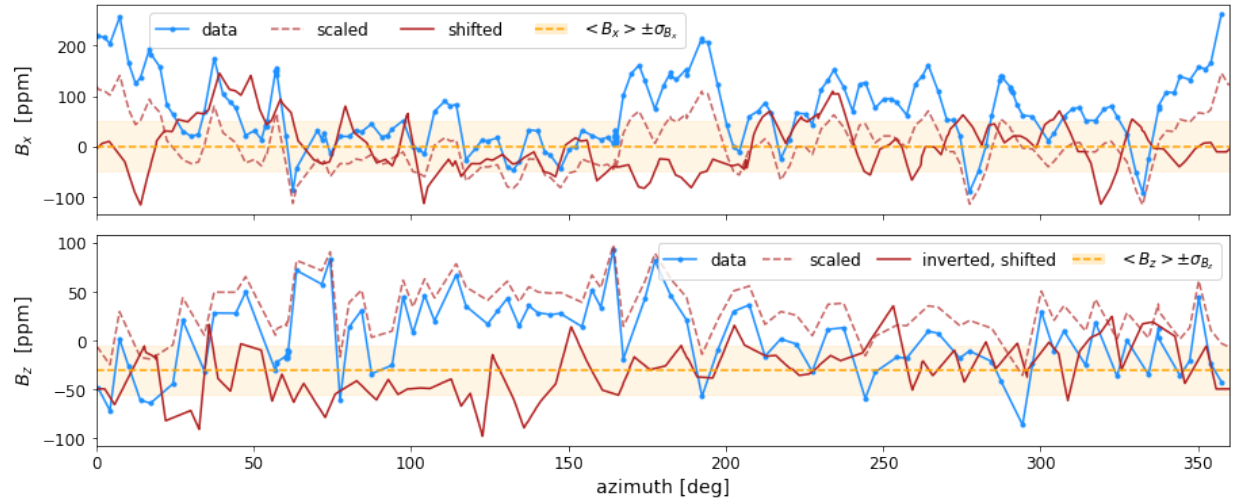


(c) Simulated radial and longitudinal field

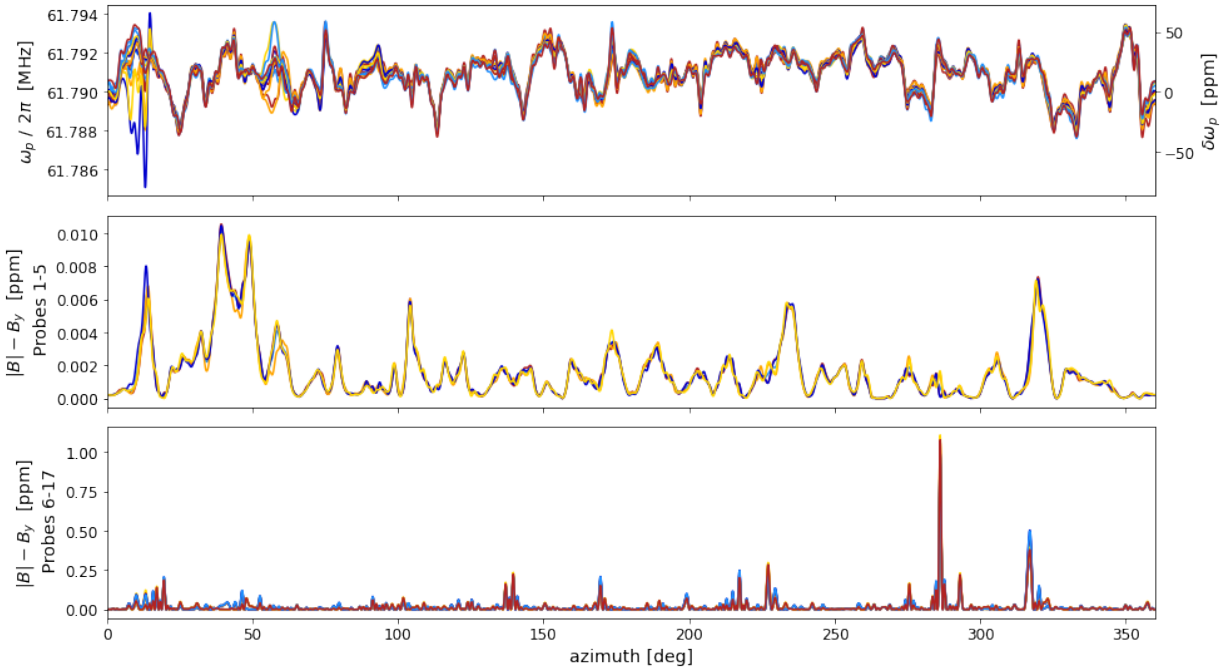
Figure C.5: Graphs of STF3

C.5 STF4

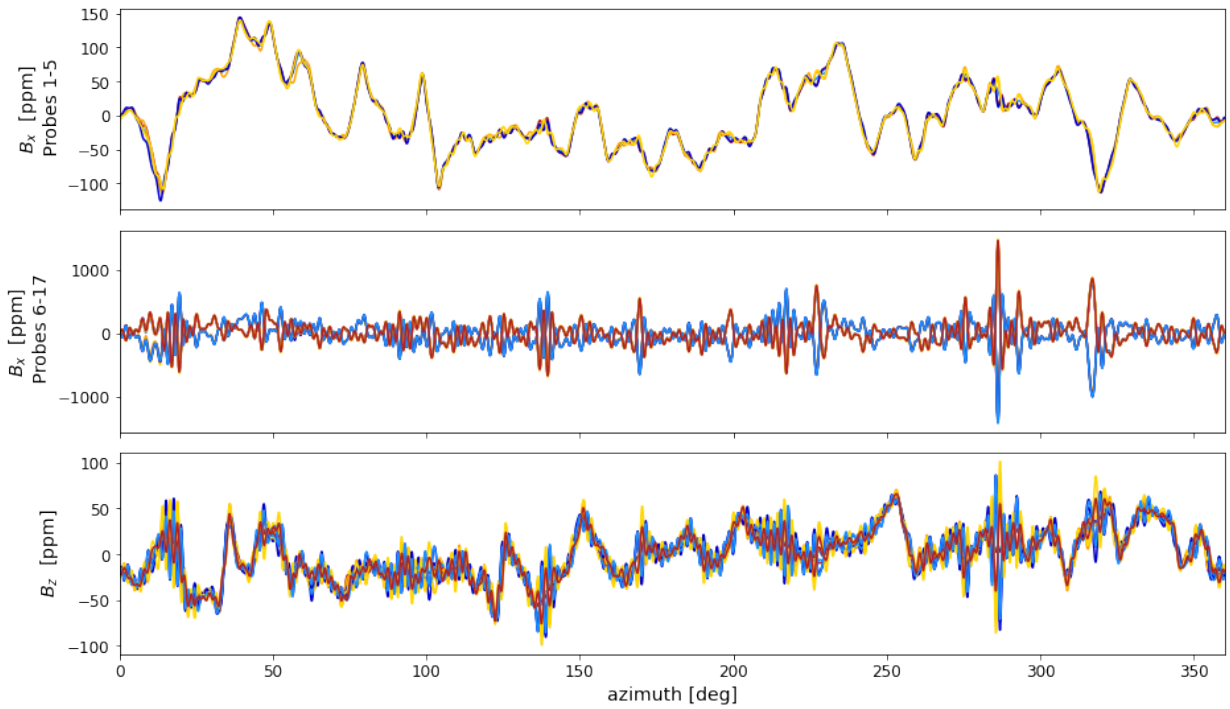
STF4 is based on trolley run 21132 (fig. C.1b).



(a) Radial and longitudinal field input



(b) Simulated field magnitude and difference to vertical field

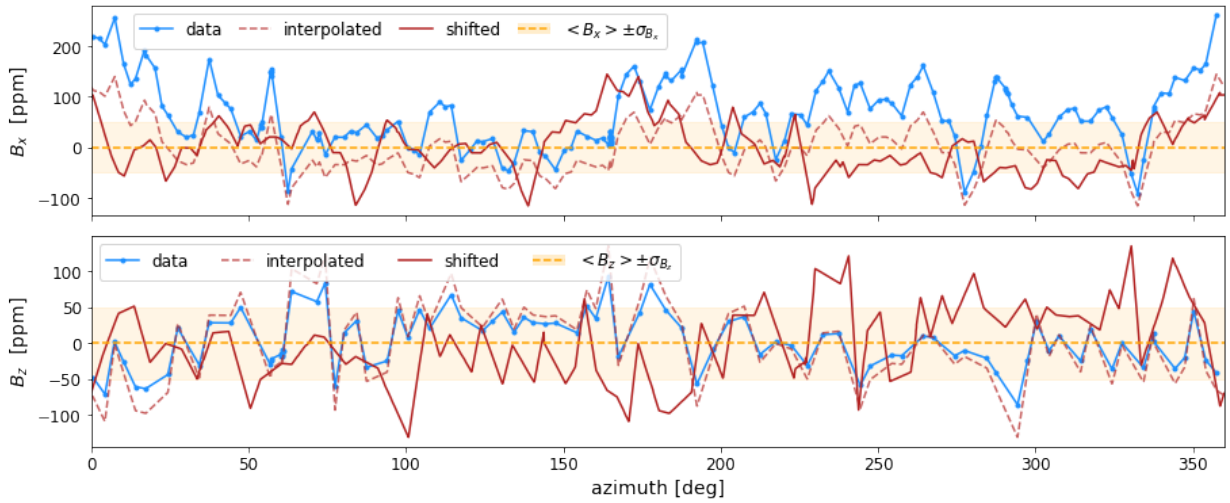


(c) Simulated radial and longitudinal field

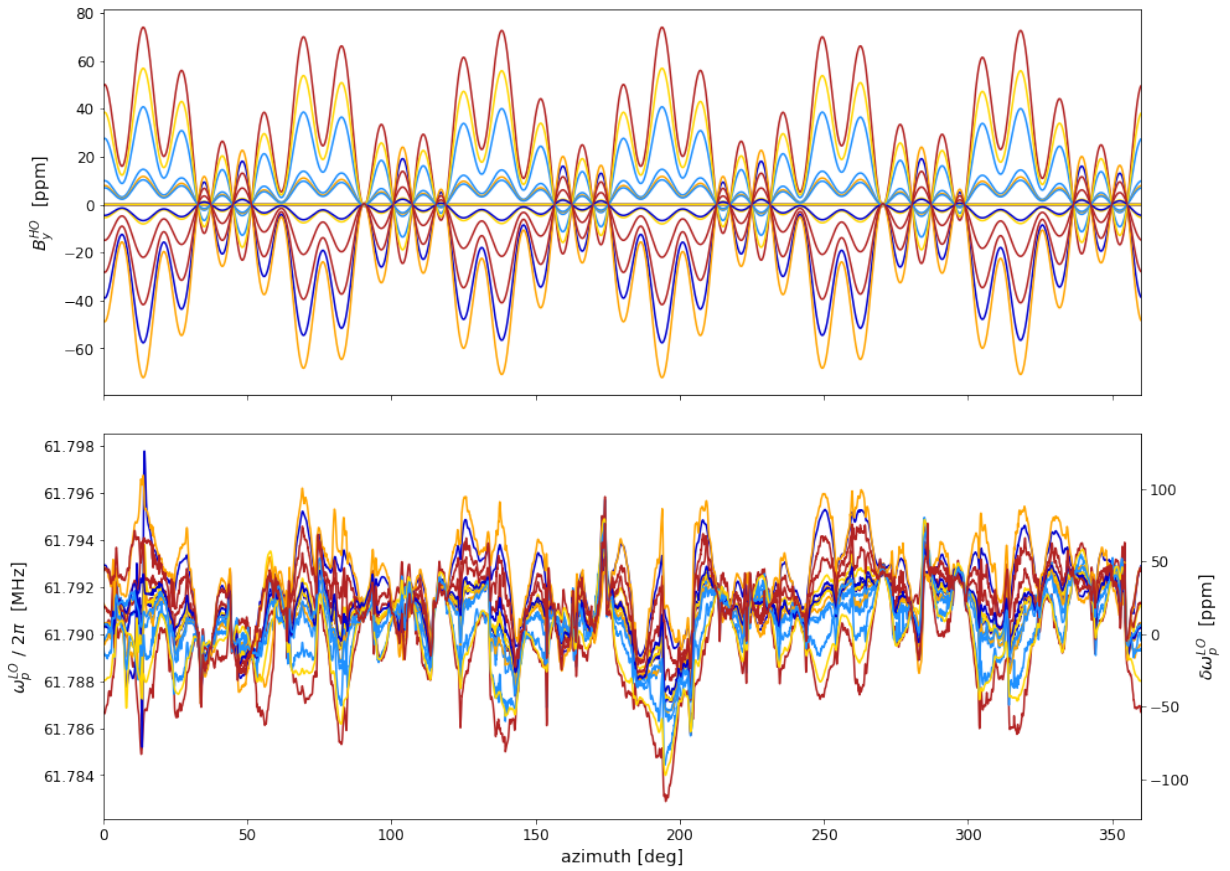
Figure C.6: Graphs of STF4

C.6 STF5

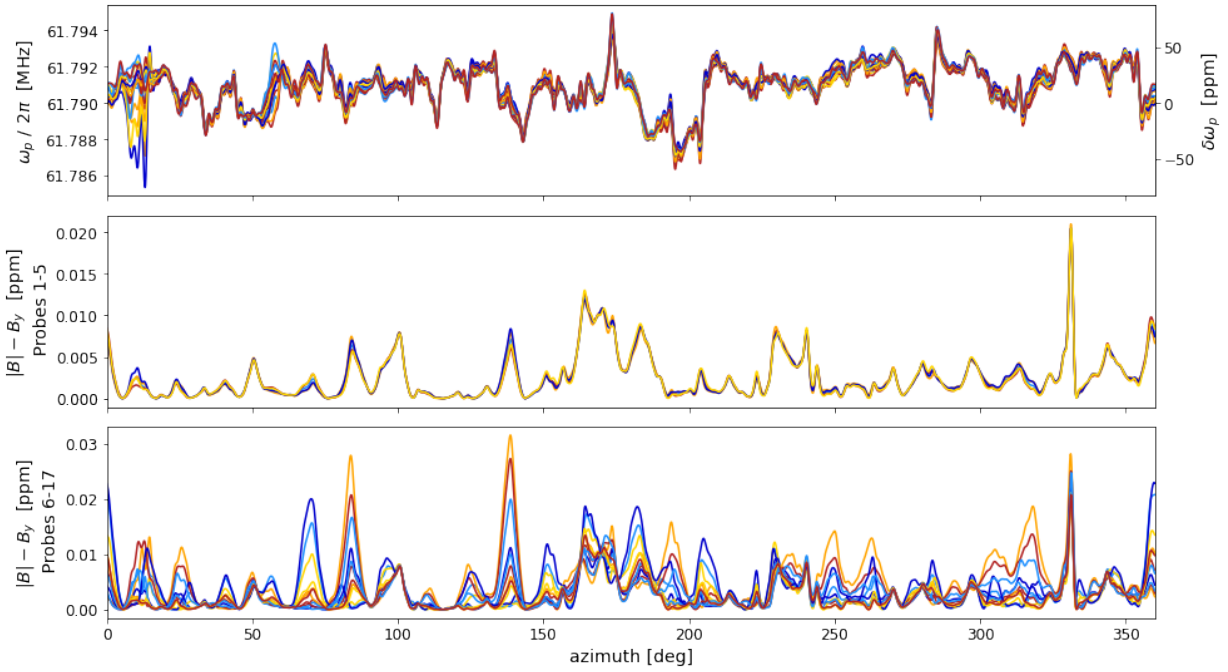
STF5 is based on trolley run 25392 (fig. C.1c).



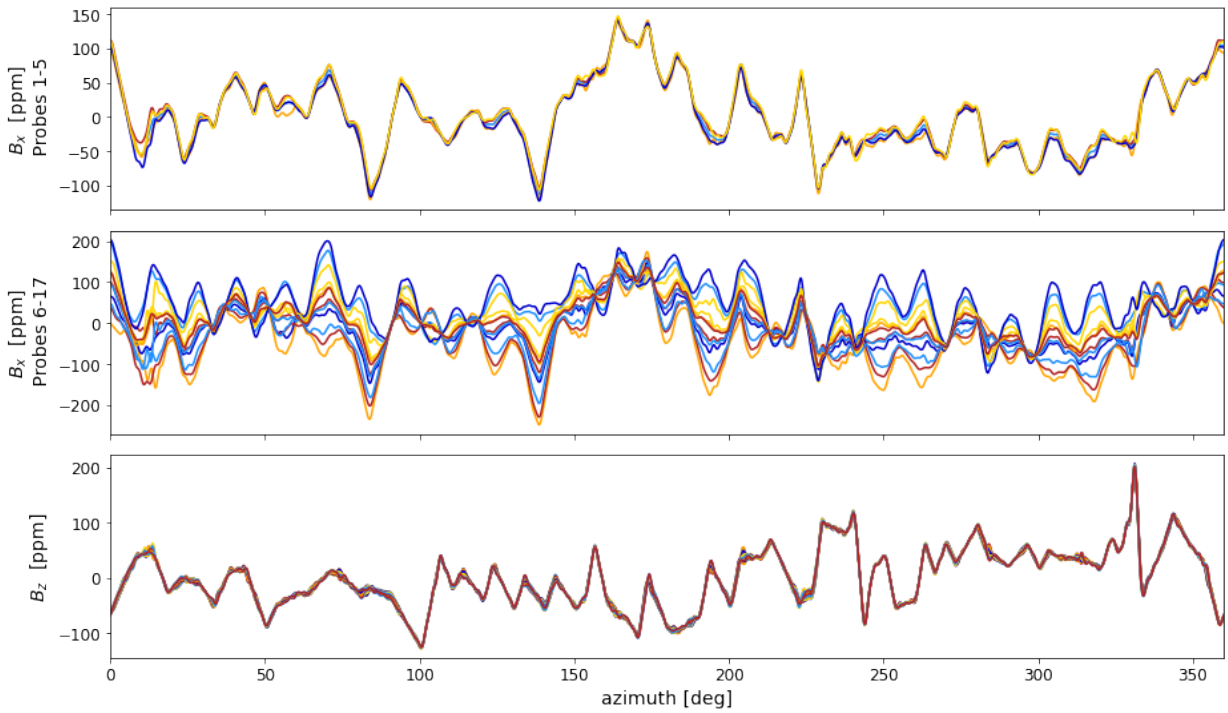
(a) Radial and longitudinal field input



(b) Enforced higher order components (top) and residual field after subtraction of higher orders (bottom)



(c) Simulated field magnitude and difference to vertical field

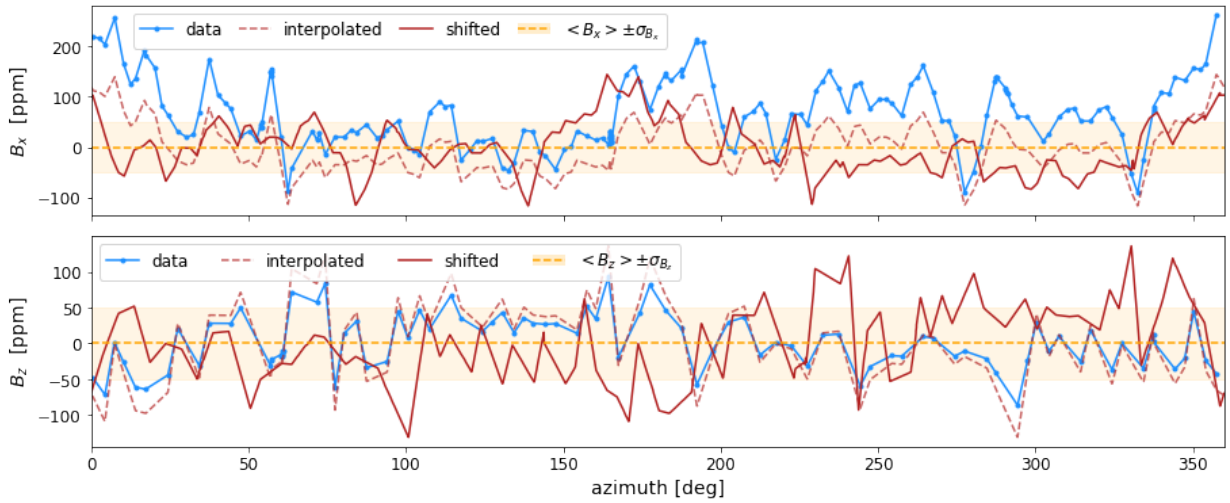


(d) Simulated radial and longitudinal field

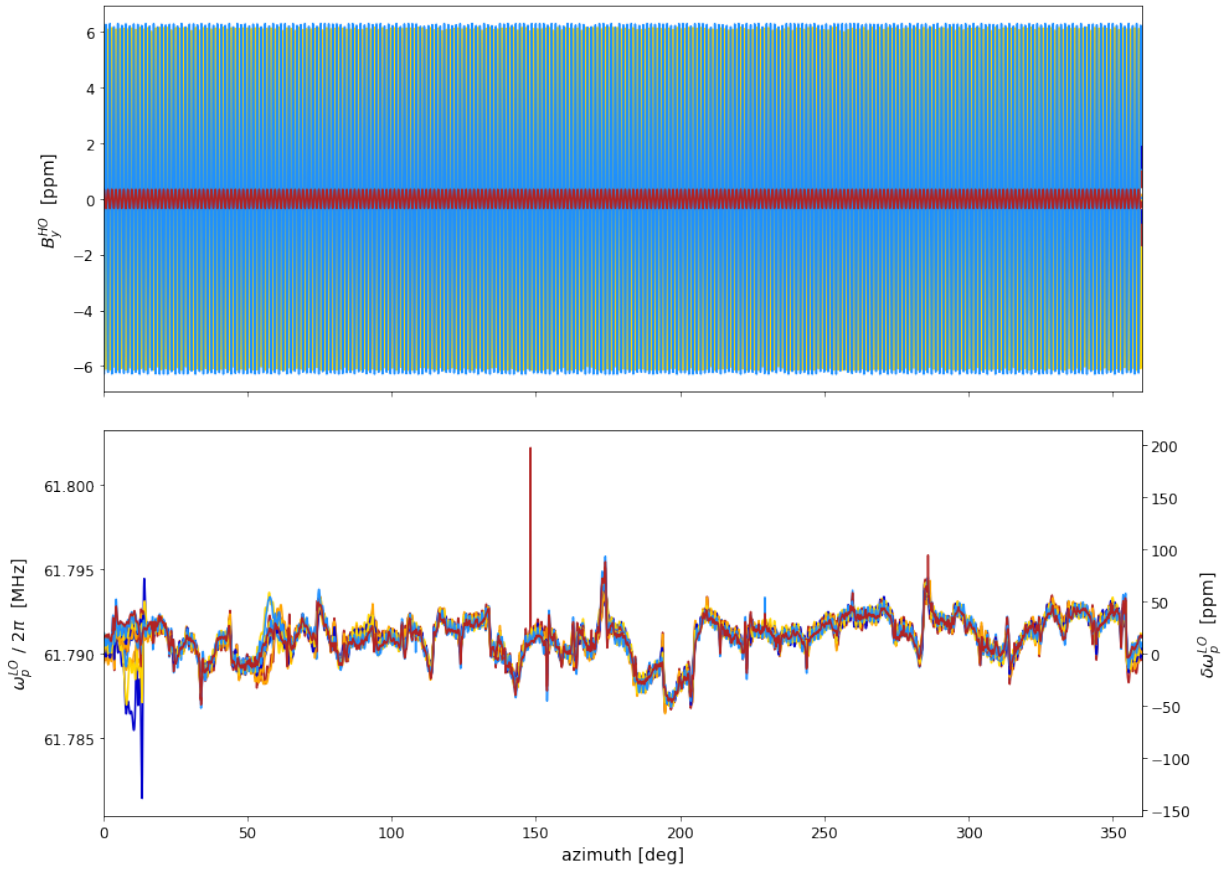
Figure C.7: Graphs of STF5

C.7 STF6

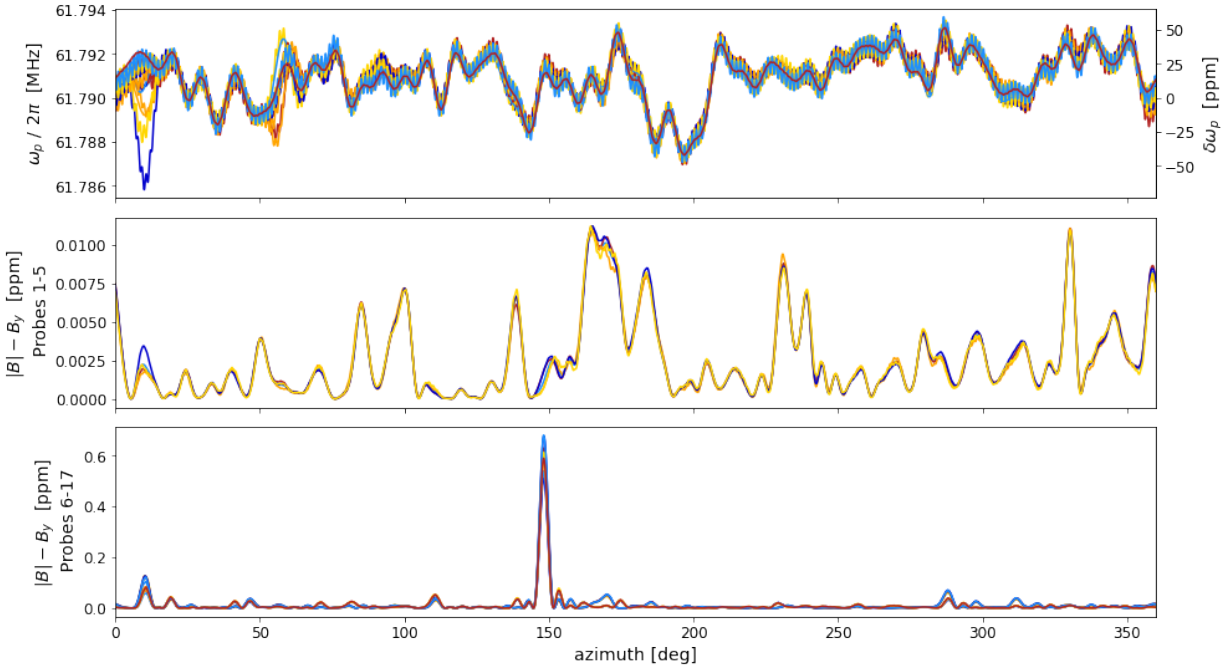
STF5 is based on trolley run 25405 (fig. C.1d).



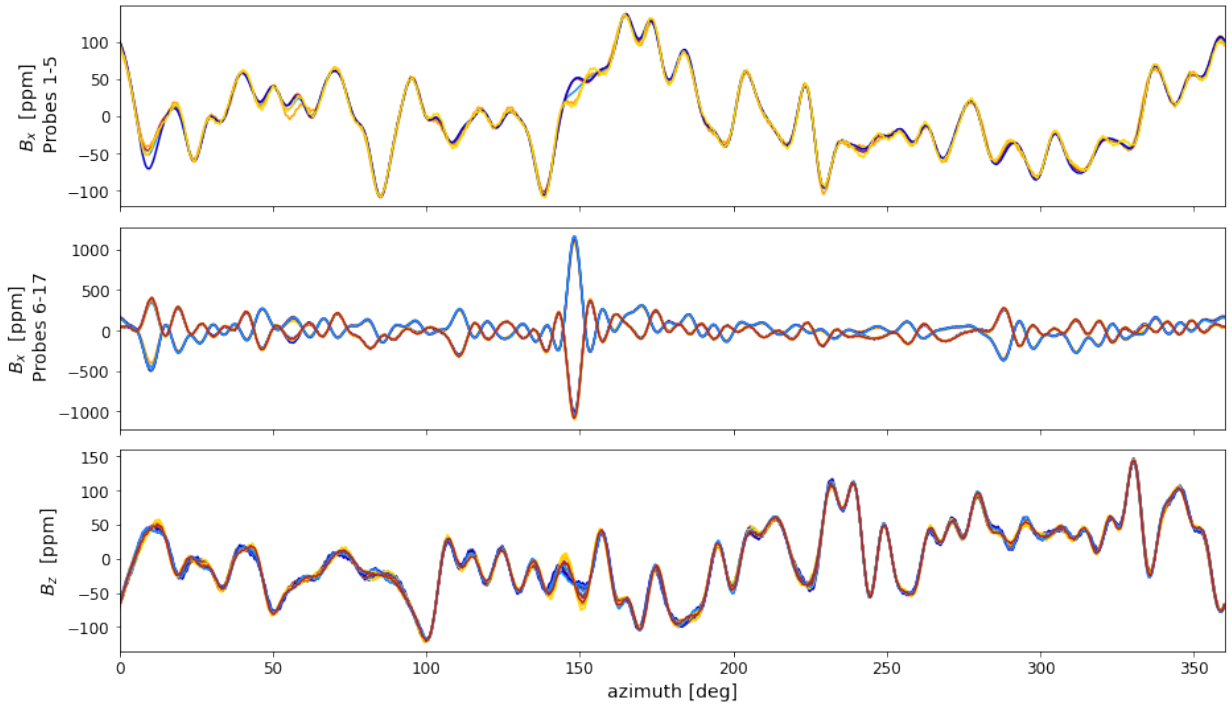
(a) Radial and longitudinal field input



(b) Enforced higher order components (top) and residual field after subtraction of higher orders (bottom)



(c) Simulated field magnitude and difference to vertical field

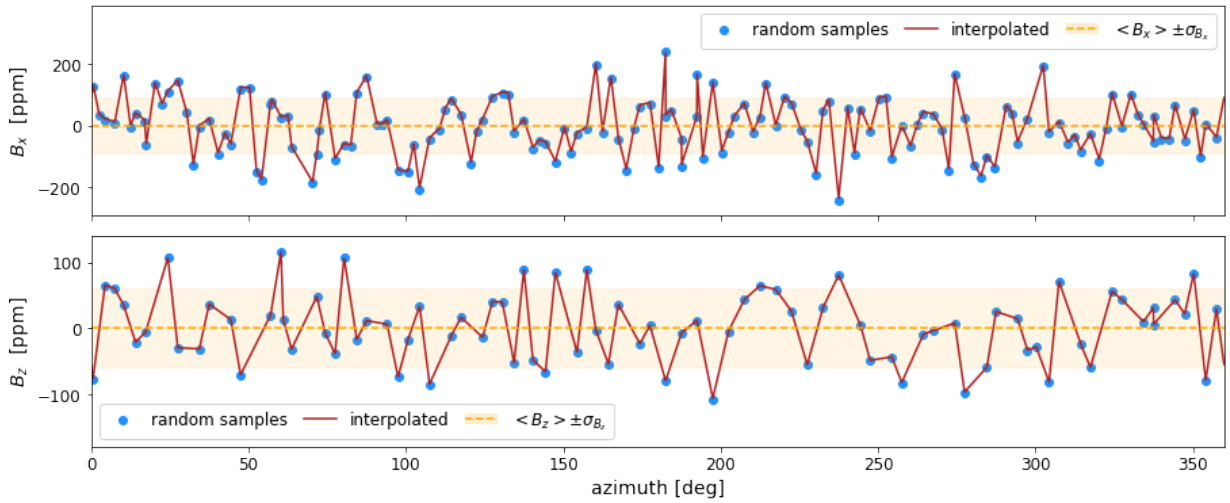


(d) Simulated radial and longitudinal field

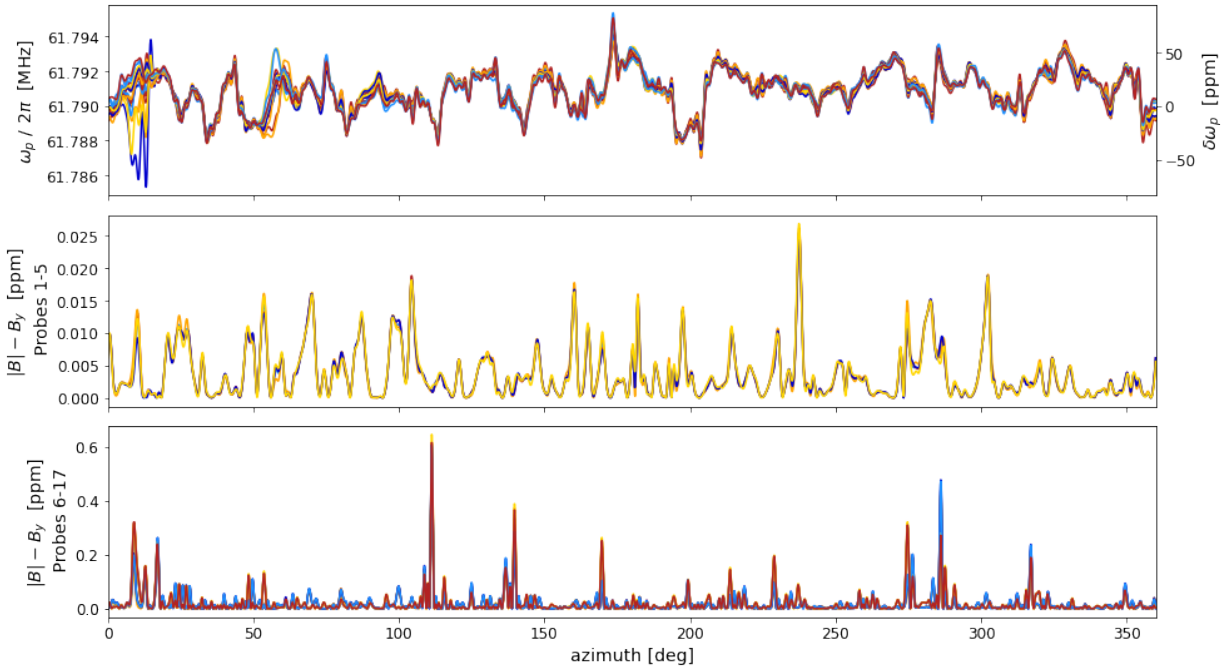
Figure C.8: Graphs of STF6

C.8 STF7

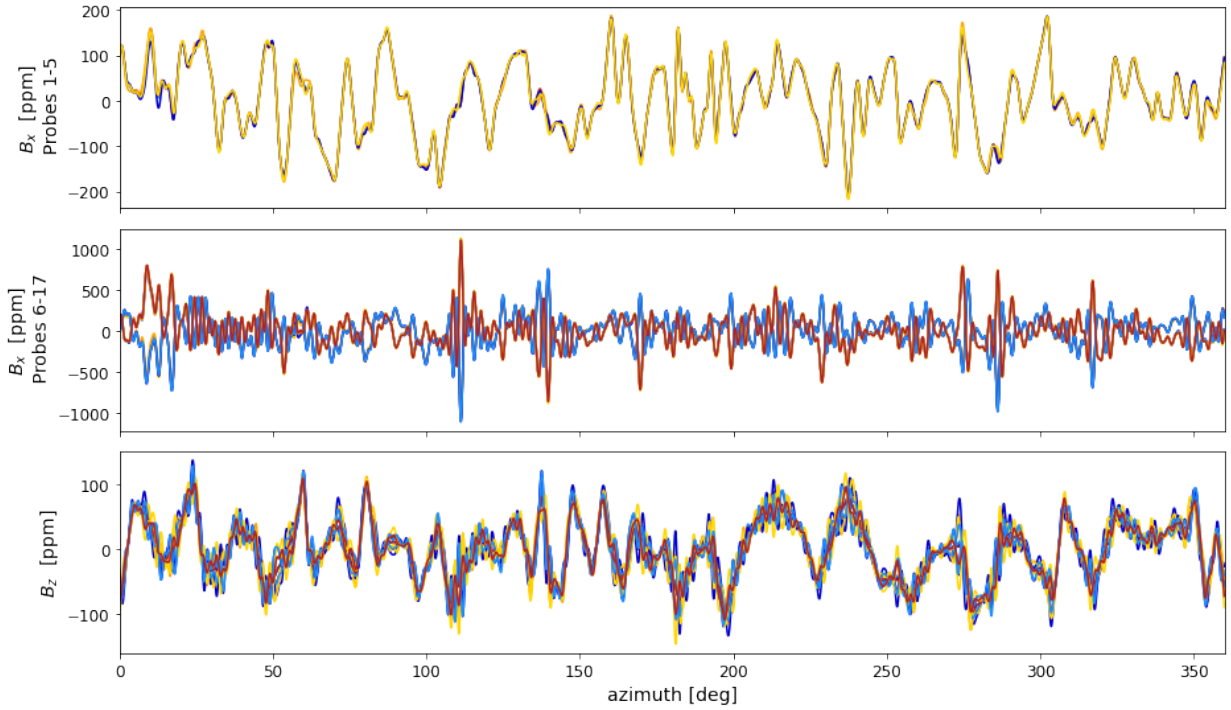
STF7 is based on trolley run 28100 (fig. C.1e).



(a) Radial and longitudinal field input



(b) Simulated field magnitude and difference to vertical field



(c) Simulated radial and longitudinal field

Figure C.9: Graphs of STF7

APPENDIX D

Run Group Uncertainty Estimates

D.1 Statistical uncertainty of sub-target run groups

Table 5.3 in section 5.2.2 lists six trolley pairs in Run-6 which contain less than the desired $600,000$ tracks. To determine whether it is acceptable to form a run group with fewer than the required tracks, we will perform an estimate of the statistical uncertainty of these run groups below.

In chapter 5 we showed that the muon weighted magnetic field of one run group $\tilde{B}_{rg}$ is calculated as the azimuthal mean of the sum of weighted multipole components:

$$\tilde{B}_{rg} = \frac{1}{N_\phi} \sum_{i=1}^{N_\phi} \sum_{j=1}^{13} k_j(\phi_i)\ m_j(\phi_i). \tag{5.13}$$

Here $N_\phi = 72$ is the number of azimuthal 5° bins, m_j is the strength of the j-th multipole moment (the Run-4/5/6 analysis uses 13), and k_j is the corresponding beam multipole projections. The higher order beam projections are functions of transverse position, and their explicit forms are listed in table 5.1. They are determined numerically from the 2D beam histograms for each azimuthal bin. Since the histograms at each azimuthal bin are shifted and scaled combinations of the histograms at the two tracker locations, we will assume that their statistical uncertainties are fully correlated across azimuth, whereas the beam projections are treated as independent from each other ("across j)[1]. Thus the statistical uncertainty on the weighted magnetic field of a run group (due to the beam profiles) can be

[1]Even for the smallest trolley pair dataset, there are still tens of thousands of tracks in each of the station histograms, and aliasing between the different beam multipole projections can be considered non-existent.

written as

$$\sigma^2_{\tilde{B}_{rg}} = \sum_{j=2}^{13} \left(\frac{1}{N_\phi} \sum_{i=1}^{N_\phi} |m_j(\phi_i)|\ \sigma_{k_j}(\phi_i) \right)^2 \tag{D.1}$$

$$\leq \sum_j \left(\frac{1}{N_\phi} \sum_i |m_j(\phi_i)| \max_{\phi_i}(\sigma_{k_j}) \right)^2$$

$$= \sum_j \max_{\phi_i}(\sigma_{k_j})^2 \left(\frac{1}{N_\phi} \sum_i |m_j(\phi_i)| \right)^2 . \tag{D.2}$$

Note that since the dipole projection $k_1 = 1.0$ by definition, its uncertainty $\sigma_{k_1} = 0$ for all angles. To determine the statistical uncertainty of a beam projection k_j, we can express k_j as the mean of a function $f_j(\bar{x}, \bar{y})$, where $\bar{x} = x/r_0$ and $\bar{y} = y/r_0$ are the normalized transverse coordinates (normalized by the storage volume radius $r_0 = 45$mm). Then the statistical uncertainty of k_j in an azimuthal bin with center ϕ_i is just the standard error of the mean (SEM) of the respective function $f_j(\phi_i)$, as shown in table 5.1. The histogram at ϕ_i is the sum of the shifted and scaled station histograms, and therefore the function mean $\langle f_j(\phi_i) \rangle = \langle f_j \rangle_{\phi_i}$ ($\langle\ \rangle_{\phi_i}$ denotes the average at azimuthal bin ϕ_i) is equal to a weighted average of the function means of the (shifted and scaled) station histograms,

$$k_j(\phi_i) = \langle f_j(\bar{x}, \bar{y}) \rangle_{\phi_i} = \frac{N_{12} \langle f_j^{12}(\bar{x}, \bar{y}) \rangle_{\phi_i} + N_{18} \langle f_j^{18}(\bar{x}, \bar{y}) \rangle_{\phi_i}}{N_{12} + N_{18}}, \tag{D.3}$$

where N_{12} and N_{18} are the number of tracks in the station 12 and station 18 histograms, respectively, and the total number of tracks in the run group is $N_{rg} = N_{12} + N_{18}$. For a given ϕ_i and j, the SEM is (omitting ϕ_i and j as well as switching to subscript to indicate station

number):

$$
\begin{aligned}
\sigma_k^2 = \sigma_{\langle f\rangle}^2 &= \frac{1}{N_{rg}}\sigma_f^2 = \frac{1}{N_{rg}^2}\left(\sum_{n=1}^{N_{12}}\Big(f_{12}(\bar{x}_n,\bar{y}_n) - \langle f\rangle\Big)^2 + \sum_{m=1}^{N_{18}}\Big(f_{18}(\bar{x}_m,\bar{y}_m) - \langle f\rangle\Big)^2\right) \\
&= \frac{N_{12}}{N_{rg}^2}\left(\frac{1}{N_{12}}\sum_n f_{12}^2(\bar{x}_n,\bar{y}_n) - 2\langle f\rangle f_{12}(\bar{x}_n,\bar{y}_n) + \langle f\rangle^2\right) \\
&\quad + \frac{N_{18}}{N_{rg}^2}\left(\frac{1}{N_{18}}\sum_m f_{18}^2(\bar{x}_n,\bar{y}_n) - 2\langle f\rangle f_{18}(\bar{x}_n,\bar{y}_n) + \langle f\rangle^2\right) \\
&= \frac{N_{12}}{N_{rg}^2}\Big(\langle f^2\rangle_{12} - 2\langle f\rangle\langle f\rangle_{12} + \langle f\rangle^2\Big) + \frac{N_{18}}{N_{rg}^2}\Big(\langle f^2\rangle_{18} - 2\langle f\rangle\langle f\rangle_{18} + \langle f\rangle^2\Big) \\
&= \frac{N_{12}}{N_{rg}^2}\langle f^2\rangle_{12} + \frac{N_{18}}{N_{rg}^2}\langle f^2\rangle_{18} - \frac{2}{N_{rg}^2}\langle f\rangle(N_{12}\langle f\rangle_{12} + N_{18}\langle f\rangle_{18}) + \frac{1}{N_{rg}}\langle f\rangle^2 \\
&= \frac{N_{12}}{N_{rg}^2}\langle f^2\rangle_{12} + \frac{N_{18}}{N_{rg}^2}\langle f^2\rangle_{18} - \frac{1}{N_{rg}}\langle f\rangle^2 \\
&= \frac{N_{12}}{N_{rg}^2}\langle f^2\rangle_{12} + \frac{N_{18}}{N_{rg}^2}\langle f^2\rangle_{18} - \frac{1}{N_{rg}^3}\Big(N_{12}\langle f\rangle_{12} + N_{18}\langle f\rangle_{18}\Big)^2 \\
&= \frac{N_{12}}{N_{rg}^2}\langle f^2\rangle_{12} - \frac{N_{12}^2}{N_{rg}^3}\langle f\rangle_{12}^2 + \frac{N_{18}}{N_{rg}^2}\langle f^2\rangle_{18} - \frac{N_{18}^2}{N_{rg}^3}\langle f\rangle_{18}^2 - \frac{2N_{12}N_{18}}{N_{rg}^2}\langle f\rangle_{12}\langle f\rangle_{18} \\
&= \frac{N_{12}}{N_{rg}^2}\sigma_{f_{12}}^2 + \frac{N_{12}}{N_{rg}^2}\left(1 - \frac{N_{12}}{N_{rg}^2}\right)\langle f\rangle_{12}^2 + \frac{N_{18}}{N_{rg}^2}\sigma_{f_{18}}^2 + \frac{N_{18}}{N_{rg}^2}\left(1 - \frac{N_{18}^2}{N_{rg}}\right)\langle f\rangle_{18}^2 \\
&\quad - \frac{2N_{12}N_{18}}{N_{rg}^2}\langle f\rangle_{12}\langle f\rangle_{18} \\
&= \frac{N_{12}}{N_{rg}^2}\sigma_{f_{12}}^2 + \frac{N_{18}}{N_{rg}^2}\sigma_{f_{18}}^2 + \frac{N_{12}N_{18}}{N_{rg}^3}\Big(\langle f\rangle_{12}^2 + \langle f\rangle_{18}^2 - 2\langle f\rangle_{12}\langle f\rangle_{18}\Big) \\
&= \frac{N_{12}}{N_{rg}^2}\sigma_{f_{12}}^2 + \frac{N_{18}}{N_{rg}^2}\sigma_{f_{18}}^2 + \frac{N_{12}N_{18}}{N_{rg}^3}\Big(\langle f\rangle_{12} - \langle f\rangle_{18}\Big)^2. \qquad \text{(D.4)}
\end{aligned}
$$

Eq. D.4 gives us an easy way to numerically determine σ_{k_j} for each azimuthal bin, since the extrapolation code already calculates each of the shifted and scaled station histograms. We simply calculate the function f_j^{st} for each histogram bin, and take the mean and standard deviation directly. Once $\sigma_{k_j}(\phi_i)$ has been obtained for each ϕ_i, we take the largest value for our total uncertainty estimate. Table D.1 summarizes $\max_{\phi_i}(\sigma_{k_j})$ for the six trolley pairs in question.

With the azimuthal average of each multipole strength magnitude determined from the VTRs (shown in table D.2), we can find the upper bound for each multipole component's statistical variance from eq. D.2, and the total variance as their sum, both given in table D.3. The statistical uncertainty of each of the sub-target trolley pairs (calculated as the

square root of the variance) is on the order of a few ppb, dominated by the NQ term, and as expected larger than the $<$ 1ppb desired statistical uncertainty. However, their contribution to the total statistical uncertainty of RF-dataset xyRF6 is expected to still be insignificant, given that each of these trolley pairs also has a smaller than average number of CTAGs. By multiplying the trolley pair statistical uncertainty $\sigma_{\tilde{B}_{rg}}$ with the ratio of the trolley pair's CTAGs to all CTAGs in dataset xyRF6[2], we can estimate each trolley pair's contribution to the total statistical uncertainty. Table 5.3 summarizes these results and shows that in all six cases, the contribution to the total is less than 1ppt, much below the order of magnitude that would require any change in our approach.

	$\max_{\phi_i}(\sigma_{k_j})[10^{-3}]$					
Trolley pair	6F6	6G3	6J3	6J6	6J7	6J9
k_2 (NQ)	0.749	0.517	0.635	0.503	0.862	0.629
k_3 (SQ)	0.618	0.424	0.518	0.416	0.711	0.515
k_5 (N6)	0.425	0.293	0.360	0.288	0.494	0.356
k_4 (S6)	0.402	0.279	0.339	0.271	0.460	0.336
k_7 (N8)	0.304	0.212	0.260	0.207	0.354	0.257
k_6 (S8)	0.304	0.209	0.255	0.204	0.347	0.252
k_8 (N10)	0.247	0.171	0.215	0.168	0.284	0.208
k_9 (S10)	0.252	0.172	0.208	0.170	0.288	0.210
k_{10} (N12)	0.225	0.153	0.199	0.152	0.253	0.190
k_{11} (S12)	0.227	0.154	0.190	0.153	0.259	0.188
k_{12} (N14)	0.219	0.145	0.192	0.149	0.240	0.188
k_{13} (S14)	0.226	0.152	0.200	0.151	0.252	0.184

Table D.1: Maximum statistical uncertainty of beam multipole projections (taken across azimuthal bins) in sub-target trolley pair run groups.

[2]This should be considered an upper bound, as it would represent the uncertainty contribution if rungroups were fully correlated in the RF-dataset. Rather, the statistical uncertainties of all rungroups should be added in quadrature, which will reduce each rungroup's contribution even further.

	$\langle\lvert m_j\rvert\rangle_{\phi_i}$ [ppm]					
Trolley pair	6F6	6G3	6J3	6J6	6J7	6J9
m_2 (NQ)	2.235	2.214	2.259	2.386	2.384	2.367
m_3 (SQ)	1.432	1.429	1.355	1.359	1.360	1.356
m_5 (N6)	1.260	1.263	1.137	1.236	1.236	1.254
m_4 (S6)	3.073	3.071	3.081	3.094	3.080	3.102
m_7 (N8)	1.247	1.261	1.223	1.209	1.218	1.224
m_6 (S8)	1.865	1.865	1.860	1.861	1.858	1.859
m_8 (N10)	0.754	0.744	0.742	0.740	0.744	0.748
m_9 (S10)	1.607	1.592	1.593	1.603	1.595	1.612
m_{10} (N12)	0.758	0.765	0.756	0.764	0.756	0.769
m_{11} (S12)	0.190	0.182	0.201	0.172	0.187	0.190
m_{12} (N14)	1.049	1.054	1.057	1.069	1.076	1.061
m_{13} (S14)	0.000	0.000	0.000	0.000	0.000	0.000

Table D.2: Average of multipole strength magnitudes (taken across azimuthal bins) in sub-target trolley pair run groups.

	$\sigma^2_{\tilde{B}_{rg,j}}[10^{-6}\text{ppm}]$					
Trolley pair	6F6	6G3	6J3	6J6	6J7	6J9
$j = 2$ (NQ)	2.804	1.308	2.054	1.440	4.226	2.213
$j = 3$ (SQ)	0.782	0.367	0.492	0.319	0.937	0.487
$j = 5$ (N6)	0.287	0.137	0.167	0.126	0.372	0.199
$j = 4$ (S6)	1.523	0.734	1.090	0.703	2.011	1.085
$j = 7$ (N8)	0.144	0.072	0.101	0.063	0.186	0.099
$j = 6$ (S8)	0.322	0.151	0.226	0.144	0.416	0.220
$j = 8$ (N10)	0.035	0.016	0.026	0.015	0.045	0.024
$j = 9$ (S10)	0.164	0.075	0.110	0.074	0.211	0.114
$j = 10$ (N12)	0.029	0.014	0.023	0.014	0.037	0.021
$j = 11$ (S12)	0.002	0.001	0.001	0.001	0.002	0.001
$j = 12$ (N14)	0.053	0.023	0.041	0.025	0.067	0.040
$j = 13$ (S14)	0.000	0.000	0.000	0.000	0.000	0.000
$\sigma^2_{\tilde{B}_{rg}} = \sum_{j=2}^{13} \sigma^2_{\tilde{B}_{rg,j}}$	6.145	2.898	4.331	2.924	8.510	4.503

Table D.3: Upper bound on statistical uncertainty, per multipole and total, in sub-target trolley pair run groups.

BIBLIOGRAPHY

[1] G. Aad et al. Observation of a new particle in the search for the standard model higgs boson with the atlas detector at the lhc. *Physics Letters B*, 716(1):1–29, 2012.

[2] B. Abi et al. Measurement of the positive muon anomalous magnetic moment to 0.46 ppm. *Phys. Rev. Lett.*, 126:141801, Apr 2021.

[3] Milton Abramowitz and Irene Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. United States Department of Commerce, National Bureau of Standards, 1964.

[4] M.N. Achasov et al. Measurement of the $e^+e^- \to \pi^+\pi^-$ process cross section with the SND detector at the VEPP-2000 collider in the energy region $0.525 < \sqrt{s} < 0.883$ gev. *J. High Energ. Phys.*, 2021:113, Jan 2021.

[5] D. P. Aguillard et al. Measurement of the positive muon anomalous magnetic moment to 0.20 ppm. *Phys. Rev. Lett.*, 131:161802, Oct 2023.

[6] D. P. Aguillard et al. Detailed report on the measurement of the positive muon anomalous magnetic moment to 0.20 ppm. *Phys. Rev. D*, 110:032009, Aug 2024.

[7] D. P. Aguillard et al. Measurement of the positive muon anomalous magnetic moment to 127 ppb. *Phys. Rev. Lett.*, 135:101802, Sep 2025.

[8] David P. Aguillard et al. E989 note 294: Run 2/3 field documentation: Review, jun 2023. Internal Muon g-2 collaboraton document 28082.

[9] T. Albahri et al. Beam dynamics corrections to the run-1 measurement of the muon anomalous magnetic moment at fermilab. *Phys. Rev. Accelerators and Beams*, 24(4), April 2021.

[10] T. Albahri et al. Magnetic-field measurement and analysis for the muon g-2 experiment at fermilab. *Phys. Rev. A*, 103(4), April 2021.

[11] T. Albahri et al. Measurement of the anomalous precession frequency of the muon in the fermilab muon g2 experiment. *Phys. Rev. D*, 103(7), April 2021.

[12] Ricard Alemany, Michel Davier, and Andreas Höcker. Improved determination of the hadronic contribution to the muon $(g-2)$ and to $\alpha(m_z^2)$using new data from hadronic τdecays. *Eur. Phys. J. C*, 2:123–135, March 1998.

[13] C. Alexandrou et al. Lattice calculation of the short and intermediate time-distance hadronic vacuum polarization contributions to the muon magnetic moment using twisted-mass fermions. *Phys. Rev. D*, 107:074506, Apr 2023.

[14] C. Alexandrou et al. Pion transition form factor from twisted-mass lattice qcd and the hadronic light-by-light π^0-pole contribution to the muon $g-2$. *Phys. Rev. D*, 108:094514, Nov 2023.

[15] C. Alexandrou et al. Strange and charm quark contributions to the muon anomalous magnetic moment in lattice qcd with twisted-mass fermions. *Phys. Rev. D*, 111:054502, Mar 2025.

[16] R. Aliberti et al. The anomalous magnetic moment of the muon in the standard model: an update. *Physics Reports*, 1143:1–158, 2025.

[17] Riccardo Aliberti et al. Radiative corrections and Monte Carlo tools for low-energy hadronic cross sections in e^+e^- collisions. *SciPost Phys. Comm. Rep.*, page 9, 2025.

[18] A. Anastasi et al. Combination of KLOE $\sigma\big(e^+e^- \to \pi^+\pi^-\gamma(\gamma)\big)$ measurements and determination of $a_\mu^{\pi^+\pi^-}$ in the energy range $0.10 < s < 0.95$ gev^2. *Journal of High Energy Physics*, (3), March 2018.

[19] E. Anderson et al. *LAPACK Users' Guide*. SIAM, 3rd edition, 1999.

[20] T. Aoyama et al. The anomalous magnetic moment of the muon in the standard model. *Physics Reports*, 887:1–166, December 2020.

[21] Tatsumi Aoyama, Masashi Hayakawa, Toichiro Kinoshita, and Makiko Nio. Complete tenth-order qed contribution to the muon. *Physical Review Letters*, 109(11), September 2012.

[22] Tatsumi Aoyama, Toichiro Kinoshita, and Makiko Nio. Theory of the anomalous magnetic moment of the electron. *Atoms*, 7(1), 2019.

[23] Christopher Aubin, Thomas Blum, Maarten Golterman, and Santiago Peris. Muon anomalous magnetic moment with staggered fermions: Is the lattice spacing small enough? *Phys. Rev. D*, 106:054503, Sep 2022.

[24] Christopher Aubin et al. Light quark vacuum polarization at the physical point and contribution to the muon $g-2$. *Phys. Rev. D*, 101:014503, Jan 2020.

[25] J. Bailey et al. Precision measurement of the anomalous magnetic moment of the muon. *Physics Letters B*, 28(4), 1968.

[26] J. Bailey et al. Tochnoye izmyeryeniye anomalcyrillic small soft signnogo magnktnogo momyenta myuona. *Il Nuovo Cimento A*, 9(4), 1972.

[27] J Bailey et al. New measurement of (g-2) of the muon. *Physics Letters B*, 55(4), 1975.

[28] J. Bailey et al. Final report on the cern muon storage ring including the anomalous magnetic moment and the electric dipole moment of the muon, and a direct test of relativistic time dilation. *Nuclear Physics B*, 150:1–75, 1979.

[29] Alexei Bazavov et al. Hadronic vacuum polarization for the muon $g-2$ from lattice qcd: Complete short and intermediate windows. *Phys. Rev. D*, 111, May 2025.

[30] Alexei Bazavov et al. Hadronic vacuum polarization for the muon $g-2$ from lattice qcd: Long-distance and full light-quark connected contribution. *Phys. Rev. Lett.*, 135, Jul 2025.

[31] G. W. Bennett et al. Final report of the e821 muon anomalous magnetic moment measurement at bnl. *Phys. Rev. D*, 73:072003, Apr 2006.

[32] Genessa Benton et al. Data-driven determination of the light-quark connected component of the intermediate-window contribution to the muon $g-2$, 2023. arXiv:2306.16808.

[33] David Bernecker and Harvey B. Meyer. Vector correlators in lattice qcd: Methods and applications. *The European Physical Journal A*, 47(11), November 2011.

[34] Philip R. Bevington and D. Keith Robinson. *Data Reduction and Error Analysis for the Physical Sciences*. McGraw-Hill, 3rd edition, 2003.

[35] Johan Bijnens, Nils Hermansson-Truedsson, and Antonio Rodríguez-Sánchez. Short-distance constraints for the hlbl contribution to the muon anomalous magnetic moment. *Physics Letters B*, 798:134994, November 2019.

[36] T. Blum et al. Calculation of the hadronic vacuum polarization contribution to the muon anomalous magnetic moment. *Phys. Rev. Lett.*, 121:022003, Jul 2018.

[37] T. Blum et al. Update of euclidean windows of the hadronic vacuum polarization. *Phys. Rev. D*, 108:054507, Sep 2023.

[38] T. Blum et al. Long-distance window of the hadronic vacuum polarization for the muon $g-2$. *Phys. Rev. Lett.*, 134:201901, May 2025.

[39] Thomas Blum et al. Hadronic light-by-light scattering contribution to the muon anomalous magnetic moment from lattice qcd. *Phys. Rev. Lett.*, 124:132002, Apr 2020.

[40] Thomas Blum et al. Hadronic light-by-light contribution to the muon anomaly from lattice qcd with infinite volume qed at physical pion mass. *Phys. Rev. D*, 111(1), 2025.

[41] A. Boccaletti et al. High precision calculation of the hadronic vacuum polarisation contribution to the muon anomaly, 2024.

[42] G.T. Bodwin, H.S. Chung, and J. Repond. Implementation of maxwell's equations in the reconstruction of the magnetic field in the g2 storage ring. *Journal of Instrumentation*, 14(07):P07002, jul 2019.

[43] Sz. Borsanyi et al. Hadronic vacuum polarization contribution to the anomalous magnetic moments of leptons from first principles. *Phys. Rev. Lett.*, 121, Jul 2018.

[44] Sz. Borsanyi et al. Leading hadronic contribution to the muon magnetic moment from lattice qcd. *Nature*, 593(7857), April 2021.

[45] Elia Bottalico, Robert Carey, and James Mott. The electric field correction using the cern-extended method for run-456, May 2025. Internal Muon g-2 collaboraton document 31572.

[46] Mattia Bruno, Taku Izubuchi, Christoph Lehner, and Aaron Meyer. On isospin breaking in τ decays for $(g-2)_\mu$ from lattice qcd, 2018.

[47] M. Cè et al. Window observable for the hadronic vacuum polarization contribution to the muon $g-2$ from lattice qcd. *Phys. Rev. D*, 106:114502, Dec 2022.

[48] En-Hung Chao et al. Hadronic light-by-light contribution to $(g-2)_\mu$ from lattice qcd: a complete calculation. *The European Physical Journal C*, 81(7), July 2021.

[49] En-Hung Chao et al. The charm-quark contribution to light-by-light scattering in the muon $(g-2)$ from lattice qcd. *The European Physical Journal C*, 82(8), August 2022.

[50] Saskia Charity et al. Bloch analysis note, Apr 2022. Internal Muon g-2 collaboraton document 27000.

[51] Saskia Charity et al. Field interpolation report run-456, Mar 2025. Internal Muon g-2 collaboraton document 31245.

[52] G. Charpak et al. Measurement of the anomalous magnetic moment of the muon. *Physical Review Letters*, 6(3), 1961.

[53] G. Charpak et al. A new measurement of the anomalous magnetic moment of the muon. *Physics Letters*, 1(1), 1962.

[54] S. Chatrchyan et al. Observation of a new boson at a mass of 125 gev with the cms experiment at the lhc. *Physics Letters B*, 716(1):30–61, 2012.

[55] G. Colangelo et al. Data-driven evaluations of euclidean windows to scrutinize hadronic vacuum polarization. *Physics Letters B*, 833:137313, 2022.

[56] Gilberto Colangelo et al. Remarks on higher-order hadronic corrections to the muon g2. *Physics Letters B*, 735:90–91, July 2014.

[57] Gilberto Colangelo et al. Longitudinal short-distance constraints for the hadronic light-by-light contribution to $(g2)_\mu$ with large-nc regge models. *Journal of High Energy Physics*, 2020(3), March 2020.

[58] Gilberto Colangelo, Martin Hoferichter, Massimiliano Procura, and Peter Stoffer. Dispersion relation for hadronic light-by-light scattering: two-pion contributions. *Journal of High Energy Physics*, 2017(4), April 2017.

[59] Simon Corrodi. Run-4/5/6 field calibration summary, Apr 2025. Internal Muon g-2 collaboraton document 31744.

[60] Simon Corrodi et al. Purcell review, Jun 2022. Internal Muon g-2 collaboraton document 27375.

[61] Andrzej Czarnecki, William J. Marciano, and Arkady Vainshtein. Refinements in electroweak contributions to the muon anomalous magnetic moment. *Phys. Rev. D*, 67:073006, Apr 2003.

[62] G. D'Agostini. A multidimensional unfolding method based on bayes' theorem. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 362(2):487–498, 1995.

[63] Michel Davier et al. Hadronic vacuum polarization: comparing lattice QCD and data-driven results in systematically improvable ways, 2023. arXiv:2308.04221.

[64] Michel Davier et al. Tensions in $e^+e^- \to \pi^+\pi^-(\gamma)$ measurements: the new landscape of data-driven hadronic vacuum polarization predictions for the muon $g-2$. *The European Physical Journal C*, 84(7), July 2024.

[65] C. T. H. Davies et al. Hadronic-vacuum-polarization contribution to the muon's anomalous magnetic moment from four-flavor lattice qcd. *Phys. Rev. D*, 101, Feb 2020.

[66] Christine Davies. Muon g-2 - an update from the lattice, Aug 2024. Internal Muon g-2 collaboraton document 30904.

[67] Dalibor Djukanovic et al. The hadronic vacuum polarization contribution to the muon g 2 at long distances. *Journal of High Energy Physics*, 2025:98, Apr 2025.

[68] Philipp Eller et al. A flexible event reconstruction based on machine learning and likelihood principles. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1048:168011, 2023.

[69] Saskia Charity Eva Kraegeloh. Run 4/5/6 muon weighting report, Jan 2025. Internal Muon g-2 collaboraton document 31518.

[70] X. Fan, T. G. Myers, B. A. D. Sukra, and G. Gabrielse. Measurement of the electron magnetic moment. *Phys. Rev. Lett.*, 130:071801, Feb 2023.

[71] F.J.M. Farley et al. The anomalous magnetic moment of the negative muon. *Il Nuovo Cimento A Series 10*, 45(1), 1966.

[72] F.J.M. Farley and Y.K. Semertzidis. The 47 years of muon g2. *Progress in Particle and Nuclear Physics*, 52(1), 2004.

[73] Steven Farrell et al. Novel deep learning methods for track reconstruction, 2018.

[74] Ashley Ferreira et al. Ai meets antimatter: Unveiling antihydrogen annihilations, 2024.

[75] Zoltan Fodor et al. Hadronic light-by-light scattering contribution to the anomalous magnetic moment of the muon at the physical pion mass. *Phys. Rev. D*, 111:114509, Jun 2025.

[76] S. Ganguly. Muon campus at fermilab, 2022. arXiv:2208.02889.

[77] Richard L. Garwin, Leon M. Lederman, and Marcel Weinrich. Observations of the failure of conservation of parity and charge conjugation in meson decays: the magnetic moment of the free muon. *Phys. Rev.*, 105:1415–1417, Feb 1957.

[78] Antoine Gérardin et al. Leading hadronic contribution to $(g-2)_\mu$ from lattice qcd with $N_\mathrm{f} = 2+1$ flavors of O(a) improved wilson quarks. *Phys. Rev. D*, 100:014510, Jul 2019.

[79] Antoine Gérardin et al. Lattice calculation of the π^0, η and $\eta^{'}$ transition form factors and the hadronic light-by-light contribution to the muon $g-2$. *Phys. Rev. D*, 111:054511, Mar 2025.

[80] Antoine Gérardin, Harvey B. Meyer, and Andreas Nyffeler. Lattice calculation of the pion transition form factor with $N_f = 2+1$ wilson quarks. *Phys. Rev. D*, 100:034520, Aug 2019.

[81] D. Giusti et al. Strange and charm hvp contributions to the muon (g 2) including qed corrections with twisted-mass fermions. *Journal of High Energy Physics*, 2017:157, Oct 2017.

[82] D. Giusti et al. Electromagnetic and strong isospin-breaking corrections to the muon $g-2$ from lattice QCD + QED. *Phys. Rev. D*, 99:114502, Jun 2019.

[83] D. Giusti, F. Sanfilippo, and S. Simula. Light-quark contribution to the leading hadronic vacuum polarization term of the muon $g-2$ from twisted-mass fermions. *Phys. Rev. D*, 98:114504, Dec 2018.

[84] Davide Giusti and Silvano Simula. Lepton anomalous magnetic moments in lattice qcd+qed. In *Proceedings of 37th International Symposium on Lattice Field Theory — PoS(LATTICE2019)*, volume 363, page 104, 2019.

[85] C. Gnendiger, D. Stöckinger, and H. Stöckinger-Kim. The electroweak contributions to $(g-2)_\mu$ after the higgs-boson mass measurement. *Phys. Rev. D*, 88:053005, Sep 2013.

[86] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, 4th edition, 2013.

[87] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.

[88] J. Grange et al. Muon (g-2) technical design report, 2018.

[89] Joe Grange. E989 note 57: Shift in ωp due to couplings between longitudinal and transverse magnetic field inhomogeneities (distortion of the closed orbit), Jan 2019. Internal Muon g-2 collaboraton document 2548.

[90] Samuel Grant. *Towards a search for the electric dipole moment of the muon at the Fermilab Muon g 2 experiment*. PhD thesis, University College London, 2022.

[91] David J. Griffiths. *Introduction to Electrodynamics*. Pearson, 2013.

[92] Per Christian Hansen. The truncated SVD as a method for regularization. *BIT Numerical Mathematics*, 27(4):534–553, Dec 1987.

[93] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2nd edition, 2009.

[94] Joeri Hermans, Volodimir Begy, and Gilles Louppe. Likelihood-free MCMC with amortized approximate ratio estimators. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4239–4248. PMLR, Jul 2020.

[95] Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, 2nd edition, 2002.

[96] Martin Hoferichter et al. Dispersion relation for hadronic light-by-light scattering: pion pole. *Journal of High Energy Physics*, 2018(10), October 2018.

[97] Martin Hoferichter and Thomas Teubner. Mixed leptonic and hadronic corrections to the anomalous magnetic moment of the muon. *Phys. Rev. Lett.*, 128:112002, Mar 2022.

[98] Siew Yan Hoh and Tyler Barrett. Combined pitch correction analysis for run 4, 5 and 6, May 2025. Internal Muon g-2 collaboraton document 31575.

[99] F. V. Ignatov et al. Measurement of the $e^+e^- \to \pi^+\pi^-$ cross section from threshold to 1.2 gev with the CMD-3 detector. *Phys. Rev. D*, 109:112002, Jun 2024.

[100] John David Jackson. *Classical Electrodynamics*. Wiley, 3rd edition, 1999.

[101] Savely G. Karshenboim and Vladimir G. Ivanov. The g factor of the proton. *Physics Letters B*, 566(1):27–34, 2003.

[102] Alexander Keshavarzi, Daisuke Nomura, and Thomas Teubner. $g-2$ of charged leptons, $\alpha(M_Z^2)$, and the hyperfine splitting of muonium. *Phys. Rev. D*, 101:014029, Jan 2020.

[103] On Kim et al. Reduction of coherent betatron oscillations in a muon g 2 storage ring experiment using rf fields. *New Journal of Physics*, 22(6), jun 2020.

[104] Nicholas Brennan Kinnaird. *Muon spin precession frequency extraction and decay positron track fitting in Run 1 of the Fermilab Muon g-2 experiment*. PhD thesis, Boston University, 2020.

[105] Granino A. Korn and Theresa M. Korn. *Mathematical Handbook for Scientists and Engineers: Definitions, Theorems, and Formulas for Reference and Review*. McGraw-Hill, 1961.

[106] Eva Kraegeloh. Muon weighting inputs for run 4/5/6, Apr 2024. Internal Muon g-2 collaboraton document 30261.

[107] Simon Kuberski. Muon g-2: Lattice calculations of the hadronic vacuum polarization. In *Proceedings of The 40th International Symposium on Lattice Field Theory — PoS(LATTICE2023)*, volume 453, 2024.

[108] Simon Kuberski et al. Hadronic vacuum polarization in the muon g 2: the short-distance contribution from lattice qcd. *Journal of High Energy Physics*, 2024:172, Mar 2024.

[109] Alexander Kurz, Tao Liu, Peter Marquard, and Matthias Steinhauser. Hadronic contribution to the muon anomalous magnetic moment to next-to-next-to-leading order. *Physics Letters B*, 734:144–147, June 2014.

[110] Kevin Labe. E989 note 188: Studying muon losses with recon east, Dec 2019. Internal Muon g-2 collaboraton document 17695.

[111] J. P. Lees et al. Precise measurement of the $e^+e^- \to \pi^+\pi^-(\gamma)$ cross section with the initial-state radiation method at babar. *Phys. Rev. D*, 86:032013, Aug 2012.

[112] J. P. Lees et al. Measurement of additional radiation in the initial-state-radiation processes $e^+e^- \to \mu^+\mu^-\gamma$ and $e^+e^- \to \pi^+\pi^-\gamma$ at babar. *Phys. Rev. D*, 108:L111103, Dec 2023.

[113] Christoph Lehner and Aaron S. Meyer. Consistency of hadronic vacuum polarization between lattice qcd and the r ratio. *Phys. Rev. D*, 101:074515, Apr 2020.

[114] Tian Lin et al. Lattice qcd calculation of the π^0-pole contribution to the hadronic light-by-light scattering in the anomalous magnetic moment of the muon. *Rept. Prog. Phys.*, 88(8), 2025.

[115] Gabriel López Castro, Alejandro Miranda, and Pablo Roig. Isospin breaking corrections in 2π production in tau decays and e^+e^- annihilation: Consequences for the muon $g-2$ and conserved vector current tests. *Phys. Rev. D*, 111:073004, Apr 2025.

[116] Brynn MacCoy. *Beam Dynamics Challenges in the Muon g 2 Experiment*. PhD thesis, University of Washington, 2023.

[117] Brynn MacCoy, Saskia Charity, and Simon Corrodi. E989 note 314: Run 2/3 muon weighted field documentation, Apr 2023. Internal Muon g-2 collaboraton document 28482.

[118] Brynn MacCoy et al. E989 note 281: Muon b-field weighting report for bloch team, Apr 2021. Internal Muon g-2 collaboraton document 24447.

[119] Pere Masjuan and Pablo Sanchez-Puertas. Pseudoscalar-pole contribution to the $(g_\mu - 2)$: A rational approach. *Phys. Rev. D*, 95:054026, Mar 2017.

[120] Kirill Melnikov and Arkady Vainshtein. Hadronic light-by-light scattering contribution to the muon anomalous magnetic moment reexamined. *Phys. Rev. D*, 70:113006, Dec 2004.

[121] Peter Mohr, David Newell, Barry Taylor, and Eite Tiesinga. Codata recommended values of the fundamental physical constants: 2022, 2024. arXiv:2409.03787.

[122] P. Moon and D.E. Spencer. *Field Theory Handbook: Including Coordinate Systems, Differential Equations, and their Solutions*. Springer Verlag, 2nd, corr. 3rd print edition, 1988.

[123] Léo Morel, Zhibin Yao, Pierre Cladé, and Saïda Guellati-Khélifa. Determination of the fine-structure constant with an accuracy of 81 parts per trillion. *Nature*, 588:61–65, December 2020.

[124] M. Della Morte et al. The hadronic vacuum polarization contribution to the muon $g - 2$ from lattice qcd. *Journal of high energy physics*, October 2017.

[125] S. Navas et al. Review of particle physics. *Phys. Rev. D*, 110:030001, Aug 2024.

[126] Rachel Osofsky. *Magnetic Field Determination for Run 1 of the Fermilab Muon g 2 Experiment*. PhD thesis, University of Washington, 2019.

[127] Richard H. Parker et al. Measurement of the fine-structure constant as a test of the standard model. *Science*, 360(6385):191–195, April 2018.

[128] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Addison-Wesley, 1995.

[129] William D Phillips, William E Cooke, and Daniel Kleppner. Magnetic moment of the proton in h2o in bohr magnetons. *Metrologia*, 13(4):179, oct 1977.

[130] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 3rd edition, 2007.

[131] J. J. Sakurai and Jim Napolitano. *Modern Quantum Mechanics*. Addison Wesley, 2nd edition, 1994.

[132] Julian Schwinger. On quantum-electrodynamics and the magnetic moment of the electron. *Phys. Rev.*, 73:416–417, Feb 1948.

[133] Brynn MacCoy Simon Corrodi, Tim Chupp. Field: 9 vs 12 term multipole fits and run-2/3 residuals, Mar 2023. Internal Muon g-2 collaboraton document 28572.

[134] Sebastian Spiegel and Christoph Lehner. High-precision continuum limit study of the hvp short-distance window. *Phys. Rev. D*, 111:114517, Jun 2025.

[135] G. W. Stewart and Ji guang Sun. *Matrix Perturbation Theory*. Academic Press, 1990.

[136] George Sweetmore. *Measurement of the muon spin precession frequency using the straw tracking detectors at the Fermilab Muon g-2 experiment*. PhD thesis, University of Manchester, 2023.

[137] M. J. Syphers. E989 note 152: Closed orbit distortions from quad alignment, Aug 2018. Internal Muon g-2 collaboraton document 13351.

[138] D. A. Tarazona, M. Berz, and K. Makino. Muon loss rates from betatron resonances at the muon g 2 storage ring at fermilab. *International Journal of Modern Physics A*, 34(36), 2019.

[139] David Tarazona. *Beam Dynamics Characterization and Uncertainties in the Muon g-2 Experiment at Fermilab*. PhD thesis, Michigan State University, 2021.

[140] Alexander Tewsley-Booth. *The Precision Magnetic Field Analysis for the Fermilab Muon g-2 Experiment*. PhD thesis, University of Michigan, 2019.

[141] Tammy Walton. Muon g-2 track finding reconstruction, Sep 2019. Internal Muon g-2 collaboraton document 19303.

[142] Gen Wang, Terrence Draper, Keh-Fei Liu, and Yi-Bo Yang. Muon g-2 with overlap valence fermions. *Phys. Rev. D*, 107:034513, Feb 2023.

[143] Kenneth G. Wilson. Confinement of quarks. *Phys. Rev. D*, 10, Oct 1974.